

DreamBlend: Advancing Personalized Fine-tuning of Text-to-Image Diffusion Models

Shwetha Ram, Tal Neiman, Qianli Feng, Andrew Stuart, Son Tran, Trishul Chilimbi
Amazon

{shweram, taneiman, fengq, andrxstu, sontran, trishulc}@amazon.com

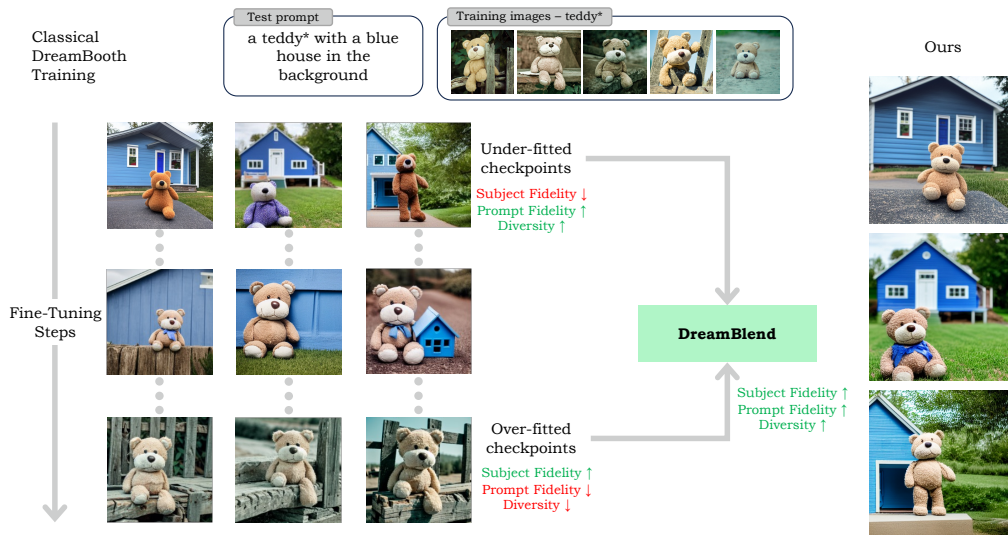


Figure 1. DreamBlend merges prompt fidelity and diversity from underfit checkpoints with subject fidelity from overfit checkpoints during image generation. The generated images have the layout of the underfit images and the subject fidelity of the overfit images, achieving better subject fidelity, prompt fidelity and diversity.

Abstract

Given a small number of images of a subject, personalized image generation techniques can fine-tune large pre-trained text-to-image diffusion models to generate images of the subject in novel contexts, conditioned on text prompts. In doing so, a trade-off is made between prompt fidelity, subject fidelity and diversity. As the pre-trained model is fine-tuned, earlier checkpoints synthesize images with low subject fidelity but high prompt fidelity and diversity. In contrast, later checkpoints generate images with low prompt fidelity and diversity but high subject fidelity. This inherent trade-off limits the prompt fidelity, subject fidelity and diversity of generated images. In this work, we propose DreamBlend to combine the prompt fidelity from earlier checkpoints and the subject fidelity from later checkpoints during inference. We perform a cross attention guided image synthesis from a later checkpoint, guided by an image

generated by an earlier checkpoint, for the same prompt. This enables generation of images with better subject fidelity, prompt fidelity and diversity on challenging prompts, outperforming state-of-the-art fine-tuning methods.

1. Introduction

Text-to-Image models like Stable Diffusion [33] enable generating images from text prompts with broad capabilities. However, users often seek personalized image generation for specific subjects in diverse contexts. This involves synthesizing novel images conditioned on text prompts using a few subject images. For example, in Fig. 1, we generate images of subject *teddy** “with a blue house in the background”. Successful image generation hinges on two criteria: subject fidelity, ensuring the *teddy** in generated images is “same” as that in input training images, and prompt fidelity, adhering to the text prompt of a blue house in the



Figure 2. Images generated by different checkpoints for prompt: ‘a *backpack** on a cobblestone street’ as a SD model is fine-tuned from 5 to 1000 steps. Early checkpoints have higher prompt fidelity and diversity but lower subject fidelity while later checkpoints have higher subject fidelity but lower prompt fidelity and diversity. At step=1000, the model reproduces the input images used for fine-tuning.

background. Additionally, diversity is desired, showcasing *teddy** in varied poses and environments with diverse elements like driveways, lawns and trees.

