
Customizable Image Synthesis with Multiple Subjects

– *Supplementary Material* –

Anonymous Author(s)

Affiliation

Address

email

1 **A Experiment details**

2 We supplement the experimental details of each method (Textual Inversion [1], DreamBooth [2],
3 Custom Diffusion [3], and Cones [4]) in this section. For better generation quality, we use Stable
4 Diffusion v2-1-base¹ as the pretrain model. For a fair comparison, we use 50 steps of DDIM [5]
5 sampler with a scale of 7.5 for all above methods. All experiments are conducted using one A-100
6 GPU.

7 **A.1 Textual Inversion**

8 We use the third-party implementation of huggingface [6] for Textual Inversion. We train each
9 subject-specific token with the recommended² batch size of 4 and a learning rate of 0.002 for 3000
10 steps. In particular, we initialize the subject-specific token with the corresponding class token. For
11 example, to customize a specific cat, we initialize the subject-specific token "<cat>" with the original
12 "cat" token.

13 **A.2 DreamBooth**

14 We use the third-party implementation of huggingface [6] for DreamBooth. Training is with a batch
15 size of 2, learning rate 5×10^{-5} , and training steps of $800 \times$ number of subjects.

16 **A.3 Custom Diffusion**

17 We use the official implementation³ for Custom Diffusion. Training is with a batch size of 2, learning
18 rate 1×10^{-5} and training steps of $250 \times$ number of subjects.

19 **A.4 Cones**

20 We use the official implementation for Cones. Training is with a batch size of 2, learning rate 4×10^{-5}
21 and training steps of $1200 \times$ number of subjects.

22 **A.5 Our Approach**

23 For our approach, We train each subject-specific residual token embedding with a batch size of 1 and
24 a learning rate of 1×10^{-6} for 3,000 steps. At inference time, the layouts are appointed by bounding
25 boxes given by the users to indicate the location of each subject. We use a positive value of +2.5
26 to strengthen the signal of the target subject and we use a negative value of -1×10^{-5} to weaken

¹<https://huggingface.co/stabilityai/stable-diffusion-2-1>

²https://github.com/rinongal/textual_inversion

³<https://github.com/adobe-research/custom-diffusion>

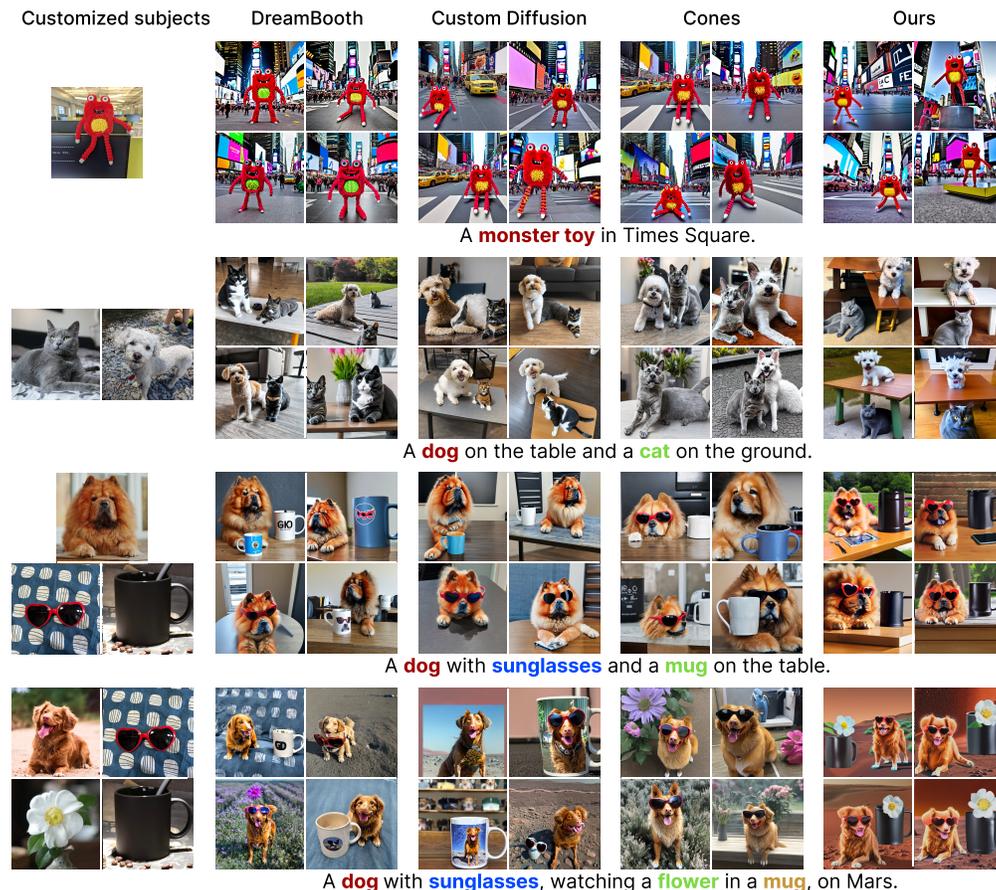
27 the signal of irrelevant subjects. Furthermore, we guide all 50 steps with the layout guidance in the
 28 whole generation process to get good customized generation results.

