

# JeDi: Joint-Image Diffusion Models for Finetuning-Free Personalized Text-to-Image Generation

Yu Zeng<sup>1</sup> Vishal M. Patel<sup>1</sup> Haochen Wang<sup>2</sup> Xun Huang<sup>3</sup>  
Ting-Chun Wang<sup>3</sup> Ming-Yu Liu<sup>3</sup> Yogesh Balaji<sup>3</sup>  
<sup>1</sup>Johns Hopkins University <sup>2</sup>TTI-Chicago <sup>3</sup>NVIDIA Research

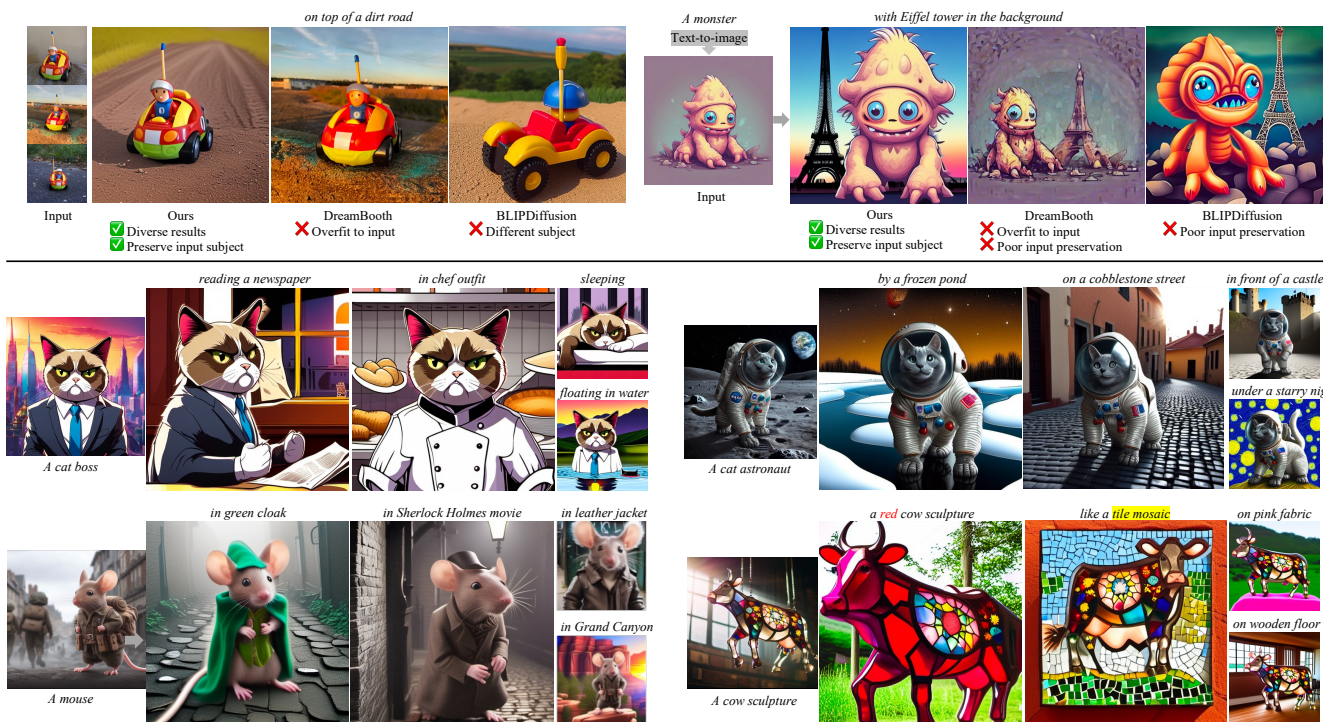


Figure 1. We present Joint-Image Diffusion (JeDi), a finetuning-free image personalization model that can operate on any number of reference images. JeDi is able to preserve the appearance of custom subjects while generating novel variations. As shown in the top row, JeDi does not suffer from the issues of overfitting and lack of diversity exhibited by the prior models. The examples in the bottom two rows demonstrate JeDi’s high-quality results on challenging personalization tasks.

## Abstract

Personalized text-to-image generation models enable users to create images that depict their individual possessions in diverse scenes, finding applications in various domains. To achieve the personalization capability, existing methods rely on finetuning a text-to-image foundation model on a user’s custom dataset, which can be non-trivial for general users, resource-intensive, and time-consuming. Despite attempts to develop finetuning-free methods, their generation quality is much lower compared to their finetuning counterparts. In this paper, we propose Joint-Image

Diffusion (JeDi), an effective technique for learning a finetuning-free personalization model. Our key idea is to learn the joint distribution of multiple related text-image pairs that share a common subject. To facilitate learning, we propose a scalable synthetic dataset generation technique. Once trained, our model enables fast and easy personalization at test time by simply using reference images as input during the sampling process. Our approach does not require any expensive optimization process or additional modules and can faithfully preserve the identity represented by any number of reference images. Experimental results show that our model achieves state-of-the-art gen-

eration quality, both quantitatively and qualitatively, significantly outperforming both the prior finetuning-based and finetuning-free personalization baselines.

