

# Image is All You Need to Empower Large-scale Diffusion Models for In-Domain Generation

Pu Cao<sup>†</sup> Feng Zhou<sup>†</sup> Lu Yang<sup>\*</sup> Tianrui Huang Qing Song  
 Beijing University of Posts and Telecommunications  
 {caopu, zhoufeng, soeaver, huangtianrui, priv}@bupt.edu.cn

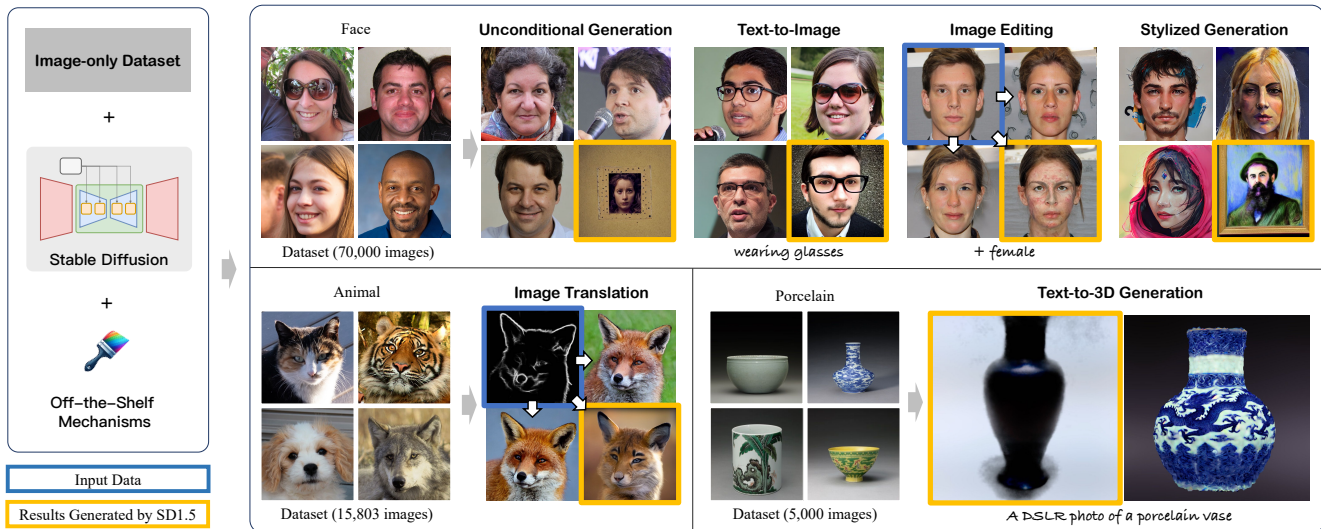


Figure 1. **Illustration of In-domain Generation.** In this work, we empower large-scale pre-trained diffusion models using only image data to perform varied generation tasks within each domain with high fidelity and controllability. We mark the input data in blue and the results generated with the original Stable Diffusion v1.5 model in orange.

## Abstract

*In-domain generation aims to perform a variety of tasks within a specific domain, such as unconditional generation, text-to-image, image editing, 3D generation, and more. Early research typically required training specialized generators for each unique task and domain, often relying on fully-labeled data. Motivated by the powerful generative capabilities and broad applications of diffusion models, we are driven to explore leveraging label-free data to empower these models for in-domain generation. Fine-tuning a pre-trained generative model on domain data is an intuitive but challenging way and often requires complex manual hyper-parameter adjustments since the limited diversity of the training data can easily disrupt the model’s origi-*

*nal generative capabilities. To address this challenge, we propose a guidance-decoupled prior preservation mechanism to achieve high generative quality and controllability by image-only data, inspired by preserving the pre-trained model from a denoising guidance perspective. We decouple domain-related guidance from the conditional guidance used in classifier-free guidance mechanisms to preserve open-world control guidance and unconditional guidance from the pre-trained model. We further propose an efficient domain knowledge learning technique to train an additional text-free UNet copy to predict domain guidance. Besides, we theoretically illustrate a multi-guidance in-domain generation pipeline for a variety of generative tasks, leveraging multiple guidances from distinct diffusion models and conditions. Extensive experiments demonstrate the superiority of our method in domain-specific synthesis and its compatibility with various diffusion-based control methods and applications.*

<sup>†</sup> The first two authors contributed to this paper equally.

<sup>\*</sup> Corresponding author: Lu Yang.

Code is available at <https://github.com/PRIV-Creation/In-domain-Generation-Diffusion>.



Figure 2. **Challenge of fine-tuning diffusion models on domain data.** We show the fine-tuning process of Stable Diffusion v1.5 with a facial image dataset (*i.e.*, FFHQ[20]). The open-world controllability of pre-trained diffusion models is gradually decreased during the fine-tuning process, with domain fidelity improved.

## 1. Introduction

In-domain generation aims to accomplish various tasks within a specific domain, often requiring tailored approaches and labeled data. For instance, pix2pix [19] was developed for image-to-image translation (*e.g.*, converting segmentation or edge images to realistic images), relying on paired images for training. Similarly, InterfaceGAN[45] enables attribute-based image editing but requires attribute-labeled images to train editing vectors. Such traditional methods[6, 12, 19, 45] depend heavily on specific generative tasks and data domains, which limits their generalization capability.

