

---

# Multiresolution Textual Inversion

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 We extend Textual Inversion to learn pseudo-words that represent a concept at  
2 different resolutions. This allows us to generate images that use the concept with  
3 different levels of detail and also to manipulate different resolutions using language.  
4 Once learned, the user can generate images at different levels of agreement to the  
5 original concept; “A photo of  $S^*(0)$ ” produces the exact object while the prompt  
6 “A photo of  $S^*(0.8)$ ” only matches the rough outlines and colors. Our framework  
7 allows us to generate images that use different resolutions of an image (e.g. details,  
8 textures, styles) as separate pseudo-words that can be composed in various ways.

## 9 1 Introduction

10 Textual Inversion [1] is a novel technique for introducing a new concept in a pre-trained text con-  
11 ditional generative model. Given a few images of a user-provided concept (e.g. an object or style),  
12 Textual inversion learns a new “word” in the embedding space to represent that object. Remarkably,  
13 these new *pseudo-words* can be composed in language to produce all kinds of creative compositions.

14 We extend Textual inversion to learn multiple pseudo-words that represent a concept at different  
15 resolutions. For example, in Fig. 1 the input concept is the re-creation, with various small objects,  
16 of Vermeer’s *Girl with a Pearl Earring*, by artist J. Perkins [2]. We learn this concept using 4  
17 re-croppings of the input. We then use a pre-trained text-to-image model to generate images using  
prompts that contain the learned pseudo-words that represent the concept at various resolutions.

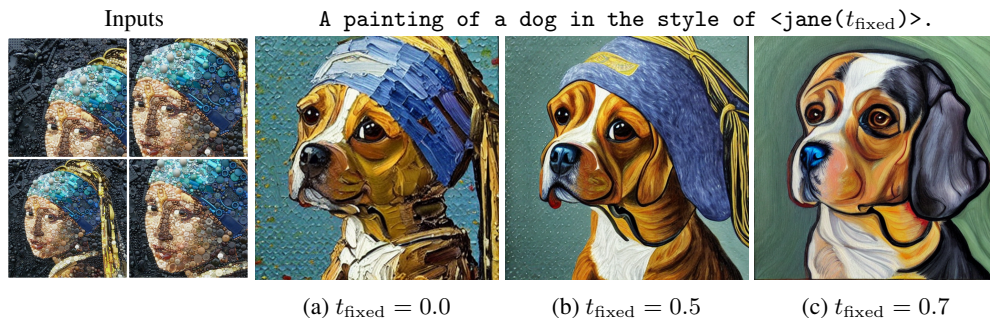


Figure 1: Multi-resolution textual inversion. We learn with our method the input concept as  $\langle \text{jane} \rangle$ . We then show images generated with the prompt: A painting of a dog in the style of  $\langle \text{jane}(t_{\text{fixed}}) \rangle$ , where  $t_{\text{fixed}}$  controls the resolution in which the pseudo-word captures the input concept. Each pseudo-word represents the concept with different levels of detail with the zero index capturing the full detailed learned concept.

19 **2 Method**

20 **Background.** Text-conditional diffusion models are trained to restore a clean image,  $x$ , based on a  
 21 corrupted observation,  $z_t(x)$ , and a text caption,  $y$ , describing the image [3, 4, 5, 6, 7]. The corruption,  
 22 which is typically additive noise, can be either in the image itself [3, 4] or in an encoding of the  
 23 image [5]. We denote with  $z_t(x)$  the diffused observation of image  $x$  at time  $t$ . Text-conditional  
 24 models are typically trained to minimize the objective:

$$J(\theta) = \mathbb{E}_{t \sim \mathcal{U}[0,1]} \mathbb{E}_{(x,y) \sim p} \mathbb{E}_{z_t \sim q_t(z_t|x,t)} \|\hat{\epsilon}_\theta(z_t(x), t, c_\theta(y)) - \epsilon(z_0, z_t)\|^2, \quad (1)$$

25 where  $p$  is the joint distribution of image-captions,  $q_t$  is the distribution of the diffused observations,  
 26  $y$  is the text-conditioning and  $\epsilon$  is the unscaled residual to  $z_0(x)$ .

27 The conditioning network,  $c_\theta(y)$ , maps the tokenized text  $y$  into a latent that is used to inform the  
 28 prediction of the denoiser network,  $\hat{\epsilon}_\theta$ . The first layer of  $c_\theta(y)$  is typically a lookup table,  $\text{emb}(y)$ ,  
 29 that maps tokens to vectors. We decompose the conditioning network as:  $c_\theta(y) = \text{enc}_\theta(\text{emb}_\theta(y))$ .