## 1. Introduction

The state-of-the-art in text-to-image generation has advanced significantly in the last two years, propelled by the emergence of large-scale diffusion models and paired image-text datasets [2, 3, 6, 22, 25, 26, 28]. Despite their superior capability in generating high-quality images well-aligned to the input text prompts, existing models cannot generate novel images depicting specific custom objects or styles that are only available as few reference images external to the training datasets. To address this important use case, various personalization methods have been developed.

The key challenge of personalized image generation is to produce distinct variations of a custom subject while preserving its visual appearance. Most existing approaches achieve this goal by finetuning a pre-trained model on a reference set of images to make it memorize the custom concept. Although these methods can yield good synthesis results, they require substantial resources and a long training time to fit the custom subject, and more than one reference image is needed to avoid overfitting. To overcome these challenges, there has been recent interest in developing *finetuning-free personalization* methods [16, 30, 32, 33]. These methods typically encode the reference images into a compact feature space, and condition the diffusion model on the encoded features. However, the encoding step results in information loss, leading to poor appearance preservation, especially for challenging unusual objects as seen in Fig. 1. Therefore, the performance of encoder-based personalization techniques is inferior to finetuning-based approaches.

In this paper, we present JeDi, a novel approach for finetuning-free personalized text-to-image generation that excels at preserving input reference content. Our core idea is to train a diffusion model to learn a joint distribution of multiple related text-image pairs that share a common subject. As illustrated in Fig. 2, this goal is achieved using two key ingredients: First, we construct a synthetic dataset of related images in which each sample contains a set of text-image pairs that share a common subject. We present a scalable approach for creating such a dataset using LLMs and pre-trained single-image diffusion models. Second, we modify the architecture of existing text-to-image diffusion models to encode relationships between multiple images in a sample set. Specifically, we adapt the self-attention layers of the diffusion U-Net so that the attention blocks corresponding to different input images are coupled. That is, the self-attention layer corresponding to each image co-attends to every other image in the sample set. The use of the coupled self-attentions at different levels of hierarchy in the

Table 1. In contrast to prior work, JeDi does not require finetuning or the use of image encoders for image personalization.

Method	Finetuning-free	Encoder-free
DreamBooth [27]	✗	✓
CustomDiffusion [15]	✗	✓
ELITE [32]	✓	✗
BLIPDiffusion [16]	Optional	✗
JeDi	✓	✓

U-Net provides a much stronger representation needed for good input preservation.

At test time, JeDi can take multiple text prompts as input and generate images of the same subject in different contexts. By simply substituting reference images as observed variables in the sampling process, JeDi can generate personalized images based on any number of reference images. We utilize guidance techniques [11] on reference images to further improve the image alignment. JeDi can achieve high-fidelity personalization results even in challenging cases involving unique subjects (Fig. 1, 7), using as few as a single reference image.

Our **key contributions** are summarized as follows:

- We propose a finetuning-free text-to-image generation method with a novel joint-image diffusion model.
- We present a simple and scalable data synthesis pipeline for generating a multi-image personalization dataset with images sharing the same subject.
- We design novel architectural and sampling techniques such as coupled self-attention and image guidance for achieving high-fidelity personalization.

## 2. Related Work

### 2.1. Text-to-Image Generation

Denosing diffusion models [8, 12, 31] formulate the image generation task as a series of progressive denoising steps. The denoising network can be trained conditioned on text embeddings to generate images from an input caption. DALL-E2 [25] achieves high-resolution text-to-image synthesis using two diffusion models: the first model converts a CLIP text embedding to a CLIP image embedding, while the second model converts the image embedding to an output image. Imagen [28] trains a cascaded diffusion model conditioned on T5 language embeddings [24]. eDiff-I [2] uses an ensemble of expert denoisers to increase the model capacity, with each expert specializing in a specific noise range. Latent diffusion models [6, 22, 26] train the diffusion model in a compact latent space of an autoencoder for efficient training and sampling.