Prior works have embedded subject identity in an input word embedding [9], layers of a pre-trained model through fine-tuning [34], or both [17]. Fine-tuning based approaches pose inherent trade-offs between subject fidelity, prompt fidelity and diversity, presented in Fig. 1. A pre-trained model excels at generating high-fidelity images for various prompts by leveraging its world knowledge. Fine-tuning enhances subject fidelity but reduces prompt fidelity and diversity due to overfitting, language drift and catastrophic forgetting. Over time, generated images closely resemble the training inputs as shown in Fig. 2.

Existing methods seek a “sweet spot” between prompt fidelity, subject fidelity, and diversity using intermediate checkpoints. We introduce *DreamBlend*, which merges the strengths of early and late checkpoints during image generation. This approach provides superior trade-offs by combining early checkpoint prompt fidelity and diversity with late checkpoint subject fidelity. A careful study of image generation in early and later checkpoints reveals the phenomenon of catastrophic attention collapse in later checkpoints. This leads to our key insight of using cross attention guidance to preserve the prompt fidelity in early checkpoint images. By synthesizing images from a later checkpoint guided by an image from an earlier one for the same prompt, we achieve improved subject fidelity, prompt fidelity and diversity, as demonstrated in Fig. 1.

In summary, our contributions are as follows: first, we study different operating points in current finetuning-based text-to-image personalization methods and observe a catastrophic attention collapse in later checkpoints that diminishes prompt fidelity and diversity; second, we propose a novel approach, *DreamBlend*, of using cross attention guidance to combine the prompt fidelity and diversity of early checkpoints with the subject fidelity of later checkpoints during image generation and find that this regularization on cross attention maps is effective at minimizing the effect of over-fitting; finally, we demonstrate that *DreamBlend*

produces superior images with enhanced subject fidelity, prompt fidelity, and diversity on challenging prompts, surpassing existing state-of-the-art fine-tuning based methods.

2. Related work

2.1. Text-to-image diffusion models

Text-to-image diffusion models are trained to generate samples from a conditional data distribution by the gradual denoising of a variable sampled from a Gaussian distribution. Recent progress in text-to-image synthesis has been fuelled by large models trained on web scale data [31–33, 36, 39, 47]. We leverage the capabilities of such models for personalized image generation.

2.2. Personalized text-to-image diffusion models

Generating specific subjects with pre-trained text-to-image models via prompt engineering is difficult unless those subjects were well-represented in the training data. Consequently, efforts have emerged to teach specific subjects post-training [1, 9, 10, 12, 16, 17, 19, 24, 34, 35, 37, 40, 42, 44–46, 48].

New subject identities can be added through input word embeddings [1, 9, 42, 45], fine-tuning model weights [34] or both [17]. While fine-tuning often provides better subject fidelity and photorealism due to greater expressive power, it suffers from loss of prompt fidelity and diversity as fine-tuning continues. This can be attributed to over-fitting and language drift, observed in both language models [18, 22] and text-to-image diffusion models [17, 34, 40]. To alleviate this, DreamBooth [34] regularizes the network with its own generated images while Custom Diffusion [17] fine-tunes only text-image cross attention weights and uses class-specific retrieved images for regularization. Perfusion [40] performs a gated rank one update inspired by ROME [26] on key and value projection matrices, locking the key matrices to the subject’s super category to reduce over-fitting. Our approach differs by using cross-attention guidance to synthesize images from an overfit checkpoint, guided by an image from an underfit checkpoint for the same prompt.

Encoder-based approaches [4, 10, 16, 24, 37, 44, 46] aim to avoid fine-tuning and storing weights for each subject. Some methods [10, 16, 37, 44] are limited to specific domains like dogs or human faces and may still require some fine-tuning for personalization [10], while others [4, 19] are more generalized with zero-shot capabilities. They often require large scale pre-training with the diffusion model in the loop, with some methods [4] harvesting the pre-training data from expert models that are fine-tuned for each subject. Fine-tuning methods provide a cost-effective means to adapt existing text-to-image models for personalized image generation. They also contribute to generating high-quality data for training models with better inference efficiency. Consequently, our focus lies in enhancing the image quality produced by these fine-tuning based approaches.