30 The authors of Textual Inversion [1] consider the task of finding an embedding,  $\text{emb}^*$  for a pseudo-  
 31 token  $y^*$  that represents a concept described by a small collection of images  $X = \{x_1, \dots, x_N\}$ . They  
 32 do so by solving the following optimization problem:

$$\min_{\text{emb}} \mathbb{E}_{t \in \mathcal{U}[0,1]} \mathbb{E}_{x \sim \mathcal{U}_X} \mathbb{E}_{z_t \sim q_t(z_t|x,t)} \|\hat{\epsilon}_\theta(z_t(x), t, \text{enc}_\theta(\text{emb})) - \epsilon(z_0, z_t)\|^2. \quad (2)$$

33 Intuitively, Textual Inversion is optimizing for an embedding that helps the model predict in the most  
 34 effective way the residual across all the images in the set and across all corruption levels.

35 **Multiresolution Textual Inversion.** The key idea of our method is that *the conditioning signal can*  
 36 *depend on the diffusion time (i.e. the noise level)*. For example, if the input image to the diffusion  
 37 model is a slightly noisy image of a cat, the text conditioning should give information on the details  
 38 of the cat (e.g. texture) and not necessarily on the image class which can easily be inferred from the  
 39 image input. Similarly, for very noisy images the conditioning should be capturing basic information  
 40 such as image class (e.g. cat vs dog) and colors. In general, we propose the following objective:

$$\min_{\{\text{emb}_0, \dots, \text{emb}_{T-1}\}} \mathbb{E}_{t \sim \mathcal{U}[0,1]} \mathbb{E}_{x \sim \mathcal{U}_X} \mathbb{E}_{z_t \sim q_t(z_t|x,t)} \|\hat{\epsilon}_\theta(z_t(x), t, \text{enc}_\theta(\text{emb}_{\lfloor t/T \rfloor})) - \epsilon(z_0, z_t)\|^2. \quad (3)$$

41 After training, we have learned a set of embeddings,  $\text{Emb}^* = \{\text{emb}_0, \dots, \text{emb}_{T-1}\}$ . Each of the  
 42 elements of the set, captures different levels of detail. The idea of having latents that describe the  
 43 image with different amounts of agreement to the input has been exploited to solve inverse problems  
 44 with GANs [8, 9, 10]. In these works, the resolution is controlled by the index of the GAN layer in  
 45 which we invert our images. In our method the agreement to the input is determined by the diffusion  
 46 time index and we can manipulate the different resolutions using language.

47 We can create embeddings indexed by continuous time by using linear interpolations of the elements  
 48 of the learned set. We present three different ways to use the learned embeddings to sample at  
 49 different levels of agreement to the concept appearing in the input images. For a visual comparison,  
 50 we show how these methods perform for a given image at different resolution levels, in Fig. 3.

51 **Fixed Resolution Sampling (Method 1).** The first method fixes the conditioning to a specific  
 52 embedding throughout the sampling. It is the most similar sampling method to Textual Inversion, in  
 53 the sense that it is using a fixed embedding for all diffusion levels. The user can control the resolution  
 54 by picking one embedding from the set  $\text{Emb}^*$ . This method allows the user to visualize what is  
 55 learned at each resolution. As we explained earlier, we expect that embeddings that correspond to  
 56 time close to  $t = 0$ , should learn details (e.g. texture) since those are more informative to denoise a  
 57 slightly noisy image than image class for example. Embeddings closer to  $t = 1$  should be related to  
 58 more coarse information about the image, e.g. what is the object, what are the colors, etc.

59 For the next two sampling methods, we change the conditioning based on the sampling time. Assume  
 60 that we want to generate an image at resolution  $t = t_{\text{fixed}}$ . We propose two sampling methods:

61 **Semi Resolution-Dependent Sampling (Method 2).** Semi Resolution-Dependent Sampling uses  
 62  $\text{emb}_t$  when the sampling time is  $t$  if  $t \geq t_{\text{fixed}}$  and no conditioning otherwise. Essentially, this forces  
 63 the model to generate images that match (in a distributional sense) the input images at noise level  $t$ .  
 64 This method is particularly useful for style-transfer and creative prompts – for some percentage of  
 65 the sampling procedure the model performs unconditional generation (with a starting point a diffused  
 66 sample that matches the original concept with noise). This idea extends the novel SDEdit paper [11]  
 67 in the sense that it allows to do guided image synthesis (but this time starting from language tokens).