### 2.2. Personalized Text-to-Image Generation

**Finetuning based methods.** Most prior works achieve image personalization by finetuning the diffusion model on a

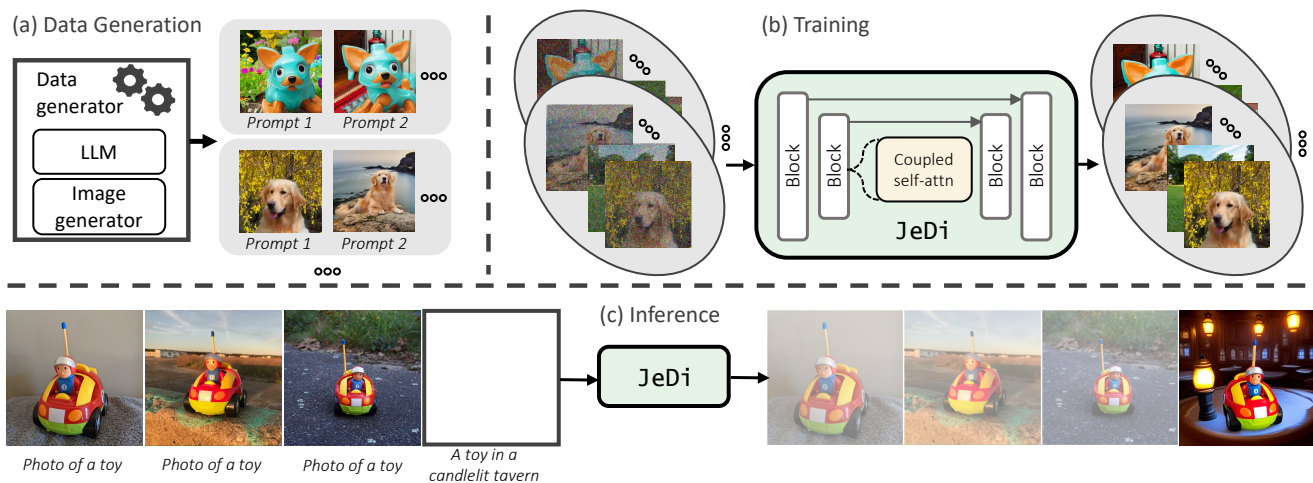


Figure 2. **Overall framework.** (a) We generate training data by using large language models and prompting pretrained single-image diffusion models. (b) During training, the JeDi model learns to denoise multiple same-subject images together, where each image attends to every image of the same subject set through coupled self-attention. (c) At inference, personalized generation is performed in an inpainting fashion where the goal is to generate the missing images of the joint-image set.

custom dataset. Dreambooth [27] finetunes the entire model weights on the reference set, with a loss on images of similar concepts as regularization. CustomDiffusion [15] optimizes only a few parameters to enable fast tuning, and combines multiple finetuned models for multi-concept personalization. Textual Inversion [9] projects the reference images onto the text embedding space through an optimization process. SVDiff [10] finetunes only the singular values of the weight matrices to reduce the risk of overfitting. These finetuning-based methods require a substantial amount of resources and long training time, and often need multiple reference images per custom subject.

**Finetuning-free methods.** To improve the efficiency of image customization, there has been recent interest in developing finetuning-free methods. These approaches typically use an image encoder to encode a reference image onto a compact feature space, and train the diffusion model conditioned on this feature vector. BLIPDiffusion [16] uses BLIP-2 [17] encoder, while FastComposer [33] uses a CLIP [23] encoder for image encoding. ELITE [32] and InstantBooth [30] use a learnable image encoder trained jointly with the diffusion model. These encoder-based methods produce reasonable results for common subjects, but often fail to generate uncommon subjects and preserve fine-grained details due to the information loss in the encoding step. In contrast, our approach directly trains a joint-image diffusion model without an encoding step, resulting in better input preservation even for challenging objects.

**Personalization dataset.** To achieve finetuning-free personalization, a training dataset comprised of same-subject image sets is required. Methods like [16, 30, 32, 33] rely on image augmentation and background removal to construct training data, which often does not provide sufficient variations for the same subject. To improve the

diversity, we present a scalable data generation approach by prompting pre-trained single-image diffusion models to produce multi-image photo collages with good variation.

## 3. Method

### 3.1. Dataset Creation

Training a model to produce a joint distribution of multiple same-subject images requires a dataset where each sample is a set of images sharing a common subject. While there exist some same-subject datasets such as CustomConcept101 [15] and DreamBooth [27], they are small in scale and lack sufficient variations desired for diffusion model training. Therefore, we create a diverse large-scale dataset of image-text pairs containing the same subject, called the Synthetic Same-Subject ( $S^3$ ) dataset, using large language models [21] and single-image diffusion models [22].