29 A.6 User Study

30 For two- to four-subject generation tasks, we design four different subject combinations for each
 31 task. This will yield 12 subject combinations in total. For each subject combination, we design
 32 four different text prompts to generate images with 5 random seeds. We conduct this procedure
 33 to all four methods. With such settings, each method generates 80 different images for each task.
 34 We give each generated image 4-8 questions for testing image alignment (2-4 questions) and text
 35 alignment (2-4 questions). The number of questions is proportional to the number of subjects used to
 36 customize the image (average 6 questions per generated image). Finally, we mess up the order of all
 37 the image-question pairs and assigned them to 25 different users for scoring, and finally summarized
 38 the results. In detail, every user needs to score $4 \times 4 \times 5 \times 6 = 480$ questions for each task and for
 39 each method.

40 B More comparisons

41 In this section, we conducted further comparisons between our approach and three other baselines.
 42 As shown in Fig. S1, regarding the generation of single subjects, the four methods exhibited similar
 43 performance. However, when dealing with semantically similar subjects, such as a dog and a
 44 cat, as well as scenarios involving three or more subjects, our approach clearly exhibited superior
 45 performance. Moreover, as shown in Fig. S2, we further showcase additional generated results,
 providing further evidence of the robustness of our approach.



46 Figure S1: More comparison of our approach and other baselines.

Table S1: Quantitative comparisons between our method (learning a residual token embedding) and learning a token embedding directly.

	Single Subject		Two Subjects		Three Subjects		Four Subjects	
	Text Alignment	Image Alignment						
Our	0.330	0.725	0.309	0.708	0.304	0.689	0.299	0.673
Token embedding	0.324	0.720	0.291	0.686	0.292	0.669	0.281	0.651

47 C More challenging cases

48 As shown in Fig. S3, we present a larger number of images generated by our approach, featuring a
 49 greater diversity of customized subjects. In comparison with other methods, we observe that when
 50 the number of customized subjects reaches four, the performance of other methods significantly
 51 deteriorates. In contrast, our approach can generate a larger number of customized subjects,
 52 exemplifying the superiority of our approach.

53 D Importance of residual token embedding

54 To demonstrate the superior generalization of the residuals, we conduct comparative experiments.
 55 As shown in Tab. S1, compared to directly updating the class embedding parameters in a single text
 56 embedding, our method, which involves updating the text encoder and calculating the average shift
 57 from the class to the specific subject based on a certain number of text templates, outperforms in both
 58 textual and visual similarity.

59 E Generated results of textual inversion

60 As shown in Fig. S4 We observe that Textual Inversion struggles with the generation of complex
 61 single subject and multiple subjects.

62 F Social impact and limitations

63 F.1 Social impact.

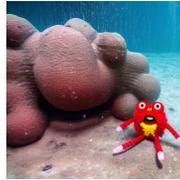
64 While training individual large-scale diffusion models remains prohibitively expensive, advancements
 65 in fine-tuning techniques have enabled individual users to customize their own models. Our
 66 technology empowers users to linearly combine their personalized single-subject models, generating
 67 high-quality images with multiple customized subjects, while maintaining significant advantages in
 68 terms of computation and storage efficiency. Furthermore, there is a growing need for more reliable
 69 detection techniques to identify and mitigate the presence of fake data.

70 F.2 Limitations.

71 Our method is limited by the inherent capabilities of the base model. Specifically, when it comes to
 72 combining more than six subjects, our method may not be able to consistently generate satisfactory
 73 results. In order to achieve the desired generation results, the provided layout by the user needs to be
 74 roughly consistent with the textual description.

Customized subjects

Ours



A **monster toy** under the water.



A **chair** under the mount fuji.



A **cat** and a **dog** on the beach.



A **teapot** and a **mug** on the grass.