68 **Fully Resolution-Dependent Sampling (Method 3).** Fully Resolution-Dependent Sampling is  
 69 similar to Semi Resolution-Dependent Sampling, but instead of doing unconditional generation for  
 70  $t < t_{\text{fixed}}$ , it fixes the embedding to  $\text{emb}_{\text{fixed}}$ . This still allows for variations from the given concept  
 71 (since  $\text{emb}_{\text{fixed}}$  only resembles it in a given resolution), but the variations are more controlled since  
 72 there is no sampling period of unconditional generation.

73 **Creating a textual interface.** Similar to Textual Inversion, we create pseudo-words that allow us  
 74 to use language to guide the image generation process with our learned embeddings. The difference  
 75 is that the embeddings for our pseudo-words might be changing based on the sampling time. We use  
 76  $S^*|t_{\text{fixed}}|$ ,  $S^*(t_{\text{fixed}})$ ,  $S^*[t_{\text{fixed}}]$  to denote tokens that are used with Fixed Resolution, Semi Resolution  
 77 and Fully Resolution Dependent Sampling respectively.

### 78 3 Experiments




Figure 2: Compositions of learned pseudo-words across resolutions. Observe that the dog (b) and cat (c) are made from small plastic objects, i.e. using the detailed structure of  $\langle \text{jane} \rangle$ , but in the shape of a dog or the toy cat. In (d) the color stripes are obtained from the cat detail.

79 We begin by comparing the different sampling methods. Results are shown in Fig. 3. Each row  
 80 shows a different sampling algorithm each each column a different resolution. With Fixed Resolution  
 81 Sampling (row 1), we can selectively extract details of the image, e.g.  $S^*_{0,0}$  extracts only the buttons.  
 82 Semi Resolution-Dependent Sampling (row 2) allows for larger variations from the input for high  
 83 values of  $t_{\text{fixed}}$ , e.g.  $S^*(0.8)$  gives a photorealistic image of person that only roughly preserves colors  
 84 and pose from the input concept. Finally, Fully Resolution-Dependent Sampling (row 3) maintains  
 85 higher agreement to the input concept along all resolutions.

86 We then show that Multiresolution Textual Inversion performs on par (or even better) with Textual  
 87 Inversion. Fig. 4 shows that our method can use the learned concepts in combinations with prompts  
 88 of all sorts – we use the prompt and images from the Textual Inversion paper.

89 Finally, Fig. 2 shows that the learned pseudo-words can be combined in arbitrary ways across  
 90 different resolutions and concepts. For example, we can extract the plastic object detailed structure of  
 91 concept  $\langle \text{jane} \rangle$  and use it to generate other objects, such as a dog painting and the cat of Fig. 2a.

### 92 4 Conclusions and Future Work

93 We showed how to learn multiple pseudo-words representing an input concept at different scales.  
 94 This expands the prompt vocabulary and can be easily used with pre-trained text conditional diffusion  
 95 models. As future work, we would like to quantitatively measure if our method outperforms Textual  
 96 Inversion and alternative methods, e.g. DreamBooth [12], for inverse problems such as inpainting.  
 97 It is also worth exploring if models trained with general diffusion processes beyond additive noise  
 98 [13, 14] can extract different aspects of scale, depending on how they corrupt their inputs.




Figure 3: Comparison of different sampling methods that be used for Multiresolution Textual Inversion. Each row shows a different sampling algorithm and each column a different resolution. With Fixed Resolution Sampling (row 1), we can selectively extract details of the image, e.g.  $S^*[0.0]$  extracts only the buttons. Semi Resolution-Dependent Sampling (row 2) allows for larger variations from the input for high values of  $t_{\text{fixed}}$ . Fully Resolution-Dependent Sampling (row 3) maintains higher agreement to the inputs along all resolutions.