Thanks to billion-scale image-text datasets [2, 37, 43, 44] and advances in large-scale natural language models [9, 35, 36], recent research on diffusion models has pushed the boundaries of AI-generated content, greatly enhancing both the diversity and controllability of generated outputs [39]. With various control mechanisms and approaches (*e.g.*, ControlNet[61], SDEdit[29]), diffusion models can also perform diverse generative tasks beyond text-to-image synthesis, alleviating the data demands of traditional generative methods. However, the generated results often fail to fully align with specific data domains, as shown in orange in Fig. 1.

While fine-tuning a pre-trained model on domain-specific data is an intuitive way to inject domain knowledge, it is challenging to balance fidelity with controllability, as illustrated in Fig. 2. For example, fine-tuning can improve the quality of face generation but may gradually reduce the effect of control prompts like `wearing hat`. In this work, we first identify this phenomenon caused by conditional guidance and unconditional guidance catastrophic forgetting during fine-tuning. These two guidances are integrated by classifier-free guidance for improved generative capability, which is employed in most text-to-image diffusion models. Then, we propose a guidance-decoupled prior preservation method to address these challenges. We decouple the conditional guidance in CFG into two parts: domain guidance, which is optimized in fine-tuning to enhance fi-

delity, and control guidance, which is preserved for controllability. With introducing multiple guidances in the synthesis process, the control guidance and unconditional guidance, which we want to preserve during personalization, are directly predicted by diffusion priors, while the domain guidance is predicted by an independent diffusion model. Additionally, we introduce an efficient domain knowledge learning mechanism, designing a null-text Diffusion Model to learn domain knowledge in a concise way.

We conduct extensive experiments to demonstrate the effects of our proposed task and validate the superiority of our approach. To summarize, our contributions are as follows:

- We present the task of aligning large-scale diffusion models with specific domains using only image data to perform a variety of generative tasks.
- We propose a guidance-decoupled prior preservation mechanism to address guidance forgetting during fine-tuning, which decouples domain guidance to learn domain knowledge with other guidance preserved. Besides, we propose an efficient learning mechanism to learn domain knowledge.
- We present the pipeline of our proposed method for in-domain generation tasks and conduct experiments and comparisons across multiple domains and tasks to demonstrate the effectiveness and advancement of our approach.

## 2. Related Works

**Text-to-image Generation.** As the scale of image-text data grows [2, 37, 43, 44], large-scale diffusion models [10, 17, 31, 39] have demonstrated astonishing generative capabilities, able to produce images based on open-world prompts. Some models, such as Imagen [42] and Stable Diffusion [39], have been explored as diffusion priors, applied in downstream tasks like controlled generation [23, 30, 49, 61] and image editing [8, 15, 21]. For instance, ControlNet [61] proposes using an additional network to receive more spatial controls to guide generation. Furthermore, diffusion models have demonstrated powerful capabilities in many tasks, such as face parsing [54–59], style transfer [24, 51, 60], image manipulation [11, 53] and super-resolution [50, 52].

**Fine-tuning Large-scale Diffusion Models.** Some existing methods study learning subjects into large-scale diffusion priors and generate variant images guided by different prompts [1, 5, 13, 26, 27, 34, 41, 46, 47]. Some methods fine-tune the part of diffusion models with one or a few images of the specific subject [5, 13, 34, 41]. Textual inversion [13] tunes the text embedding of the subject placeholder, while DreamBooth [41] further updates the weight of UNet [40]. Some methods further study training-free [1, 46], continual customization [47], or multi-subjects




Figure 3. **Illustration of Fine-tuning Process.** We demonstrate the guidance catastrophic forgetting during fine-tuning process.

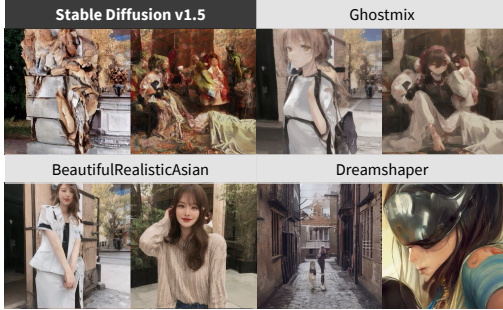


Figure 4. **Unconditional Guidance Drift.** The unconditional generation results of fine-tuned diffusion models reflect the visual pattern of training datasets, which would cause inaccurate noise estimation.

personalization [26, 27] and layout-guided generation [14]. However, these methods concentrate on the consistency of given subjects and training efficiency, tailored for the situation with small-scale image datasets and limited training process.

**Fidelity and Controllability Tradeoff.** In the field of generation, the tradeoff between fidelity and controllability has received widespread attention. In GANs, many inversion methods address this problem through regularization [3, 38], latent code discriminator [48], and invertibility prediction [4, 32]. In the personalization task, thanks to the simple customization process and excellent generalization of diffusion priors, this problem is easily solved by regularization [41], parameter efficient tuning [18], orthogonal fine-tuning [34], and so on. However, a longer training process is required for domain knowledge learning and the existing techniques struggle to solve this problem, highlighting an urgent need for a new approach.