Fig. 3 illustrates our data generation process. We first start with a list of common objects and prompt ChatGPT to generate a text description for each object in the list. Then, we use the pre-trained SDXL [22] model to generate a dataset of same-subject photo collages by appending the text “*photos of the same*” to each of the text prompt generated in the previous step. We observed that by prompting the SDXL model this way, it can generate photo collages of the same subject with varying poses. However, the generated images usually contain a close-up view of an object in a simple background. To increase the data diversity, we employ a post-processing step that performs background augmentation on the generated objects.

Given a generated photo collage, we first run object detection [18] and segmentation [14] to separate out object instances and extract foreground region. We discard pairs of instances with CLIP [23] image scores lower than 0.95

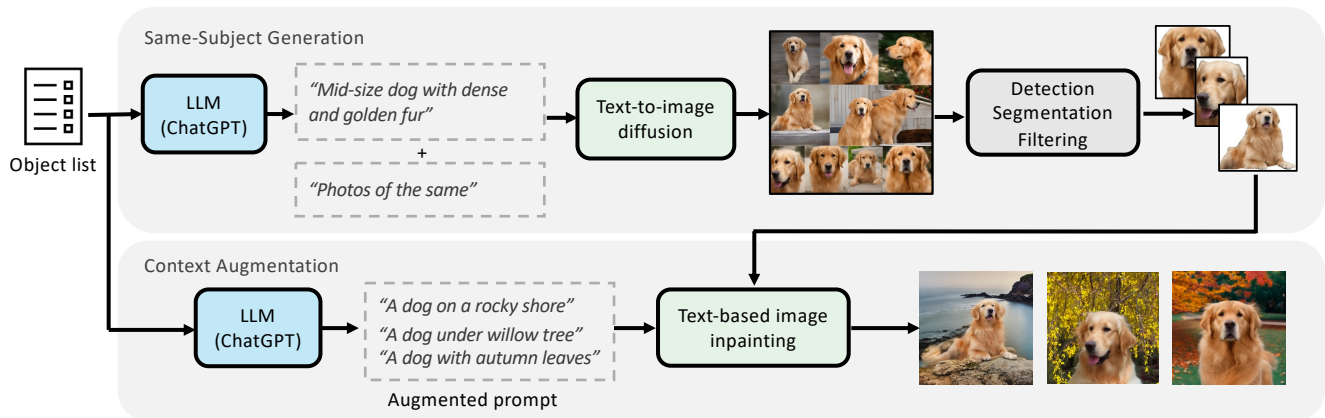


Figure 3. **Data generation process.** We construct Synthetic Same-Subject ( $S^3$ ) dataset by first prompting the pretrained text-to-image diffusion models to generate same-subject photo collages, and then increasing the diversity using text-based background inpainting.



Figure 4. **Samples from the synthetic same-subject ( $S^3$ ) dataset.** Each column denotes different images from one joint-data sample.

as they may not contain the same subject. We then paste the object at a random location in a blank image, and use the stable diffusion inpainting model to inpaint the background based on a new prompt related to the same object name. In addition, we use InstructPix2Pix [4] to stylize the generated samples with a probability of 0.5 to increase style variation using randomly selected style prompts. Fig. 4 shows some examples of the text-image data generated through this process. The generated samples have consistent subjects with good diversity and pose variations.

### 3.2. Joint-Image Diffusion

The goal of training a joint-image diffusion model is to generate multiple related images sharing the same subject. Conventional diffusion models [2, 20, 26, 28], however, can only generate individual images independently as the network architecture does not have any connections between different samples in a batch. We found that with a simple modification, a single-image diffusion model can be adapted to a joint-image model which can generate images having related content (such as the same custom subject).

More specifically, given a set of same-subject noisy input images, we modify the attention layers of the U-Net to

fuse together the self-attention features for different images in the same set. As illustrated in Fig. 6, a coupled self-attention layer has features at each spatial location attending to every other location across all images in the set. Since the U-Net architecture has attention layers at various resolutions, the use of coupled self-attentions at multiple resolutions makes the generated image set to have consistent high-level semantic features as well as low-level attributes. Fig. 5 visualizes the pixel-wise correspondences and attention heat maps of coupled self-attention layers. We observe that the co-attended regions across different images form the right correspondences across all resolutions.

After a coupled self-attention layer, the output is fed to a regular cross-attention layer, which aligns the visual feature of each image to the corresponding text prompt. The coupled self-attention layer can be implemented by simply adding two reshaping operations before and after a regular self-attention, thus enabling simple and easy adaptation from pre-trained single-image diffusion models.