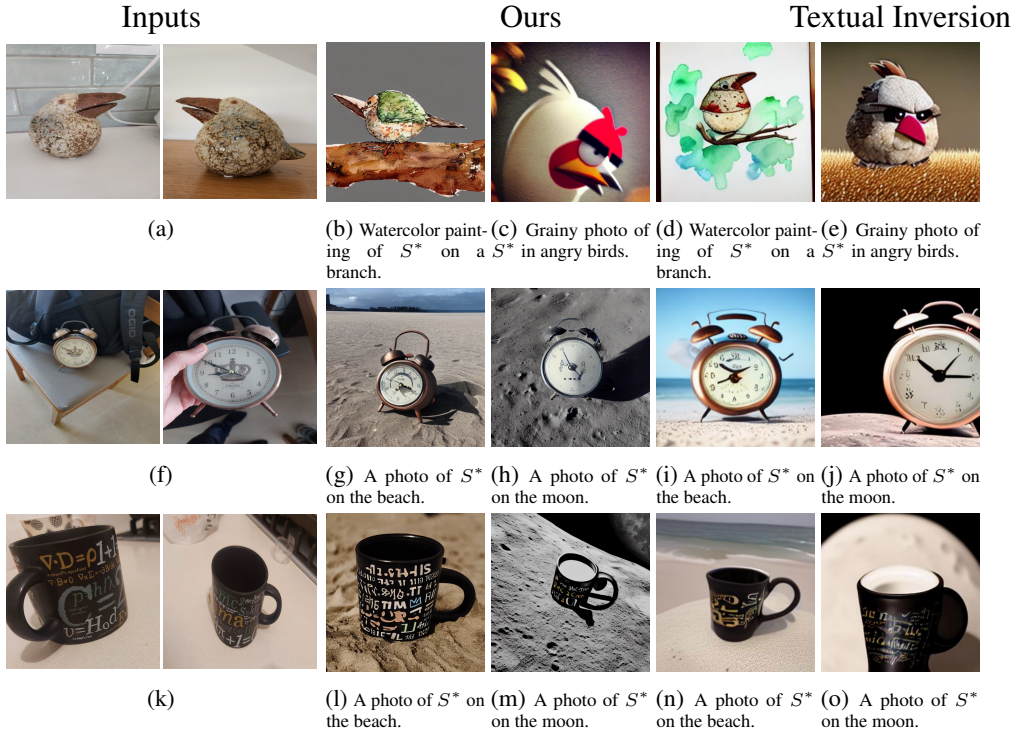


Figure 4: Personalized image generation. We show that with our method we can use the pseudo-words for the learned concept to create personalized images as if it was a normal word token. Our method performs on par and sometimes better than Textual Inversion on this task.

99 **References**

- 100 [1] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and  
101 Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using  
102 textual inversion. *arXiv preprint arXiv:2208.01618*, 2022.
- 103 [2] Jane Perkins, an artist in found things, Homepage: <https://janeperkins.co.uk/> Related article:  
104 <https://mandalaoftheday.com/2015/03/05/64-buttons-and-things-mandala/> .
- 105 [3] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical  
106 text-conditional image generation with clip latents, 2022.
- 107 [4] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed  
108 Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim  
109 Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image  
110 diffusion models with deep language understanding, 2022.
- 111 [5] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer.  
112 High-resolution image synthesis with latent diffusion models, 2021.
- 113 [6] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020.
- 114 [7] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and  
115 Ben Poole. Score-based generative modeling through stochastic differential equations, 2020.
- 116 [8] Giannis Daras, Joseph Dean, Ajil Jalal, and Alexandros G. Dimakis. Intermediate layer  
117 optimization for inverse problems using deep generative models. In *ICML*, 2021.
- 118 [9] Giannis Daras, Yuval Dagan, Alexandros G. Dimakis, and Constantinos Daskalakis. Score-  
119 guided intermediate layer optimization: Fast langevin mixing for inverse problems, 2022.
- 120 [10] Niklas Smedemark-Margulies, Jung Yeon Park, Max Daniels, Rose Yu, Jan-Willem van de  
121 Meent, and Paul Hand. Generator surgery for compressed sensing, 2021.
- 122 [11] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano  
123 Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations, 2021.
- 124 [12] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman.  
125 Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation, 2022.
- 126 [13] Arpit Bansal, Eitan Borgnia, Hong-Min Chu, Jie S Li, Hamid Kazemi, Furong Huang, Micah  
127 Goldblum, Jonas Geiping, and Tom Goldstein. Cold diffusion: Inverting arbitrary image  
128 transforms without noise. *arXiv preprint arXiv:2208.09392*, 2022.
- 129 [14] Giannis Daras, Mauricio Delbracio, Hossein Talebi, Alexandros G. Dimakis, and Peyman  
130 Milanfar. Soft diffusion: Score matching for general corruptions, 2022.