### 3. Method

#### 3.1. Guidance Catastrophic Forgetting

In this section, we delve into the phenomenon of catastrophic forgetting in the training process of diffusion models from the perspective of guidance, allowing us to address this issue at its root.

**Classifier-free Guidance (CFG).** To achieve a better trade-

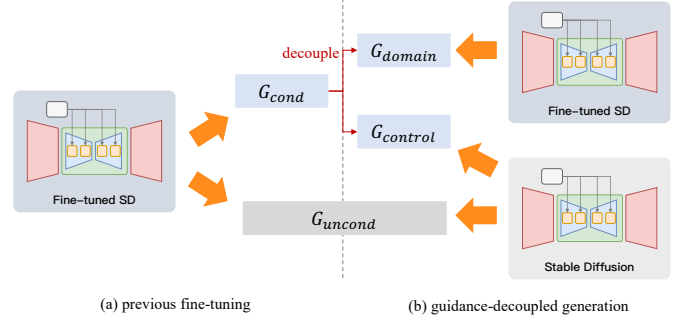


Figure 5. **Conditional Guidance Decoupling.** We compare the guidance estimation between previous customization methods and ours. We decouple the conditional guidance into domain guidance and control guidance while predicting the control guidance and unconditional guidance using the original diffusion model to keep them unchanged.

off between quality and diversity, CFG [16] is utilized to combine the conditional and unconditional score estimates with a hyperparameter  $w$ :

$$\tilde{\epsilon}_{\theta}(z_t, c) = (1 + w)\epsilon_{\theta}(z_t, c) - w\epsilon_{\theta}(z_t) \quad (1)$$

In large-scale diffusion model training, suppose that an open-world training dataset denotes as  $\{z_i, c_i\}$ , where  $z_i$  is the training image and  $c_i$  is the correspondent text, the training objective is to minimize:

$$\mathbb{E} \left( \underbrace{\mathcal{L}(p_{\theta}(z|c), p(z|c))}_{\text{conditional guidance}} + \underbrace{\mathcal{L}(p_{\theta}(z), p(z))}_{\text{unconditional guidance}} \right) \quad (2)$$

**Guidance Catastrophic Forgetting.** During domain data fine-tuning with data  $\{z_i, y_i\}$ , these two guidances inevitably deviate, which we illustrated in Fig. 3. Notably,  $p(z|y)$  represents the probability conditioned on a given domain, which does not imply that text labels of domain data are required. For conditional guidance, the condition  $y$  in fine-tuning phase is too much simpler than the condition in the pre-trained phase  $c$ , where the latter is expected to be control generation. For example, although we leverage only facial images in training, we also desire to generate images controlled by text prompts. For unconditional guidance, we suppose that the pre-trained model could estimate the accurate  $p_{\theta_{sd}}(z_t|\emptyset)$  since it is trained on open-world data. Hence, fine-tuning on domain data would make  $p_{\theta^*}(z_t|\emptyset)$  inaccuracy, illustrated in Fig. 4. The generative results from fine-tuned models reflect the characteristics of the training datasets.

#### 3.2. Guidance-Decoupled Prior Preservation

To address the forgetting issue, we propose a guidance-decoupled prior preservation mechanism. We decouple con-

ditional guidance into domain guidance and control guidance, as illustrated in Fig 5(b). The domain guidance steers the denoising process to generate domain-aligned images and the control guidance takes charge of handling open-world controls in generation. For example, to generate ‘a man wearing glasses’ within the face domain, domain guidance tells how the human face looks, and control guidance ensures the synthetic faces follow ‘man’ and ‘wearing glasses’. After such decoupling, control guidance and unconditional guidance can be predicted by diffusion priors and domain guidance can be learned by an additional model.

**Conditional Guidance Decoupling.** Denoting a sets of conditions as  $C = \{c_1, \dots, c_k\}$  with manually defined intensities  $w = \{w_1, \dots, w_k\}$ , we can use multiple UNets  $\epsilon_\theta$  to predict conditional and unconditional:  $\epsilon_{\theta_0}(z_t) = \nabla_z \log p(z_t)$  and  $\epsilon_{\theta_i}(z_t, c_i) = \nabla_z \log p(z_t|c_i)$ . The reverse process in each timestamp is as follows:

$$\hat{\epsilon}(z_t, C) = \epsilon_{\theta_0}(z_t) + \sum_{i=1}^K w_i(\epsilon_{\theta_i}(z_t, c_i) - \epsilon_{\theta_0}(z_t)) \quad (3)$$

Hence, we can decouple the original conditional guidance into two components to improve domain alignment and open-world control, respectively. As shown in Eq. (6), we can use different diffusion models to predict each guidance, allowing us to fine-tune a UNet copy to learn domain knowledge and utilize the pre-trained model to control the generation.