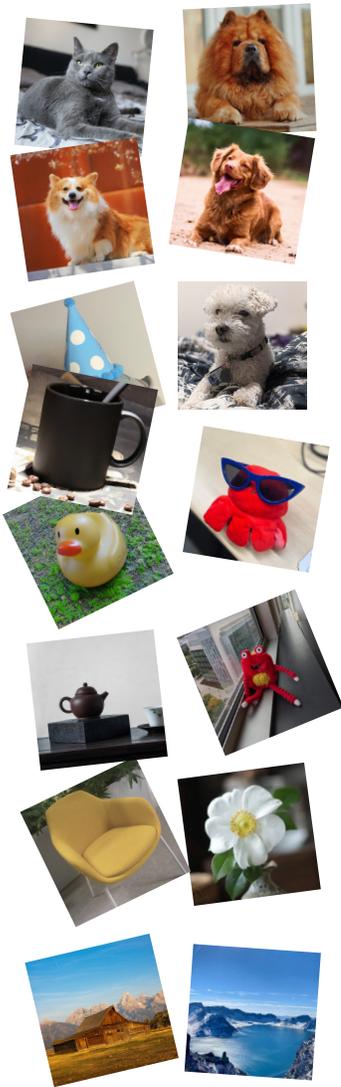
A **teapot** and a **mug** with a **flower** on the table.



A **dog** wearing **sunglasses**, sitting in a **mug**, on the table.

Figure S2: More results of our approach.

Customized subjects



A **cat** with **sunglasses**, sitting on a **chair**, with a **barn** in the background.



A **dog** with **sunglasses**, sitting in a **mug**, with a **lake** in the background.



A **dog** sitting next to a **chair** with **sunglasses** on it, with a **lake** in the background.



A **duck toy**, a **mug**, a **teapot**, and a **sunglasses** on the table.



A **dog** with **sunglasses**, watching a **flower** in a **mug**, on Mars.



A **dog** with **sunglasses** and a **hat**, sitting next to a **monster toy**, with a **lake** in the background.



A **monster toy**, a **dog**, and a **teapot**, with a **barn** in the background.



A **dog** with **sunglasses** and a **dog** with **sunglasses** on the grass.



A **dog** with **sunglasses** and a **hat**, sitting next to a **duck toy**, with a **lake** in the background.



A **dog** with **sunglasses** and a **hat**, sitting next to a **monster toy**, with a **lake** in the background.



A **duck toy**, a **teapot**, and a **mug**, with a **barn** in the background.



A **monster toy**, a **teapot**, and a **mug** with a **flower** in it, with a **barn** in the background.



A **monster toy** wearing **sunglasses**, a **hat**, sitting next to a **dog**, with a **barn** in the background.



A **dog** wearing **sunglasses**, a **hat**, sitting next to a **dog** wearing **sunglasses**, with a **barn** in the background.



A **dog** wearing **sunglasses**, a **hat**, sitting next to a **dog** wearing **sunglasses**, with a **lake** in the background.

Figure S3: More results of multi-subject generation.

Customized subjects



Textual Inversion



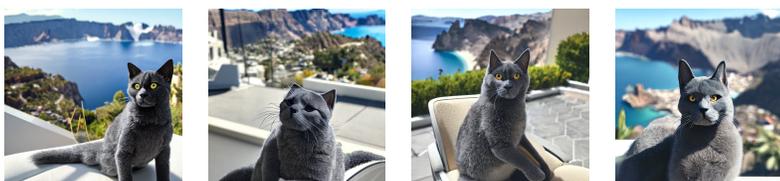
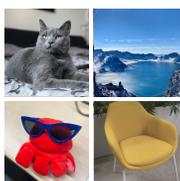
A **monster toy** in Times Square.



A **candle** and a **teapot** on the table.



A **monster toy** next to a **duck toy**, with a **barn** in the background.



A **cat** is wearing **sunglasses** and sitting on a **chair**, with a **lake** in the background.

Figure S4: Generated results of Textual Inversion.

75 **References**

- 76 [1] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or.
77 An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint*
78 *arXiv:2208.01618*, 2022.
- 79 [2] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman.
80 Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *arXiv preprint*
81 *arXiv:2208.12242*, 2022.
- 82 [3] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept
83 customization of text-to-image diffusion. *arXiv preprint arXiv:2212.04488*, 2022.
- 84 [4] Zhiheng Liu, Ruili Feng, Kai Zhu, Yifei Zhang, Kecheng Zheng, Yu Liu, Deli Zhao, Jingren Zhou, and
85 Yang Cao. Cones: Concept neurons in diffusion models for customized generation. *arXiv preprint*
86 *arXiv:2303.05125*, 2023.
- 87 [5] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Int. Conf. Learn.*
88 *Represent.*, 2021.
- 89 [6] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig
90 Davaadorj, and Thomas Wolf. Diffusers: State-of-the-art diffusion models. [https://github.com/
91 huggingface/diffusers](https://github.com/huggingface/diffusers), 2022.