2.3. Image editing with diffusion models

Image editing aims to modify specific regions of an input image while preserving the rest. Approaches like SDEdit [25] introduce stochastic noise followed by denoising, which can unintentionally alter non-targeted areas. Using a spatial mask from another model [23] or the diffusion model itself [7] to alleviate this often leads to content in the mask region being ignored and blending artifacts. In contrast to stochastic methods like DDPM [14] and SDEdit [25], deterministic DDIM inversion [38] first inverts an image for subsequent editing. Several schemes [27, 29, 43] have been designed to achieve a more editable reconstruction. Recent works [6, 20] propose personalized editing of real images using personalized text-to-image models. Choi et al. [6] combine methods from [27] and [13], while Li et al. [20] iterate image inpainting with [21], guided by spatial segmentation mask. Our work uses DDIM inversion for the initial latent and cross attention guidance for layout control.

Recent research [3, 8, 11, 13, 29] shows that text-image cross attention maps significantly influence generated image layout. Perfusion and Custom Diffusion only update the text-image cross attention weights. While this affects the cross attention maps generated, they do not directly manipulate cross attention. Attend-and-Excite [3] encourages generation of all subjects in the text prompt by enhancing attention values for the most neglected subject token at each time step. Prompt-to-Prompt [13] directly swaps the attention maps from source image generation into target image generation, for text-driven image editing. Photoswap [11] advocates swapping the self attention maps in addition to cross attention maps for superior layout control. In our work, we employ a cross attention guidance regularization to align target cross attention maps with reference maps from an underfit model. Compared to image editing, we care less about preserving the exact details in the underfit reference image and this formulation allows slight layout changes for improved subject fidelity.

3. Method

During fine-tuning of a pre-trained text-to-image diffusion model, early checkpoints produce images with high prompt fidelity and diversity but low subject fidelity. In contrast, later checkpoints yield images with high subject fidelity but lower prompt fidelity and diversity. Early checkpoints are under-fitted, lacking sufficient subject learning, while later ones are over-fitted to the subject appearance, pose and environments in the input images. A naïve way to combine the best of both would involve overlaying a subject from a later checkpoint onto an image from an earlier checkpoint. This method would likely introduce artifacts due to mismatches in pose, lighting, and other factors. However, this thought experiment inspires new methodologies aimed at integrating the advantages of both under-fitted and over-fitted checkpoints effectively.

In text-image diffusion models, text conditioning is achieved through cross-attention layers spanning various scales of the denoising network, significantly influencing the generated image layout [3, 8, 13, 28, 29, 48]. Each word in the input text prompt is tokenized and encoded using a text encoder model like CLIP [30]. These text embeddings form the keys $K^{N \times d}$ for cross-attention, where N is the number of tokens and d is the embedding dimension. Queries $Q^{M \times d}$ represent intermediate image features at each cross-attention layer, where M represents image patches treated as tokens, sharing the embedding dimension d with text embeddings. The attention map $A^{M \times N} = \text{softmax}(QK^T/\sqrt{d})$, tells us how much each text token attends to each image patch. These maps are reshaped to $N \times H \times W$ and visualized in Fig. 3, where $H \times W = M$ correspond to image spatial dimensions.

In Fig. 3, we visualize the text-image cross attention maps during fine-tuning of a StableDiffusion model on the images of *backpack** from Fig. 2, represented by the rare text token “sks” following DreamBooth [34, 41]. All images stem from same initial latent and prompt “a sks backpack on a cobblestone street”, using same classifier-free guidance and 50 steps of DDIM forward process. At step 5, the model achieves high prompt fidelity and the text-image cross attention for different words in the prompt is focused on relevant parts of the image. By step 250, the model has learnt the subject well and generates images with higher subject fidelity. This is evident as the “sks” token’s attention concentrates on a crucial feature of the backpack: its logo. At the same time, the attentions for other words like “cobble”, “stone” and “street” are also beginning to focus on *backpack**. By step 1000, over-fitting causes all text tokens to excessively focus on the subject. We call this phenomenon catastrophic attention collapse.