However, Eq. (6) assumes that all conditions are independent. This assumption leads to the independent maximization of the probability of each condition, which can cause confusion between the conditions. This issue is analyzed further in Appendix A. Our analysis reveals that a simple yet effective approach to mitigate this issue is to use more detailed descriptions (e.g., a photo of a face, wearing glasses) rather than overly simplified conditions (e.g., wearing glasses). This approach helps reduce confusion and better aligns the guidance with the intended domain.

**Unconditional Guidance Rectification.** Through multi-guidance generation, an intuitive solution to address unconditional guidance drift involves rectifying it using prior diffusion models. We can consider the original pre-trained model as an unconditional guidance predictor, effectively mitigating guidance drift after fine-tuning. This rectification operation can also used in other customized diffusion models, including open-sourced models, as shown in Sec 4.5.

### 3.3. Efficient Domain Knowledge Learning

**Null-text Diffusion Model.** The additional fine-tuned diffusion model aims to learn the domain guidance by a do-

main image dataset. To make it concise, we construct a null-text diffusion model to learn domain guidance.

In LDMs, cross-attention is used to fuse text feature and image feature with residual connection, which follows:

$$\mathcal{F}'_{img} = W(\underbrace{\text{Softmax}(\frac{Q_{img}\mathcal{K}_{text}^T}{\sqrt{d}})}_{\text{text feature injection}} \times \mathcal{V}_{text}) + \mathcal{F}_{img} \quad (4)$$

where  $Q_{img}$  is transformed from input image feature  $\mathcal{F}_{img}$ ,  $\mathcal{K}_{text}^T, \mathcal{V}_{text}$  is transformed from text feature,  $d$  is the num of channels, and  $W$  is a linear projection. Since textual information is not used in our fine-tuned model, we can convert the former term into a fixed embedding  $E$ , and we have:

$$\mathcal{F}'_{img} = E + \mathcal{F}_{img} \quad (5)$$

**Initializing  $E$  via Textual Features Optimization.** To better inherent generative capability from diffusion priors, accurate textual features represent a good initializing point in training. Therefore, we draw inspiration from Textual Inversion[13] and optimize the text features  $F^T$ , the output of text encoder, and use them as the initial values for  $E = W_V(F^T)$ . Specifically, we first attain  $F^T$  by feeding a photo of <domain name> into the text encoder and optimize  $F^T$  by denoising loss. This training process converges in just several minutes, providing a good initial value that accelerates the subsequent fine-tuning process.

### 3.4. In-domain Generation with Multi-Guidance

During domain knowledge learning, we fine-tune a domain diffusion model  $\epsilon_{\theta_d}$  to provide domain guidance and employ the Stable Diffusion  $\epsilon_{sd}$  to provide preserved guidance, as illustrated in Fig 5. For control guidance, we can also flexibly adapt existing techniques, including Stable Diffusion, ControlNet, or any open-sourced customized models, to achieve a variety of applications. The reverse process suffices:

$$\underbrace{(1 - w_d - w_c)\epsilon_{sd}(z_t)}_{\text{unconditional guidance}} + \underbrace{w_d\epsilon_{\theta_d}(z_t)}_{\text{domain guidance}} + \underbrace{w_c\epsilon_{\theta_c}(z_t, c)}_{\text{control guidance}}$$

where  $c$  is control signal with model  $\epsilon_{\theta_c}$ . For instance, to generate images conditioned by a canny image in the animal domain, we utilize the diffusion model fine-tuned on an animal dataset as  $\epsilon_{\theta_d}$  and the original diffusion model equipped with canny-to-image ControlNet model as  $\theta_c$ . Experimentally,  $w_d$  presents stability in each generative task, similar to the guidance scale used in widely-used diffusion models (e.g., using  $w_0 = 7$  for Stable Diffusion  $\epsilon_{sd}$ ). And we determine the initial value of  $w_c$  and  $w_d$  by  $w_d = \frac{w_0}{2} \cdot \frac{\|\epsilon_{sd}\|}{\|\epsilon_{\theta_d}\|}$  and  $w_c = \frac{w_0}{2} \cdot \frac{\|\epsilon_{sd}\|}{\|\epsilon_{\theta_c}\|}$  to ensure  $w_0 \cdot \|\epsilon_{sd}\| = w_d \|\epsilon_{\theta_d}\| + w_c \|\epsilon_{\theta_c}\|$ . We can also employ

this generation process in other diffusion-based generative tasks, like image editing and 3D generation.

## 4. Experiments

### 4.1. Experimental Settings

**Datasets.** In this study, we conduct experiments on three domains: face (*i.e.*, FFHQ [20]), animal (*i.e.*, AFHQv2 [7]), and porcelain (*i.e.* a collected dataset), chosen for their diversity and high-resolution (larger than  $512 \times 512$ ). FFHQ is a high-quality human face dataset, containing 70,000 images. Images come from Flickr and are aligned into  $1024 \times 1024$ . We further resize images into  $512 \times 512$  without crop. AFHQv2 is an animal face dataset with a resolution of  $512 \times 512$ , containing 15803 images. The porcelain dataset contains 5,000 images, collected from the internet.

