



# A Text Based Image Generator Using Artificial Intelligence

<sup>1</sup>Mrs. Janaki Ramal, <sup>2</sup>Karthikeyan S, <sup>3</sup>Kumaresh M, <sup>4</sup>Vinu Vimal Padua V

<sup>1</sup>Assistant Professor, Department of Computer Science and Engineering, Jeppiaar Institute of Technology, Chennai, India

<sup>1</sup>janakiramal@jeppiaarinstitute.org

<sup>2,3,4</sup>UG Scholar, Department of Computer Science and Engineering, Jeppiaar Institute of Technology, Chennai, India

<sup>2</sup>keyan4356@gmail.com, <sup>3</sup>kumareshmusk@gmail.com, <sup>4</sup>vinuvimal@gmail.com

## Abstract

Despite unprecedented ability in imaginary creation, large text-to-image models are further expected to express customized concepts. Existing works generally learn such concepts in an optimization-based manner, yet bringing excessive computation or memory burden. In this paper, we instead propose a learning-based encoder for fast and accurate concept customization, which consists of global and local mapping networks. In specific, the global mapping network separately projects the hierarchical features of a given image into multiple “new” words in the textual word embedding space, i.e., one primary word for well-editable concept and other auxiliary words to exclude irrelevant disturbances (e.g., background). In the meantime, a local mapping network injects the encoded patch features into cross attention layers to provide omitted details, without sacrificing the editability of primary concepts. We compare our method with prior optimization-based approaches on a variety of user-defined concepts, and demonstrate that our method enables more high-fidelity inversion and robust editability with a significantly faster encoding process. Our code will be publicly available at <https://github.com/kumaresh-rgb>.

## 1. INTRODUCTION

Recently, large-scale diffusion models [3, 25, 28, 30] have demonstrated impressive superiority in text-to-image generation. By training with billions of image-text pairs, large text-to-image diffusion models have exhibited excellent semantic

understanding ability, and generate diverse and photo-realistic images being accordant to given text prompts. Owing to their unprecedentedly creative capabilities, they have been applied to various tasks, such as image editing [14, 19], data augmentation [20], and even artistic creation [22]. However, despite diverse and general generation, users may expect to create imaginary instantiations with undescribable personalized concepts [15], e.g., “corgi” in Fig. 1. To



Figure 1. Given an input image, customized text-to-image generation learns a pseudo-word ( $S^*$ ) in word embedding space to represent the target concept. With  $S^*$ , one can synthesize or edit the concept with text prompts flexibly. The time to learn a new concept is listed at the bottom, and our method learns the new concept faster. This end, many recent studies have been given to customized text-to-image generation [13, 15, 29], which aims to learn a specific concept from a small set of user-provided images (e.g., 3~5 images). Then, users can flexibly compose the learned concepts into new scenes, e.g., A  $S^*$  wearing

sunglasses in Fig. 1. Given a small image set depicting the target concept, Textual Inversion [13] learned a new pseudo-word (i.e.,  $S^*$ ) in the well-editable embedding space of text encoder to represent the user-defined concept, yet failed to capture image-specific details. DreamBooth [29] finetuned the entire diffusion model to accurately align the target concept with a unique identifier. Custom Diffusion [15] balanced the fidelity and memory by selectively finetuning K, V mapping parameters in cross attention layers. Since text prompts are usually not sufficient to precisely describe a specific object [13], customized generation is more practical for specific concept synthesis and editing. Albeit flexible generation has been achieved by [13, 15, 29], the computational efficiency remains a challenging issue to obtain the textual embedding of a visual concept. Existing methods usually adopt the per-concept optimization formulation, which requires several or tens of minutes to learn a single concept. From Fig. 1, the most efficient Custom Diffusion [15] still takes around 6 minutes to learn one concept, which is infeasible for online applications. In contrast, in the field of GAN inversion, efficient learning-based methods [27] have been proposed to accelerate the optimization process. They trained an encoder to infer the latent code directly, which only needs one step forward inference. Driven by the above analysis, we propose a learning-based encoder for Encoding visual concepts into Textual Embeddings, termed as ELITE. As shown in Fig. 2, our ELITE adopts a pre-trained CLIP image encoder [24] for feature extraction, followed by a global mapping network and a local mapping network to encode visual concepts into textual embeddings. Firstly, we train a global mapping network to map the CLIP image features into the textual word embedding space of the CLIP text encoder, which is analogous to [13]. Since a given image contains both the subject and irrelevant disturbances, encoding them as a single word embedding severely degrades the editability of subject concept. Thus, we propose to separately learn them with a well-editable primary word and several auxiliary words. Using the hierarchical features from CLIP intermediate layers, the word learned from the deepest features naturally links to the primary concept (i.e., the subject), while auxiliary words learned from other features describe the irrelevant disturbances (e.g., background). Following [15], the parameters of K, V projection in cross attention layers are finetuned during global training for a better inversion. When deploying to customized generation, we only use the primary word to avoid editability degradation from auxiliary words. Usually, a visual concept is