At step 1000, the model becomes highly over-fitted, mapping all latents to one of the input images of *backpack** used for fine-tuning. In contrast, an intermediate model









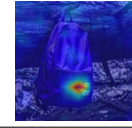
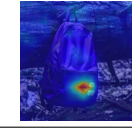
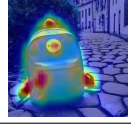
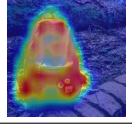
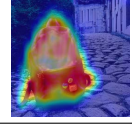
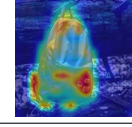
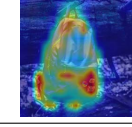

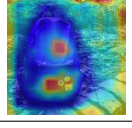
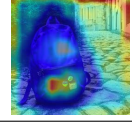
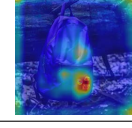
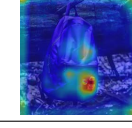



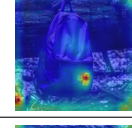




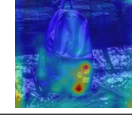
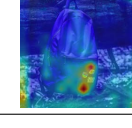
steps \ tokens	step 5	step 250	step 250 + CAG	step 1000	step 1000 + CAG
					
sks					
backpack					
cobble					
stone					
street					

Figure 3. Attention Guidance and Attention Collapse: Images generated at 5, 250, and 1000 steps of DreamBooth fine-tuning, with text-image cross attention maps for *backpack** in Fig. 2. The prompt “a sks backpack on a cobblestone street” features the rare token “sks” as *backpack**. Step 250 + CAG: Cross attention guidance (CAG) from step 5 image is effective. The resulting image maintains layout of step 5 image, while preserving subject fidelity. Step 1000 + CAG: By step 1000, over-fitting leads to catastrophic attention collapse, focusing attention of all tokens mostly on the subject. CAG becomes ineffective as the model maps all latents to one of the input images.

like the one at step 250 retains semantic understanding of concepts like “cobble”, “stone” and “street” from its world knowledge, although their appearances have morphed closer to concepts seen in the input images, as evident from the bush in the step 5 image being replaced by the tree bark of *backpack** input images in the step 250 image. Our key finding is that this step 250 model can be guided to generate an image that follows the layout of the image generated by the early step 5 model, with cross attention guidance. This result is shown in Fig. 3: observe that step 250 + CAG (Cross Attention Guidance) image follows the layout of the step 5 image. Moreover, the generated backpack is close to our *backpack** as the step 250 model used for image generation has learnt the concept of *backpack** well. This enables us to generate an image with both high prompt fidelity and high subject fidelity.

To achieve this, we implement guided image synthesis

where an early checkpoint serves as guidance model G and a later checkpoint as edit model E . Starting with a text prompt P , we first generate a reference image using G , conditioned on text features c derived from P . This image prioritizes prompt fidelity but may lack subject fidelity. To utilize the layout of this image as a guide for our final image, we store cross attention maps at each timestep.

Next, we generate the final output image using E , starting from the same initial latent used for G and performing cross attention guidance at every step of the diffusion process. At each step, we update the latent in a direction that encourages the current cross attention maps to be close to the reference cross attention maps obtained from G . This is achieved by a regularization loss R , defined to be the absolute difference of the reference and current cross attention maps and a scalar α , that controls the amount of update. These steps are summarized in Algorithm 1.

Algorithm 1 Cross Attention Guided Image Synthesis

Input: Text features c , Guidance model G , Edit model E **Output:** Personalized Image I_o

- 1: Step 1: Store reference cross attention maps from G
 - 2: $t \leftarrow 0$
 - 3: $l \leftarrow l_i \in N(0, 1)$
 - 4: $A_{ref} \leftarrow \emptyset$
 - 5: **while** $t \leq T$ **do**
 - 6: $\epsilon, \text{crossAttn} \leftarrow G(l, t, c)$
 - 7: $l \leftarrow DDIM(l, t, \epsilon)$
 - 8: $A_{ref}[t] \leftarrow \text{crossAttn}$
 - 9: **end while**
 - 10: Step 2: Synthesize final image from E , with cross attention guidance
 - 11: $t \leftarrow 0$
 - 12: $l \leftarrow l_i$
 - 13: **while** $t \leq T$ **do**
 - 14: $-, \text{crossAttn} \leftarrow E(l, t, c)$
 - 15: $R \leftarrow |A_{ref}[t] - \text{crossAttn}|$
 - 16: $l \leftarrow l - \alpha * \nabla_l(R)$
 - 17: $\epsilon, - \leftarrow E(l, t, c)$
 - 18: $l \leftarrow DDIM(l, t, \epsilon)$
 - 19: **end while**
 - 20: $I_o \leftarrow VAE_{decode}(l)$
-

4. Experiments and results

4.1. Dataset and evaluation metrics

We use the standard DreamBooth benchmark dataset and evaluation metrics [34], used in many works [4, 19, 34]. This comprises of 30 subjects from different categories like pets, toys and other objects and 25 prompts for each subject. Subject fidelity and prompt fidelity are important criteria for successful text-to-image personalization. Following prior works [4, 17, 19, 34, 35], we use CLIP [30] image similarity (CLIP-I) and DINO [2] for subject fidelity and CLIP text similarity (CLIP-T) for prompt fidelity. CLIP-I and DINO are the average pairwise cosine similarities between CLIP and DINO embeddings of input and generated images, respectively. CLIP-T is the average cosine similarity between CLIP embeddings of the text prompt and generated images.