**Settings.** We train the null-text domain diffusion model based on Stable Diffusion v1.5 (SD1.5) at the resolution of 512. For embedding learning phase, 3,200 images are used to attain initial embeddings with batch size of 32 and 100 steps for all domains. For all experiments, we maintain a constant learning rate of  $1e-5$  with the optimizer of AdamW and a batch size of 32. These experiments are conducted on 8 NVIDIA A100 GPUs, with each running until convergence. All controlling models, including ControlNet and open-sourced customized diffusion models, are borrowed from their official releases based on the corresponding version of Stable Diffusion without any training. We employ the DPM-Solver++ scheduler [28] with 100 steps across all models. For more experimental settings, please refer to Appendix C.

### 4.2. Results of In-domain Image Generation

We present the unconditional generation results in Fig. 6 to illustrate the degree of domain alignment. As shown, our results closely match the distribution of the domain datasets, highlighting the effectiveness of our approach in aligning the model outputs with domain data. For conditional generation, we showcase the qualitative results in Fig. 7. We conduct experiments on both text-conditioned and spatial-conditioned generation, utilizing ControlNet models for the latter task. The results clearly demonstrate that our method not only achieves high fidelity with respect to the target domains but also retains the generative controllability seen in pre-trained models, ensuring that the model can be guided by text prompts or spatial conditions as intended. Quantitative results and comparisons can be found in Tab. 1 and Sec. 4.4.

### 4.3. Results of Other In-domain Generation Tasks

Apart from the basic generative tasks, our proposed method also benefits other in-domain generation tasks. In this part,

we further demonstrate the effectiveness of our method in a wider range of tasks.

**Image Generation under Complex Conditions.** As discussed in Sec. 3.2, multiple guidances can be leveraged to control the generation process. In this section, we present the generation results under complex conditions in Fig. 8. Despite the introduction of complex conditions, the generated images maintain high fidelity to the target domain, demonstrating both the robustness and quality of our method. Additionally, our generative framework allows us to leverage additional open-source models for controlling the generation process, such as using DreamShaper to control the style.

**Image Editing.** In Fig. 9, we implement SDEdit[29] on the face domain, where the input image is first encoded into latent space and added noise and then we use a similar generation pipeline presented in Sec. 3.4. Detailed settings can be found in Appendix C.3.1. Our results are more aligned with the distribution of real human face images, demonstrating better visual quality.

**Text-to-3D Generation.** In Fig. 10, we leverage our model on porcelain domain for text-to-3d porcelain generation and compare it with classic text-to-3d model [33] on SD1.5. Specifically, we begin by generating porcelain images based on the provided prompt, which are then converted into 3D models using an off-the-shelf image-to-3D method [25]. Subsequently, we fine-tune the resulting 3D models using our porcelain model, guided by the SDS [33] loss. Lastly, we directly convert the NeRF models into meshes. Detailed settings and additional ablation studies on 3D generation can be found in Appendix C.3.2. Our results demonstrate superiority in color richness, fine detail, and alignment with the dataset.

### 4.4. Comparisons to Fine-tuning Techniques

To validate the superiority of our proposed method, we compare our method with several fine-tuning baselines, including vanilla fine-tuning, parameter-efficient fine-tuning, and some training techniques borrowed from other tasks. For detailed experimental settings, please refer to Appendix C.2.

**Baselines.** To establish comprehensive baselines, we conduct experiments on several training techniques, including vanilla fine-tuning, low-rank fine-tuning [18], textual inversion [13], custom diffusion [22], and OFT mechanism [34]. To train diffusion models without text annotation, we adopt prompt templates, utilizing formats such as a photo of [V] <domain name>. These additional tokens [V] are omitted in compared methods where text embedding is not trained.

**Evaluation Setting.** For unconditional generation, we calculate FID on 50,000 images. For text-guided generation,

	Face	Animal	Porcelain
Domain Dataset			
Ours (based on SD1.5)			
SD1.5			
SDXL1.0			

Figure 6. **Visualization of Domain Alignment.** Our results are more aligned to given domain dataset. For SD1.5 and SDXL 1.0, we generate images by text a photo of <domain name>.




Figure 7. **Results of conditional in-domain image generation.** We demonstrate text-conditioned and spatial-conditioned generation results to verify that our proposed methods can effectively preserve the control capability from the pre-trained model. For spatial-conditioned generation, we leverage the off-the-shelf ControlNet model without any extra training.



Figure 8. **Generative Results under Complex Conditions.** We employ multiple conditions and models to demonstrate the performance of our proposed method in achieving more detailed control over the generation process.

Table 1. **Quantitative results of unconditional generation** We evaluate FID and human preference on two datasets. Our method achieves excellent generation results among baselines. †: reproduced results in  $512 \times 512$  resolution. **Bold**: the best result. Underline: the second best result.