Fig. 2 (b) illustrates the training process. We start by creating noisy same-subject data by adding isotropic Gaussian noise, and train the joint-image diffusion model to denoise the data. Ideally, there is no limit on the size of each image set of the same subject. In our experiments, we randomly set the size to 2, 3, or 4 during training. The training loss of JeDi is very similar to that of a regular diffusion model. We use  $\epsilon$ -prediction and a simplified training objective introduced in [12]. The loss function is as follows,

$$L = \mathbb{E}_{\epsilon \sim \mathcal{N}(0,1), t \sim [1,T]} [\|\epsilon - \epsilon_{\theta}(\mathbf{x}_t)\|_2^2], \quad (1)$$

where  $\epsilon_{\theta}$  represents the network parameterized by  $\theta$ ,  $T$  is the number of diffusion steps,  $\mathbf{x}_t$  is the  $t$ -step noised image set of size  $N$ , i.e.  $\mathbf{x} = [x^1, x^2, \dots, x^N]$ . We omit the text and timestep conditioning in  $\epsilon_{\theta}(\cdot)$  for simplicity. After training, the JeDi model can take multiple text prompts as input and generate images containing the same subject.

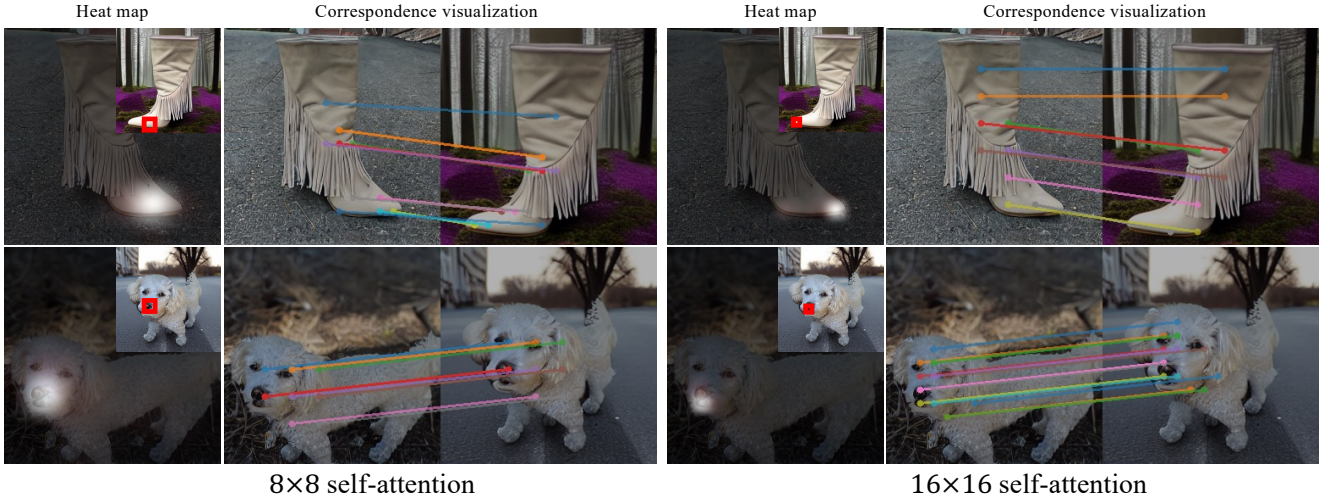


Figure 5. **Visualization of the coupled self-attentions.** For both scales (8x8 and 16x16), the correspondence map (Corr.) shows the connections with the highest weights between elements in the two images. The heatmap visualizes the distribution of the attention weights in an image for a specific element in another image (marked with a red box). We observe that similar regions in different images are co-attended in the coupled self-attention layers.




Figure 6. **Attention architecture.** In a coupled self-attention block, features corresponding to each spatial location attends to every location across all images in the image set. Following the coupled self-attention block, features of each individual image attend to their respective text embedding in cross-attention layers.

### 3.3. Personalized Text-to-Image Generation

**Personalization as inpainting.** While the joint-image diffusion model discussed in the previous section can generate same-subject images, it does not input a reference image that needs to be personalized. In this section, we propose to solve the input image personalization problem by casting it as an inpainting task. That is, given a few text-image pairs as reference, the task of generating a new personalized sample can be viewed as inpainting the missing images of a joint-image set containing reference images (Fig. 2 (c)).