4.2. Description of methods

We compare to two methods that fine-tune weights of the text-to-image diffusion model for personalization. DreamBooth [34] uses a unique rare token to represent a subject, such as “a sks teddy,” and fine-tunes all model weights. Custom Diffusion [17] fine-tunes only text-to-image cross attention weights together with an input word embedding. We apply our approach on models trained with classical DB approach, applying full fine-tune for SDv1.5 and LoRA [15] for SDXL. Results for other fine-tuning methods are

in Fig. 8. For comprehensive analysis, we train all models for a large number of 1000 steps, sampling every 5th step up to 50 steps and every 25th step thereafter, resulting in 48 operating points. For our approach, we designate models at steps 100 and 200 as edit models and include all models with lower steps, alongside the pre-trained model, as guidance models, totaling 28 operating points. This dense sampling allows us to study trade offs at different operating points. For each method, we generate 10 images per prompt at each operating point. The best operating point is automatically chosen based on F1 score between CLIP-T and DINO scores.

For completeness, we also compare to non-fine-tuning methods. Textual Inversion [9] is an inversion-based method that learns an input word embedding to represent a subject. BLIP-Diffusion [19] and IP-Adapter [46] are encoder-based methods offering zero-shot personalized generation. AnyDoor [5] is a reference-based method that tackles the more challenging problem of placing a specific subject into a specified background image. To evaluate AnyDoor on the DreamBooth benchmark, we used StableDiffusion to generate background images and CLIPSeg [23] for segmentation masks.

4.3. Qualitative evaluation

In Fig. 4, we present qualitative results. Across different subjects and prompts, DreamBlend is able to generate images preserving the layout of the underfit reference image as well as the identity of the subjects. In Fig. 5, we present a qualitative comparison with DreamBooth and Custom Diffusion on some challenging prompts. In each case, we manually select the best results for each method through visual inspection. Our method is able to generate images that achieve higher subject fidelity, prompt fidelity and diversity compared to the baselines. In Fig. 6, we present a comparison with non-fine-tuning methods, revealing that these approaches fall short in subject fidelity and photorealism, particularly for complex subjects. Additional results are in supplementary materials.

4.4. Quantitative evaluation

Tab. 1 shows the metrics averaged over all subjects and prompts of DB benchmark. Our approach achieves the best DINO, CLIP-I and CLIP-T scores, showing notable improvement in DINO and CLIP-T. As DINO is not trained to ignore differences between images that might have similar text descriptions, it is better at capturing subtle differences in subject fidelity, as also noted in DreamBooth [34].

In Fig. 7, we visualize how DreamBlend advances the pareto front for two example subjects, averaged over all 25 prompts. As each subject has a different fine-tuning trajectory, it is not possible to average metrics for different subjects after a certain number of fine-tuning steps and we



Figure 4. Cross Attention Guided Image Synthesis: Across various subjects and prompts, our approach successfully preserves the layout of the reference underfit image as well as the identity of the input subject. Images generated by the Overfit (Edit) and Underfit (Guidance) models used, are shown for reference.

present the result for each subject separately. Compared to all densely sampled operating points of both DB and CD, DreamBlend achieves better trade-offs, enabling better subject fidelity and prompt fidelity. Results for more subjects are in supplementary materials.