Method	#param.	Unconditional Generation						Conditional Generation			
		Face		Animal		Porcelain		Text-Guided		Spatial-Guided	
		FID↓	Pref.↑	FID↓	Pref.↑	FID↓	Pref.↑	Align.↑	Pref.↑	Align.↓	Pref.↑
Textual Inversion [13]	< 1K	70.62	0.3%	44.75	2.5%	112.76	1.7%	<b>0.97</b>	0.2%	0.22	1.5%
Custom Diffusion [22]	44M	40.98	1.6%	49.98	1.8%	98.77	0.6%	0.74	1.4%	<b>0.20</b>	1.7%
OFT [34]	23M	27.50	2.9%	38.74	1.4%	68.84	1.2%	0.77	3.2%	<u>0.21</u>	<u>9.6%</u>
Fine-tuning	859M	<u>12.37</u>	<u>6.4%</u>	<u>23.92</u>	<u>12.4%</u>	<u>62.76</u>	<u>10.8%</u>	0.24	4.5%	0.22	4.9%
Fine-tuning + LoRA[18]	0.19M	23.76	1.3%	24.19	1.7%	73.46	0.3%	0.76	<u>6.7%</u>	0.22	9.1%
Ours	810M	<b>6.57</b>	<b>87.5%</b>	<b>18.82</b>	<b>80.2%</b>	<b>56.46</b>	<b>85.4%</b>	<u>0.89</u>	<b>84.0%</b>	<b>0.20</b>	<b>73.2%</b>




Figure 9. **Image editing on face domain.** We follow SDEdit[29] to edit facial images equipped with our proposed methods.



Figure 10. **3D generation on porcelain domain.** We use DreamFusion [33] based on SD1.5 for comparison. We present two views of the generated objects, including their mesh edges. \* denotes the prompt ‘a DSLR photo of’.

we generate 1,000 images over 10 attribute-related prompts on the facial domain and employ facial predictors to ascertain whether these attributes are accurately represented in

the generated images. Besides, we sample 200 images from CelebA-HQ dataset to extract canny information as the spatial condition and compute the canny discrepancy. We also report human preference (denoted as ‘Pref’) by collecting the win rates of generative images across all comparisons. Detailed settings are presented in Appendix C.2.

**Qualitative Comparisons.** We conducted a qualitative comparison of our method against previous training techniques, as depicted in Fig 11. Methods like textual inversion, custom diffusion, and LoRA, which adhere to a parameter-efficient form, exhibit noticeably lower fidelity compared to others. Vanilla fine-tuning, in particular, demonstrates a decline in controllability during extended training periods. This results in a diminished capacity for text-guided generation and the emergence of unnatural artifacts, especially when used in conjunction with ControlNet. In contrast, our method excels in both fidelity and control capabilities.

**Quantitative Comparisons.** We conduct the evaluation of unconditional generation and report the results in Tab 1. Our approach demonstrates state-of-the-art generation quality, evidenced by a significantly lower FID compared to baseline methods and the highest human preference rating. For text-guided generation, Textual Inversion shows a higher alignment with the text prompts than our method. However, this is achieved at the cost of image quality, owing to its reliance on a very limited set of learnable parameters. This compromise in quality is clearly evident in the comparisons presented in Fig 11. In contrast, in spatial-guided generation, where semantic features are integrated into the UNet, the alignment degree among different methods is similar. Notably, our results not only align well with the control signals but also achieve the highest ratings in human preference assessments.

#### 4.5. Effects of Unconditional Guidance Rectification

In previous methods, unconditional guidance is predicted by the customized model, leading to inaccurate estimation of probability in sampling. In our framework, the uncon-



Figure 11. **Comparisons to existing training methods.** We evaluate in-domain image generation in unconditional, text-guided, and spatial-guided cases. We fix the random seed to illustrate the controllability for the former two.

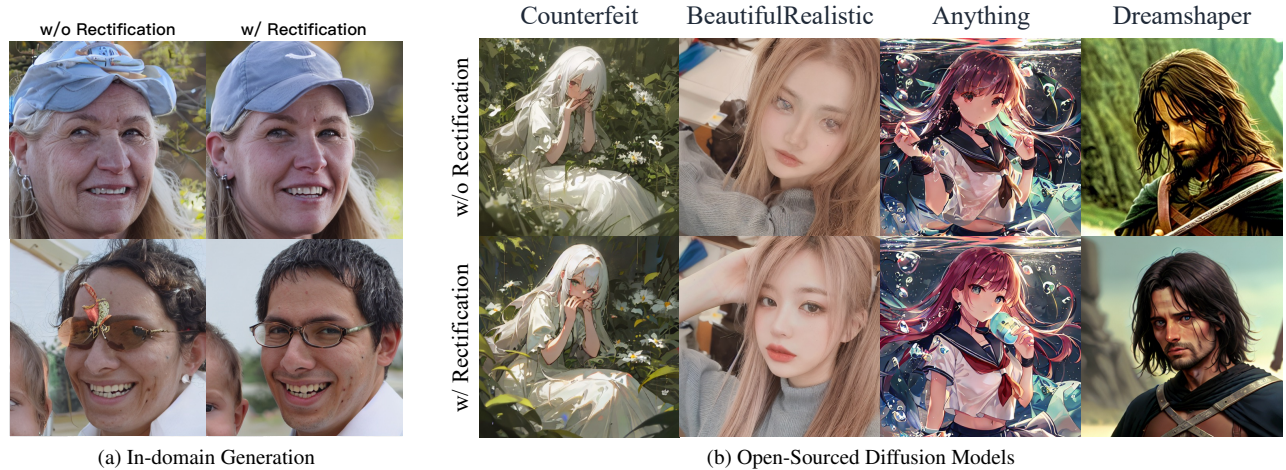


Figure 12. **Effects of Unconditional Guidance Rectification.** (a) Effects in Facial Domain Generation. (b) Unconditional guidance rectification also improves the open-sourced diffusion models.