worth more than one word, and describing it with a single word may result in the inconsistency of local details [13]. For higher fidelity of the learned concept without sacrificing its editability, we further propose a local mapping network to inject finer details. From Fig. 2, the local mapping network encodes the CLIP features into the textual feature space (i.e., the output space of the text encoder), while keeping their spatial structure. The obtained textual feature embeddings are injected through additional cross attention layers, and the output feature is summed with the global part to improve the local details. Experiments demonstrate that our method can encode the target concept efficiently and faithfully, while keeping control and editing abilities. The contributions of this work can be summarized as: • We propose a learning-based encoder ELITE for fast and accurate customized text-to-image generation. It adopts a global mapping and a local mapping to directly encode the visual concept into textual embeddings. • Multi-layer features are adopted in global mapping network to learn a well-editable primary word embedding, while the local mapping network improves the consistency of details significantly.

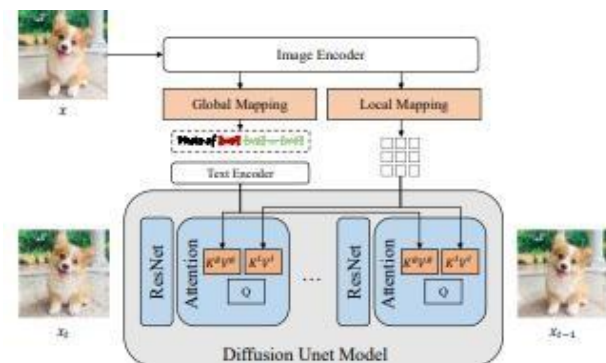


Figure 2. Inference pipeline of our proposed ELITE. Given a user-provided image  $x$ , our ELITE extracts hierarchical features with CLIP image encoder. Then, it uses global and local mapping networks to encode a visual concept into textual word embeddings (i.e.,  $w_0$  and  $w_1 \dots N$ ) and textual feature embeddings, respectively. The embeddings are injected into generation by cross attention. Note that, when deployed into customized generation, only the primary word  $w_0$  is used for better editability. • Experiments show that our encoder can faithfully recover the target image with higher visual fidelity, and also enable more robust editing.

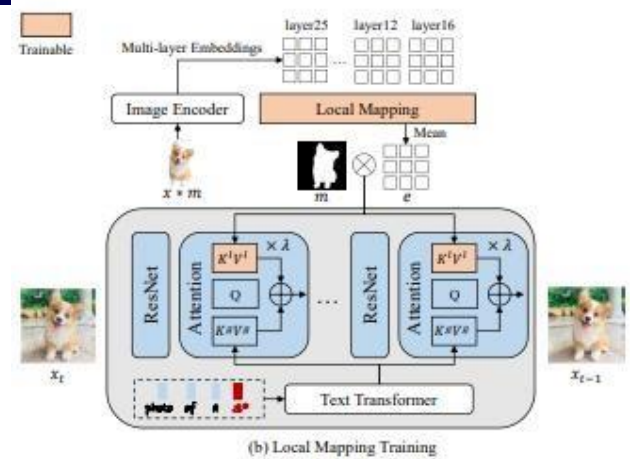
## 2. RELATED WORK

### 2.1. Text-to-Image Generation

Deep generative models have achieved tremendous success on text-conditioned image generation [3, 6, 9, 10, 12, 17, 21, 25, 26, 28, 30,



31, 34, 35] and attracted intensive recent attention. They can be categorized into three groups: GAN-based, VAE-based, and diffusion-based models. Albeit GAN-based [17, 31, 34] and VAE-based models [6, 9, 12,26,35] synthesize images with promising quality and diversity, they still cannot match user descriptions very well. Diffusion-based models demonstrate more unprecedentedly high-quality and controllable imaginary generation and are broadly applied in text-to-image generation [3, 10, 21, 25, 28, 30]. GLIDE [21] firstly introduces diffusion models into text-to-image generation with classifier-free guidance. DALLE-2 [25], Imagen [30], and LDM [28] employ pretrained large-scale text encoders to provide more controllable guidance signals. For a trade-off between efficiency and high resolution, they either introduce super-resolution models [25,30], or encode images into the low-dimensional latent space [28]. Despite their superior performance on general synthesis, they still struggle to express some specific or user-defined concepts [13, 15, 29], e.g., “corgi” in Fig. 1. Our method focuses on making pre-trained diffusion models to learn these new concepts efficiently.