Method	Backbone	DINO	CLIP-I	CLIP-T
Real Images	-	0.774	0.885	-
Textual Inversion [9]	SD1.5	0.569	0.780	0.255
BLIP-Diffusion [19] (ZS)	SD1.5	0.594	0.779	0.300
DreamBooth [34]	SD1.5	0.659	0.805	0.296
CustomDiffusion [17]	SD1.5	0.627	0.789	0.286
DreamBlend (Ours)	SD1.5	0.675	0.808	0.308
AnyDoor [5]	SD2.1	0.588	0.753	0.289
IP-Adapter [46]	SDXL	0.584	0.786	0.303
DB-LoRA	SDXL	0.669	0.794	0.312
DreamBlend (Ours)	SDXL	0.678	0.795	0.313

Table 1. Quantitative evaluation on CLIP-I, CLIP-T and DINO

Study	Overall	Diversity
over DB	61.11 ($3e^{-12}$, [58.1, 64.1])	61.82 ($3e^{-09}$, [58.1, 65.6])
over CD	70.16 ($2e^{-20}$, [66.2, 74.1])	72.70 ($2e^{-35}$, [69.5, 75.9])

Table 2. Human preference study (in % of preference), with the Exact Binomial p-values and 95% Confidence Intervals.

4.5. Human preference study

We conducted human preference studies comparing our approach to DB and CD baselines. Two studies were performed, assessing overall preference and diversity, for each baseline. In the overall preference study, users chose between an image generated by our method and a baseline method for the same text prompt, considering both subject and prompt fidelity. In the diversity study, users selected the more diverse collection of four images between our method and a baseline. The results of human preference study are presented in Tab. 2, where our method is preferred over baselines. We perform one-sample binomial test and results are statistically significant with very low p-values and lower bound of 95% confidence intervals always greater than 50%. More details in supplementary materials.

5. Discussion

Generalization: Our approach is applicable to different fine-tuning methods like DreamBooth, Custom Diffusion and LoRA, different text-to-image diffusion models that feature a cross attention mechanism, as well as personalized editing of real images, as shown in Fig. 8.

Effect of cross attention guidance: An over-fitted model tends to generate concepts it has seen in input images of the subject, reducing prompt fidelity and diversity. Cross attention guidance encourages it to follow the layout of the reference underfit image instead. Fig. 9 shows the effect of varying cross attention guidance scale α in Algorithm 1.

Choice of guidance and edit models: Optimal performance in the guidance model is achieved by selecting an



Figure 5. Comparison with fine-tuning methods: Our approach successfully generates images with better subject fidelity, prompt fidelity and diversity on challenging prompts. In our results, autumn leaves in *backpack** images are more visible, *candle** is floating on water in all three images while maintaining subject fidelity and *stuffed animal** is on a pink fabric in all three images, exhibiting different poses.



Figure 6. Comparison to non-fine-tuning methods TI (Textual Inversion), BLIP-D (BLIP-Diffusion), IP-A (IP-Adapter) and AnyDoor.

underfit checkpoint with some subject resemblance, ensuring successful edits while maintaining prompt fidelity. Conversely, the edit model benefits from choosing a checkpoint that has learnt the subject without experiencing catastrophic attention collapse [3]. Fig. 10 presents these trade-offs quantitatively for an example subject *candle**.

Limitations: Like other text-to-image personalization methods, we share the challenge of inheriting failures when




Figure 7. Image alignment (DINO) - text alignment (CLIP-T) space spanned by densely sampled operating points of DreamBooth (gray), Custom Diffusion (red) and our method (green) for two example subjects. Our method advances the pareto front and enables generation of images closer to top right corner [1,1] of the image-text alignment space, inaccessible to existing methods.

the pre-trained model fails to generate an image with high prompt fidelity for a text prompt as shown in the top row of Fig. 11. Yet, in cases where the model initially succeeds but that knowledge is gradually lost in the fine-tuning, our approach combines benefits at different operating points for effective image synthesis. Similar to other fine-tuning based approaches, our success hinges on optimal selection of operating points. If the edit model is too over-fit, cross attention guidance becomes ineffective, as shown in Fig. 3. If the subject in the guidance image is too different from the actual subject, DreamBlend may fail to perform a successful



Backbone: SDXL. Prompt: a *vase** on top of a dirt road



Backbone: SD1.4. Prompt: a *dog** with a *bowl** on a beach



Backbone: SD1.5. Prompt: a *candle** on a cobblestone street




Personalized editing of real image: cat to *cat**

Figure 8. DreamBlend applied on different backbones, different fine-tuning techniques, real image editing. Top to bottom: SDXL fine-tuned with LoRA for *vase**, SDv1.4 fine-tuned with Custom Diffusion for multiple concepts *dog** and *bowl**, SDv1.5 fine-tuned with Custom Diffusion for *candle**, personalized editing of real image with DDIM inversion.