ditional guidance is rectified by straightforwardly predicting by diffusion priors. We compare generation results in Fig 12a. Our method significantly improves image quality, generating fewer artifacts and attaining photorealistic results. Besides, we utilize four widely-used open-sourced models. The comparative results are presented in Fig 12b. Without rectification, the generation quality shows slight degradation, as evidenced by less detailed eyes and less realistic hair textures. The implementation of guidance decoupling substantially enhances generative performance via the rectification of unconditional guidance.

## 5. Conclusion

In this research, we address the task of enhancing large-scale diffusion models for in-domain generation using only

image data. Our goal is to enable these models to generate images that faithfully represent the given domains, achieving high fidelity, diversity, and controllability. Moreover, our approach is designed to be compatible with various control methods, allowing the model to perform a wide range of generative tasks. To tackle this challenge, we propose a guidance-decoupled prior preservation mechanism that separates conditional guidance into two components: domain guidance for alignment and control guidance for maintaining the models controllability. Additionally, we introduce an efficient learning approach by incorporating a null-text diffusion model, ensuring both simplicity and effectiveness. We then outline the generative pipeline of our method to achieve diverse in-domain generation tasks. We conduct experiments across three domains to demonstrate the effectiveness of our proposed method.

## 6. Acknowledgements

This work was supported by the Young Scientists Fund of NSFC (Grant No. 62406035).

## References

- [1] Moab Arar, Rinon Gal, Yuval Atzmon, Gal Chechik, Daniel Cohen-Or, Ariel Shamir, and Amit H. Bermano. Domain-agnostic tuning-encoder for fast personalization of text-to-image models. pages 1–10, 2023. 2
- [2] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Cocosuff: Thing and stuff classes in context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1209–1218, 2018. 2
- [3] Pu Cao, Lu Yang, Dongxu Liu, Zhiwei Liu, Shan Li, and Qing Song. Lsap: Rethinking inversion fidelity, perception and editability in gan latent space. *arXiv preprint arXiv:2209.12746*, 2022. 3
- [4] Pu Cao, Lu Yang, Dongxu Liu, Xiaoya Yang, Tianrui Huang, and Qing Song. What decreases editing capability? domain-specific hybrid refinement for improved gan inversion. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 4240–4249, 2024. 3
- [5] Hong Chen, Yipeng Zhang, Simin Wu, Xin Wang, Xuguang Duan, Yuwei Zhou, and Wenwu Zhu. Disenbooth: Identity-preserving disentangled tuning for subject-driven text-to-image generation. *arXiv preprint arXiv:2305.03374*, 2023. 2
- [6] Kevin Chen, Christopher B Choy, Manolis Savva, Angel X Chang, Thomas Funkhouser, and Silvio Savarese. Text2shape: Generating shapes from natural language by learning joint embeddings. In *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III 14*, pages 100–116. Springer, 2019. 2
- [7] Yunje Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8188–8197, 2020. 5
- [8] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-based semantic image editing with mask guidance. *arXiv preprint arXiv:2210.11427*, 2022. 2
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186, 2019. 2
- [10] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 2
- [11] Wenkai Dong, Song Xue, Xiaoyue Duan, and Shumin Han. Prompt tuning inversion for text-driven image editing using diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7430–7440, 2023. 2
- [12] Rao Fu, Xiao Zhan, Yiwen Chen, Daniel Ritchie, and Srinath Sridhar. Shapecrafter: A recursive text-conditioned 3d shape generation model. *Advances in Neural Information Processing Systems*, 35:8882–8895, 2022. 2
- [13] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022. 2, 4, 5, 7, 13
- [14] Yuchao Gu, Xintao Wang, Jay Zhangjie Wu, Yujun Shi, Yunpeng Chen, Zihan Fan, Wuyou Xiao, Rui Zhao, Shuning Chang, Weijia Wu, et al. Mix-of-show: Decentralized low-rank adaptation for multi-concept customization of diffusion models. *Advances in Neural Information Processing Systems*, 36:15890–15902, 2023. 3
- [15] 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
- [16] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 3
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2
- [18] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022. 3, 5, 7
- [19] 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
- [20] 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*, pages 4401–4410, 2019. 2, 5
- [21] 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. 2
- [22] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1931–1941, 2023. 5, 7, 13
- [23] Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. Gligen: Open-set grounded text-to-image generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22511–22521, 2023. 2
- [24] Zuoxin Li, Fuqiang Zhou, Lu Yang, Xiaojie Li, and Juan Li. Accelerate neural style transfer with super-resolution. *Multimedia Tools and Applications*, 79:4347–4364, 2020. 2