We design the inpainting model by modifying the input layer of the diffusion U-Net so that it can be conditioned on reference images. More specifically, the input to the diffusion model is a concatenated list of noisy images, reference images and a binary mask indicating whether the reference image is used or not. When the binary mask is all 0’s, the reference image is used. On the other hand, when the binary mask is all 1’s, the reference image is an empty black image (indicating the missing images that need to be gen-

erated). During training, for every image in the joint-image set, we use the reference image with a probability of 0.5 i.e., we assign the binary mask to 0 with probability 0.5. The training loss can then be written as follows,

$$L = \mathbb{E}_{\epsilon \sim \mathcal{N}(0,1), t \sim [1,T], \mathbf{m} \sim \binom{N}{0.5}} [\|\epsilon - \epsilon_{\theta}(x_t, \hat{\mathbf{x}}, \mathbf{M})\|_2^2], \quad (2)$$

where  $\mathbf{M}$  is the spatially repeated tensor of a binomial variable  $\mathbf{m}$ ;  $\hat{\mathbf{x}} = \mathbf{x} \odot \mathbf{M}$  denotes the reference images, with the unknown elements set to zero.

During inference, we utilize the replacement trick [19, 31] in which the known part of the joint-image set is replaced with the forward diffusion response of the clean reference image. Let  $\hat{\mathbf{x}} = [\hat{x}^1, \hat{x}^2, \dots, \hat{x}^n]$  be the reference images, which are the known elements in the image set  $\mathbf{x} = [\hat{x}^1, \dots, \hat{x}^n, x^{n+1}, \dots, x^N]$  to be generated. During sampling, at each diffusion step  $t$ , we only keep the backward diffusion output for the unknown elements  $x^{n+1}, \dots, x^N$  while replacing the known part with the forward diffusion output, i.e. the noised real images  $\hat{x}_t^1, \dots, \hat{x}_t^n$ .

**Image guidance.** Classifier-free guidance is a popular technique used in single-image diffusion models to make the image generations more aligned to the input conditioning [11]. To improve the faithfulness of the generated samples to the input reference images, we use image guidance in addition to the text guidance during sampling. The score function with the use of image guidance is as follows:

$$\tilde{\epsilon}(x_t, \hat{\mathbf{x}}, \mathbf{M}) = \epsilon^0 + \lambda[\epsilon_{\theta}(x_t, \hat{\mathbf{x}}, \mathbf{M}) - \epsilon^0] \quad (3)$$

where  $\epsilon^0 = \epsilon_{\theta}(x_t, \mathbf{0}, \mathbf{M})$  represents the unconditional score when the text prompt and all reference images are set to null;  $\lambda$  is the guidance scale. We find that the use of image guidance can significantly improve the fidelity to the input reference images.



Figure 7. **Personalized text-to-image generations.** We show the generations obtained by our JeDi model on challenging uncommon input subjects shown in the left column. JeDi accurately preserves the reference image content while being faithful to the text prompt.

## 4. Experiments

**Dataset.** We construct the Synthetic Same-Subject dataset ( $S^3$  dataset) for training as described in Sec. 3.1. After CLIP filtering, we obtain  $1.6M$  sets of images with each set containing 2-4 images. We also include the video frames from WebVid10M [1] and rendered multi-view images from Objaverse [7] during training, as the frames from the video and the rendered images from the asset usually have the same subject. We use the original video caption or asset caption as text prompts for all images obtained from the same video/asset. Additionally, we include the single-image data from LAION aesthetic dataset [29] and use a set size of 1 for these images. To evaluate our models, we use the test dataset proposed in DreamBooth [27]. DreamBooth test set contains 30 real-world subjects with 4-6 images and 25 prompts for each subject.

**Evaluation metrics.** The two main evaluation criteria for personalized text-to-image generation include (1) alignment between generated images and the input text prompts, and (2) faithfulness of the generations to the input reference images. We use the CLIP image-text similarity (CLIP-T)

between the generated images and the input captions for (1). For (2), we follow prior works [16, 27, 32] and use the cosine similarity of CLIP [23] image embedding (CLIP-I) and DINO [5] image embedding (DINO) between the generated images and the reference images. DINO is considered to be a preferred metric for measuring image similarity as it is sensitive to the appearance variations of different images in the same concept class. Additionally, we also report CLIP-I and DINO scores only on the foreground masked images, i.e., images with foreground objects cutout using object detection and segmentation [13, 18]. This helps remove the background variations when computing the image similarity scores to better reflect the faithfulness to the reference subject. We call these metrics MCLIP-I and MDINO.

**Implementation details.** We implement our method based on StableDiffusion V1.4 to enable fair comparison with prior approaches. We train the model with batch size 2048 and a learning rate  $5e-5$ . We initialize the weights using the pre-trained StableDiffusion model. For each batch, we randomly sample image sets from  $S^3$ , WebVid10M, Objaverse and LAION datasets with equal probability. We ran-




Figure 8. **Visual comparison with finetuning-free methods.** JeDi can faithfully preserve the details of input content even for challenging uncommon objects (row 2 and 3). ELITE and BLIPD fail in such cases and can get good results only in common object classes (row 1).