## 2.2. GAN Inversion

GAN inversion refers to projecting real images into latent codes, such that images can be faithfully reconstructed and edited with pre-trained GAN models. Generally, there are mainly two types of GAN inversion algorithms in the literature: (i) optimization-based: directly optimize latent code to minimize the reconstruction error [5, 8, 18], and (ii) encoder-based: train an encoder to invert an image into latent space [1, 2, 23, 27, 33]. The optimization-based methods [5,8,18] usually require hundreds of iterations to obtain promising results, while encoder-based methods greatly accelerate this process via one feed-forward pass only. To improve image fidelity without compromising editability, [33] embeds the omitted high-frequency information into latent codes. Inspired by this, our ELITE also proposes a local mapping network that encodes the concept images into textual feature space to improve details consistency.

## 2.3. Diffusion-based Inversion

For text-to-image diffusion models, their inversion can be performed in two types of latent spaces: (i) Textual Word Embedding (TWE) space [11,13,15,29] or (ii) Imagebased Noise Map (INM) space [7, 19, 32]. In this work, our method performs inversion in TWE space, which has been widely used to invert customized concepts. From the perspective of GAN inversion, existing inversion methods in the TWE space follow the optimization-based formulation. Textual inversion [13] and DreamArtist [11] directly optimize the embedding of new “words” using a few userprovided images, yet failing to recover the target image (i.e., low-fidelity). DreamBooth [29] finetunes the whole pre-trained text-to-image model to learn high-fidelity new concepts with a unique identifier. To enable fast tuning, Custom Diffusion [15] only updates the key and value mapping parameters in the cross-

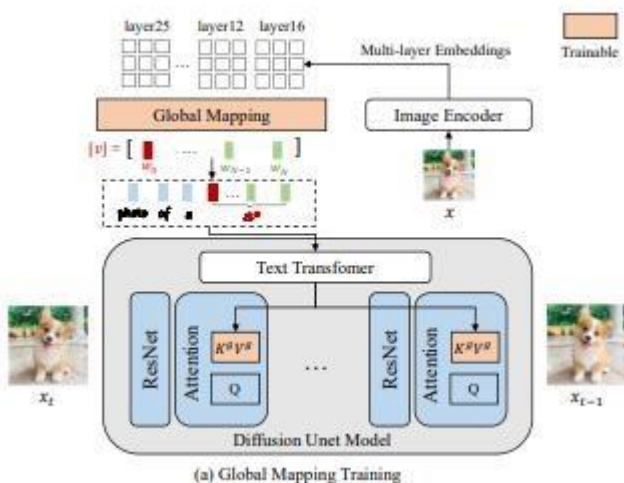


Figure 3. Training pipelines of our proposed ELITE. It consists of two stages: (a) a global mapping network is first trained to encode a concept image into multiple textual word embeddings, where one primary word ( $w_0$ ) for well-editable concept and other auxiliary words ( $w_1 \dots N$ ) to exclude irrelevant disturbances. (b) To address the information loss during global mapping, a local mapping network is further trained, which projects the foreground object into textual feature space to provide local details. In this stage, only the well-editable primary word ( $w_0$ ) is used.

attention layers with better performance. Albeit these optimization-based methods have obtained high fidelity and editability, their learning still requires many iterations and multiple user-provided images. In contrast, our method follows the encoder-based formulation to accelerate the process. Meanwhile, using global and local mapping, our method faithfully learns a new concept with one single image. Overall, our two core designs significantly improve efficiency in terms of time and memory.