Figure 9. Effect of Cross Attention Guidance Scale α



(a) Different guidance models with step 200 edit model (b) Different edit models with step 25 guidance model

Figure 10. Effect of using different checkpoints as guidance and edit models for *candle**, numbers denote fine-tuning steps.

edit, as shown in bottom row of Fig. 11. By leveraging two checkpoints for inference, our method may appear to have higher storage requirements. However, given that current



a *teapot** on top of pink fabric



a *stuffed animal** in the jungle

Figure 11. Failure Cases. Top: SD and underfit model fail to generate images that follow the prompt. Bottom: Shape of the subject in underfit image is too different from the actual subject for a successful edit.

fine-tuning methods also necessitate storing checkpoints for sampling and selection, our storage needs are comparable in practice.

6. Conclusions

We presented DreamBlend, an approach that combines prompt fidelity and diversity from earlier checkpoints and subject fidelity from later checkpoints during image generation. It is straightforward and efficient, requiring only inference-time adjustments to existing techniques. It generates images with better subject fidelity, prompt fidelity and diversity, advancing the pareto front and surpassing state-of-the-art fine-tuning methods. Importantly, it successfully generates high fidelity images for challenging prompts, on which existing approaches struggled. For future work, we will explore utilizing this idea of a regularization on cross attention maps to combat over-fitting and reconstructability-editability trade-offs in other scenarios.

7. Societal impact

Fine-tuning text-to-image diffusion models has democratized the creation of personalized visuals such as pets, furniture, or self-portraits, making it accessible compared to training large models from scratch. Technologies like DreamBlend improve image fidelity across diverse contexts, enhancing creative potential in various fields. However, this advancement also brings associated risks, including concerns about copyright, privacy, authenticity, and the potential for fake media to propagate misinformation. To responsibly deploy these technologies in production settings, it is essential to implement safeguards and mitigation strategies, such as detecting AI-generated content. Moreover, disseminating knowledge about these technologies' inner workings is critical for promoting responsible innovation.

References

- [1] Yuval Alaluf, Elad Richardson, Gal Metzer, and Daniel Cohen-Or. A neural space-time representation for text-to-image personalization. *ACM Transactions on Graphics (TOG)*, 42(6):1–10, 2023. [2](#)
- [2] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. [5](#)
- [3] Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–10, 2023. [3](#)
- [4] Wenhu Chen, Hexiang Hu, Yandong Li, Nataniel Ruiz, Xuhui Jia, Ming-Wei Chang, and William W Cohen. Subject-driven text-to-image generation via apprenticeship learning. *Advances in Neural Information Processing Systems*, 36, 2024. [3](#), [5](#)
- [5] Xi Chen, Lianghua Huang, Yu Liu, Yujun Shen, Deli Zhao, and Hengshuang Zhao. Anydoor: Zero-shot object-level image customization. *arXiv preprint arXiv:2307.09481*, 2023. [5](#), [6](#)
- [6] Jooyoung Choi, Yunjey Choi, Yunji Kim, Junho Kim, and Sungroh Yoon. Custom-edit: Text-guided image editing with customized diffusion models. *arXiv preprint arXiv:2305.15779*, 2023. [3](#)
- [7] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-based semantic image editing with mask guidance. *arXiv preprint arXiv:2210.11427*, 2022. [3](#)
- [8] Dave Epstein, Allan Jabri, Ben Poole, Alexei Efros, and Aleksander Holynski. Diffusion self-guidance for controllable image generation. *Advances in Neural Information Processing Systems*, 36:16222–16239, 2023. [3](#)
- [9] 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](#), [5](#), [6](#)
- [10] Rinon Gal, Moab Arar, Yuval Atzmon, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Encoder-based domain tuning for fast personalization of text-to-image models. *ACM Transactions on Graphics (TOG)*, 42(4):1–13, 2023. [2](#), [3](#)
- [11] Jing Gu, Yilin Wang, Nanxuan Zhao, Tsu-Jui Fu, Wei Xiong, Qing Liu, Zhifei Zhang, He Zhang, Jianming Zhang, Hyun-Joon Jung, et al. Photoswap: Personalized subject swapping in images. *Advances in Neural Information Processing Systems*, 36, 2024. [3](#)
- [12] Shaozhe Hao, Kai Han, Shihao Zhao, and Kwan-Yee K Wong. Vico: Plug-and-play visual condition for personalized text-to-image generation. *arXiv preprint arXiv:2306.00971*, 2023. [2](#)
- [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. [3](#)
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. [3](#)
- [15] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. [5](#)
- [16] Xuhui Jia, Yang Zhao, Kelvin CK Chan, Yandong Li, Han Zhang, Boqing Gong, Tingbo Hou, Huisheng Wang, and Yu-Chuan Su. Taming encoder for zero fine-tuning image customization with text-to-image diffusion models. *arXiv preprint arXiv:2304.02642*, 2023. [2](#), [3](#)
- [17] 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. [2](#), [5](#), [6](#)
- [18] Jason Lee, Kyunghyun Cho, and Douwe Kiela. Countering language drift via visual grounding. *arXiv preprint arXiv:1909.04499*, 2019. [2](#)
- [19] Dongxu Li, Junnan Li, and Steven Hoi. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. *Advances in Neural Information Processing Systems*, 36, 2024. [2](#), [3](#), [5](#), [6](#)
- [20] Tianle Li, Max Ku, Cong Wei, and Wenhu Chen. Dreamedit: Subject-driven image editing. *arXiv preprint arXiv:2306.12624*, 2023. [3](#)
- [21] 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. [3](#)
- [22] Yuchen Lu, Soumye Singhal, Florian Strub, Aaron Courville, and Olivier Pietquin. Countering language drift with seeded iterated learning. In *International Conference on Machine Learning*, pages 6437–6447. PMLR, 2020. [2](#)
- [23] Timo Lüddecke and Alexander Ecker. Image segmentation using text and image prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7086–7096, 2022. [3](#), [5](#)
- [24] Yiyang Ma, Huan Yang, Wenjing Wang, Jianlong Fu, and Jiaying Liu. Unified multi-modal latent diffusion for joint subject and text conditional image generation. *arXiv preprint arXiv:2303.09319*, 2023. [2](#), [3](#)
- [25] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jianjun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. [3](#)
- [26] Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. In *Neural Information Processing Systems*, 2022. [2](#)
- [27] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6038–6047, 2023. [3](#)