Figure 9. **Visual comparison with finetuning methods.** Finetuning methods memorize the input images (DB outputs) or result in poor input preservation (CD outputs) when the number of reference images is small. On the other hand, JeDi can generate images that are faithful both to the input images and the text prompt.

domly choose the image set size between 2-4 except for LAION, where the set size is always 1. Our model takes 36 hours to train on 32 A100 GPUs for 140K steps.

#### 4.1. Comparison

We compare our method with the state-of-the-art finetuning-free methods - BLIPD [16] and ELITE [32], along with two finetuning-based methods - DreamBooth (DB) [27] and CustomDiffusion (CD) [15]. For each subject, we use one input reference image for finetuning-free approaches and three for finetuning-based approaches.

**Comparison to finetuning-free methods.** Fig. 8 shows

the visual comparison of our results to BLIPD and ELITE. It can be seen that our method can faithfully capture the visual features of the input reference image, including both semantic attributes and low-level details. However, the generations from BLIPD and ELITE can only roughly resemble the color patterns and semantic features of the input. We also observe that BLIPD and ELITE can produce reasonable results for common object classes such as dogs. This is because their encoder can easily recognize the common object categories (such as the dog breed) which makes the personalized generation easier. However, for unique uncommon objects, their results tend to be much worse, e.g. the

Table 2. **Quantitative comparisons.** All finetuning-based methods use 3 input images, while the finetuning-free methods use 1 input image. JeDi model with 1 input image outperforms all finetuning-free baselines, while the JeDi model with 3 input images outperforms all finetuning baselines. JeDi obtains a much higher masked DINO scores, which suggests that we achieve stronger input subject preservation compared to other baselines.

	Method	CLIP-T (↑)	CLIP-I (↑)	MCLIP-I (↑)	DINO (↑)	MDINO (↑)
Finetuning-based	DreamBooth [27]	0.2812	0.8135	0.8683	0.6341	0.7115
	Custom Diffusion [15]	0.3015	0.7952	0.8640	0.6343	0.7109
Finetuning-free	BLIP Diffusion [16]	0.2934	0.7899	0.8620	0.5855	0.6692
	ELITE [32]	0.2961	0.7924	0.8615	0.5922	0.6805
	JeDi (1 input)	<b>0.3040</b>	0.7818	0.8764	0.6190	0.7510
	JeDi (3 inputs)	0.2932	<b>0.8139</b>	<b>0.9011</b>	<b>0.6791</b>	<b>0.8037</b>

Table 3. **Effect of varying the S<sup>3</sup> dataset size.**

	Dataset size			
	0	0.6M	1.2M	1.6M
DINO (↑)	0.8895	0.7501	0.7443	0.7467
MDINO (↑)	0.8931	0.8639	0.8572	0.8636
CLIP-T (↑)	0.2524	0.3020	0.3017	0.3015

Table 4. **Ablation of different joint-image diffusion designs.**

	CLIP baseline	JeDi	(+) CLIP emb	w/o IG
DINO (↑)	0.3411	0.7501	0.7432	0.4652
MDINO (↑)	0.4394	0.8639	0.8617	0.5922
CLIP-T (↑)	0.3325	0.3020	0.3041	0.3259

toy in the second row and the image in the third row. Note that even for the common classes such as the dog example in the first row, the generations from BLIPD and ELITE can miss some input features (different haircuts despite being the same breed). In contrast, our method eliminates information loss caused by the encoder and results in much better preservation of the custom concept. This is also reflected in the quantitative comparison Table 2, where our method outperforms BLIPD and ELITE by a large margin.

**Comparison to finetuning-based methods.** Fig. 9 shows the visual comparison of our approach with DreamBooth (DB) and CustomDiffusion (CD). When the number of reference images is limited, it is challenging for finetuning-based methods to avoid overfitting and generate novel variations of the input subject. From Fig. 9, we see that DB often directly copies the input image due to overfitting, while the images generated by CD do not faithfully preserve the features of the input subject. In contrast, our method creates proper variations of the reference subjects without changing its key visual features. Even without any expensive finetuning, our method outperforms DB and CD in quantitative comparisons when we provide the same number of reference images (JeDi-3), as shown in Table 2.

## 4.2. Ablation Study

**Size of S<sup>3</sup> dataset.** Table 3 reports the results of training our model using different numbers of synthetic images. Dataset size 0 refers to training the model on only videos and multi-view images. This setting yields the best image alignment (DINO and MDINO) and the worst text alignment as the model learns a shortcut to ignore the text prompts and copy the input images. The performance of columns 2-4 are roughly similar, which shows that we do not obtain much gains by increasing the size of the synthetic dataset.