- [25] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9298–9309, 2023. 5
- [26] Zhiheng Liu, Ruili Feng, Kai Zhu, Yifei Zhang, Kecheng Zheng, Yu Liu, Deli Zhao, Jingren Zhou, and Yang Cao. Cones: Concept neurons in diffusion models for customized generation. *arXiv preprint arXiv:2303.05125*, 2023. 2, 3
- [27] Zhiheng Liu, Yifei Zhang, Yujun Shen, Kecheng Zheng, Kai Zhu, Ruili Feng, Yu Liu, Deli Zhao, Jingren Zhou, and Yang Cao. Customizable image synthesis with multiple subjects. *Advances in neural information processing systems*, 36:57500–57519, 2023. 2, 3
- [28] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*, 2022. 5, 13
- [29] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 2, 5, 7
- [30] 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. 2
- [31] 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. 2
- [32] Gaurav Parmar, Yijun Li, Jingwan Lu, Richard Zhang, Jun-Yan Zhu, and Krishna Kumar Singh. Spatially-adaptive multilayer selection for gan inversion and editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11399–11409, 2022. 3
- [33] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 5, 7
- [34] Zeju Qiu, Weiyang Liu, Haiwen Feng, Yuxuan Xue, Yao Feng, Zhen Liu, Dan Zhang, Adrian Weller, and Bernhard Schölkopf. Controlling text-to-image diffusion by orthogonal finetuning. *Advances in Neural Information Processing Systems*, 36:79320–79362, 2023. 2, 3, 5, 7, 13
- [35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 2
- [36] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020. 2
- [37] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022. 2
- [38] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Transactions on graphics (TOG)*, 42(1):1–13, 2022. 3
- [39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2
- [40] 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. 2
- [41] 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. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22500–22510, 2023. 2, 3, 13
- [42] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022. 2
- [43] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 2
- [44] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in neural information processing systems*, 35:25278–25294, 2022. 2
- [45] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representation learned by gans. *IEEE transactions on pattern analysis and machine intelligence*, 44(4):2004–2018, 2020. 2
- [46] Jing Shi, Wei Xiong, Zhe Lin, and Hyun Joon Jung. Instantbooth: Personalized text-to-image generation without test-time finetuning. pages 8543–8552, 2024. 2
- [47] James Seale Smith, Yen-Chang Hsu, Lingyu Zhang, Ting Hua, Zsolt Kira, Yilin Shen, and Hongxia Jin. Continual diffusion: Continual customization of text-to-image diffusion with c-lora. *arXiv preprint arXiv:2304.06027*, 2023. 2
- [48] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021. 3
- [49] Andrey Voynov, Kfir Aberman, and Daniel Cohen-Or. Sketch-guided text-to-image diffusion models. In *ACM SIG-*

- GRAPH 2023 Conference Proceedings*, pages 1–11, 2023. [2](#)
- [50] Yujin Wang, Lingen Li, Tianfan Xue, and Jinwei Gu. Reconstruct-and-generate diffusion model for detail-preserving image denoising. *arXiv preprint arXiv:2309.10714*, 2023. [2](#)
- [51] Zhizhong Wang, Lei Zhao, and Wei Xing. Stylediffusion: Controllable disentangled style transfer via diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7677–7689, 2023. [2](#)
- [52] Chanyue Wu, Dong Wang, Yunpeng Bai, Hanyu Mao, Ying Li, and Qiang Shen. Hsr-diff: Hyperspectral image super-resolution via conditional diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7083–7093, 2023. [2](#)
- [53] Chen Henry Wu and Fernando De la Torre. A latent space of stochastic diffusion models for zero-shot image editing and guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7378–7387, 2023. [2](#)
- [54] Lu Yang, Qing Song, Zhihui Wang, and Ming Jiang. Parsing r-cnn for instance-level human analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 364–373, 2019. [2](#)
- [55] Lu Yang, Qing Song, Zhihui Wang, Mengjie Hu, Chun Liu, Xueshi Xin, Wenhe Jia, and Songcen Xu. Renovating parsing r-cnn for accurate multiple human parsing. In *European Conference on computer vision*, pages 421–437. Springer, 2020.
- [56] Lu Yang, Qing Song, Xueshi Xin, and Zhiwei Liu. Quality-aware network for face parsing. *arXiv preprint arXiv:2106.07368*, 2021.
- [57] Lu Yang, Zhiwei Liu, Tianfei Zhou, and Qing Song. Part decomposition and refinement network for human parsing. *IEEE/CAA Journal of Automatica Sinica*, 9(6):1111–1114, 2022.
- [58] Lu Yang, Qing Song, Zhihui Wang, Zhiwei Liu, Songcen Xu, and Zhihao Li. Quality-aware network for human parsing. *IEEE Transactions on Multimedia*, 2022.
- [59] Lu Yang, Wenhe Jia, Shan Li, and Qing Song. Deep learning technique for human parsing: A survey and outlook. *International Journal of Computer Vision*, 132(8):3270–3301, 2024. [2](#)
- [60] Serin Yang, Hyunmin Hwang, and Jong Chul Ye. Zero-shot contrastive loss for text-guided diffusion image style transfer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22873–22882, 2023. [2](#)
- [61] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. [2](#)