- [28] Chong Mou, Xintao Wang, Jiechong Song, Ying Shan, and Jian Zhang. Dragondiffusion: Enabling drag-style manipulation on diffusion models. *arXiv preprint arXiv:2307.02421*, 2023. 3
- [29] Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, and Jun-Yan Zhu. Zero-shot image-to-image translation. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. 3
- [30] 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. 3, 5
- [31] 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
- [32] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. 2
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1, 2
- [34] 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, 5, 6
- [35] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Wei Wei, Tingbo Hou, Yael Pritch, Neal Wadhwa, Michael Rubinstein, and Kfir Aberman. Hyperdreambooth: Hypernetworks for fast personalization of text-to-image models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6527–6536, 2024. 2, 5
- [36] 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
- [37] Jing Shi, Wei Xiong, Zhe Lin, and Hyun Joon Jung. Instant-booth: Personalized text-to-image generation without test-time finetuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8543–8552, June 2024. 2, 3
- [38] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 3
- [39] Quan Sun, Qiyang Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yueze Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative pretraining in multimodality. *arXiv preprint arXiv:2307.05222*, 2023. 2
- [40] Yoad Tewel, Rinon Gal, Gal Chechik, and Yuval Atzmon. Key-locked rank one editing for text-to-image personalization. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. 2
- [41] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, Dhruv Nair, Sayak Paul, William Berman, Yiyi Xu, Steven Liu, and Thomas Wolf. Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers>, 2022. 3
- [42] Andrey Voynov, Qinghao Chu, Daniel Cohen-Or, and Kfir Aberman. p+: Extended textual conditioning in text-to-image generation. *arXiv preprint arXiv:2303.09522*, 2023. 2
- [43] Bram Wallace, Akash Gokul, and Nikhil Naik. Edict: Exact diffusion inversion via coupled transformations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22532–22541, 2023. 3
- [44] Qixun Wang, Xu Bai, Haofan Wang, Zekui Qin, and Anthony Chen. Instantid: Zero-shot identity-preserving generation in seconds. *arXiv preprint arXiv:2401.07519*, 2024. 2, 3
- [45] 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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15943–15953, 2023. 2
- [46] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023. 2, 3, 5, 6
- [47] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gungjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2(3):5, 2022. 2
- [48] Yanbing Zhang, Mengping Yang, Qin Zhou, and Zhe Wang. Attention calibration for disentangled text-to-image personalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4764–4774, 2024. 2, 3