**Joint-image diffusion model.** In Table. 4, we report the ablation study of different design choices in the training of JeDi. The first column shows the results of a CLIP encoder baseline, which is a single-image diffusion model conditioned on the CLIP image features of the reference image. Our JeDi model yields a much better image alignment than the CLIP encoder baseline, which demonstrates the advantages of using the joint-image model over the image encoders. We also find that adding CLIP image embedding as extra conditional input to JeDi (+ CLIP emb) does not improve the performance as shown in column 3. This implies that the joint-image model already captures the information extracted in the image embedding. The last column reports the results without image guidance (w/o IG). By comparing the second and the last columns, we can see that image guidance is crucial to obtaining good personalization results.

## 5. Conclusion

This paper presents JeDi, a novel approach for finetuning-free personalized text-to-image generation using a joint-image diffusion model. We show how a single image diffusion U-Net can be adapted to learn a joint image distribution using coupled self-attention layers. To train the joint-image diffusion model, we construct a synthetic dataset called S<sup>3</sup>, in which each sample contains a set of images sharing the same subject. The experimental results show that the proposed JeDi model outperforms the previous approaches both quantitatively and qualitatively in benchmark datasets.



## References

- [1] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *International Conference on Computer Vision*, 2021. 6
- [2] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022. 2, 4
- [3] James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, Wesam Manassra, Prafulla Dhariwal, Casey Chu, Yunxin Jiao, and Aditya Ramesh. Improving image generation with better captions. <https://api.semanticscholar.org/CorpusID:264403242>. 2
- [4] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2023. 4
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *International Conference on Computer Vision*, 2021. 6
- [6] Xiaoliang Dai, Ji Hou, Chih-Yao Ma, Sam Tsai, Jialiang Wang, Rui Wang, Peizhao Zhang, Simon Vandenhende, Xiaofang Wang, Abhimanyu Dubey, et al. Emu: Enhancing image generation models using photogenic needles in a haystack. *arXiv preprint arXiv:2309.15807*, 2023. 2
- [7] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2023. 6
- [8] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Conference on Neural Information Processing Systems*, 2021. 2
- [9] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit Haim Bermano, Gal Chechik, and Daniel Cohen-or. An image is worth one word: Personalizing text-to-image generation using textual inversion. In *International Conference on Learning and Representation*, 2022. 3
- [10] Ligong Han, Yinxiao Li, Han Zhang, Peyman Milanfar, Dimitris Metaxas, and Feng Yang. Svdiff: Compact parameter space for diffusion fine-tuning. *arXiv preprint arXiv:2303.11305*, 2023. 3
- [11] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS Workshop on Deep Generative Models and Downstream Applications*, 2021. 2, 5
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Conference on Neural Information Processing Systems*, 2020. 2, 4
- [13] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 6
- [14] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 3
- [15] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2023. 2, 3, 7, 8
- [16] Dongxu Li, Junnan Li, and Steven CH Hoi. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. *arXiv preprint arXiv:2305.14720*, 2023. 2, 3, 6, 7, 8
- [17] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 3
- [18] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 3, 6
- [19] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 5
- [20] 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. 4
- [21] OpenAI. Chatgpt. <https://openai.com/blog/chatgpt>. Accessed: 2023. 3
- [22] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 2, 3
- [23] 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. 3, 6
- [24] 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. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020. 2
- [25] 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. 2
- [26] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 2, 4

- [27] 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 *IEEE Conference on Computer Vision and Pattern Recognition*, 2023. [2](#), [3](#), [6](#), [7](#), [8](#)
- [28] 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. In *Conference on Neural Information Processing Systems*, 2022. [2](#), [4](#)
- [29] 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. [6](#)
- [30] Jing Shi, Wei Xiong, Zhe Lin, and Hyun Joon Jung. Instantbooth: Personalized text-to-image generation without test-time finetuning. *arXiv preprint arXiv:2304.03411*, 2023. [2](#), [3](#)
- [31] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *Proc. ICLR*, 2021. [2](#), [5](#)
- [32] Yuxiang Wei, Yabo Zhang, Zhilong Ji, Jinfeng Bai, Lei Zhang, and Wangmeng Zuo. Elite: Encoding visual concepts into textual embeddings for customized text-to-image generation. *arXiv preprint arXiv:2302.13848*, 2023. [2](#), [3](#), [6](#), [7](#), [8](#)
- [33] Guangxuan Xiao, Tianwei Yin, William T Freeman, Frédo Durand, and Song Han. Fastcomposer: Tuning-free multi-subject image generation with localized attention. *arXiv preprint arXiv:2305.10431*, 2023. [2](#), [3](#)