### 3. METHODOLOGY

Given a pretrained text-to-image model  $\theta$  and an image  $x$  indicates the target concept (usually an object), customized text-to-image generation aims to learn a pseudoword ( $S^*$ ) in word embedding space to describe the concept faithfully, while keeping the editability. To achieve fast and accurate customized text-to-image generation, we propose a method ELITE to encode the visual concept into textual embeddings. As illustrated in Fig. 2, our ELITE adopts a global mapping network to encode visual concepts into the textual word embedding space. The obtained word embeddings can be composed with texts for customized generation. To address the information loss in word embedding, we further proposed a local mapping network to encode the visual concept into textual feature space to improve the consistency of the details. In the following, we begin by presenting an overview of the text-to-image model utilized in our approach (Sec. 3.1). Then, we will introduce the details of the proposed global mapping network (Sec. 3.2) and local mapping network (Sec. 3.3).

#### 3.1. Preliminary

In our experiments, we employ the Stable Diffusion [28] as our text-to-image model, which is trained on large-scale data and consists of two components. First, the autoencoder ( $E(\cdot), D(\cdot)$ ) is trained to map an image  $x$  to a lower dimensional latent space by the encoder  $z = E(x)$ . While decoder  $D(\cdot)$  learns to map the latent code back to the image, such that  $D(E(x)) \approx x$ . Then, the conditional diffusion model  $\theta(\cdot)$  is trained on the latent space to generate latent codes based on text condition  $y$ . To train the diffusion model, simple mean-squared loss is adopted,  $LLDM := E_{z \sim E(x), y \sim N(0,1), t} \| \theta(z_t, t, \tau\theta(y)) - z_t \|^2$ , (1) where  $\tau\theta(\cdot)$  denotes the unscaled noise,  $t$  is the time step,  $z_t$  is the latent noised to time  $t$ , and  $\tau\theta(\cdot)$  represents the pretrained CLIP text encoder [24]. During inference, a random Gaussian noise  $z_T$  is iteratively denoised to  $z_0$ , and the final image is obtained through the decoder  $x_0 = D(z_0)$ . To incorporate text information in the process of image synthesis, cross attention is adopted in Stable Diffusion. Specifically, the latent image

feature  $f$  and text feature  $\tau\theta(y)$  are first transformed by the projection layers to obtain the query  $Q = WQ \cdot f$ , key  $K = WK \cdot \tau\theta(y)$  and value  $V = WV \cdot \tau\theta(y)$ .  $WQ, WK$ , and  $WV$  are weight parameters of query, key, and value projection layers, respectively. Attention is conducted by a weighted sum over value features,  $Attention(Q, K, V) = \text{Softmax} \left( \frac{QK^T}{\sqrt{d}} \right) V$ , (2) where  $d$  is the output dimension of key and query features. The latent image feature is then updated with the attention block output.

#### 3.2. Global Mapping

Following [13, 15], we choose the textual word embedding space of CLIP text encoder as the target for inversion, which has been demonstrated with great editing capacity. To improve the computation efficiency, we propose a global mapping network that encodes the given concept image into word embeddings directly. As illustrated in Fig. 3(a), to facilitate the embedding learning, the pretrained CLIP image encoder  $\psi\theta(\cdot)$  is adopted as feature extractor, and our global mapping network  $Mg(\cdot)$  projects the CLIP features as word embeddings  $v, v = Mg \circ \psi\theta(x)$ . (3) where  $v \in \mathbb{R}^{N \times d}$ ,  $N$  is the number of words and  $d$  is the dimension of word embedding. Global average pooling is employed to obtain the final word embedding.

Since the image  $x$  contains both desired object concept and other irrelevant concepts (e.g., background), encoding them into one word

(i.e.,  $N = 1$ ) results in an entangled word embedding  $v$  with poor editability. To obtain a more informative and editable word embedding, we adopt a multi-layer approach to learn  $N$  words from the image  $x$  separately ( $N > 1$ ). Specifically, we select  $N$  layers from CLIP image encoder, and each layer  $\psi_{L_i} \theta(\cdot)$  learns one word  $w_i$  independently. All words  $[w_0, \dots, w_N]$  are concatenated together to form the textual word embedding  $v$ . Following [13], we introduce a pseudo-word  $S^*$  to represent the new concept, and associate its embedding with  $v$ . To train our global mapping, we adopt Eqn. (1), and regularize the obtained word embeddings,  $L_{global} = LLDM + \lambda_{global} \|v\|_1$ , (4) where  $\lambda_{global}$  is a trade-off hyperparameter. To condition the new concept on generation, we randomly sample a text from the CLIP ImageNet templates [24] as text input, such as a photo of a  $S^*$ . The full template list is provided in the Suppl. Besides, following [15], the key and value projection layers in cross attention are finetuned during training, and the obtained new projections are denoted as  $K_g = W_g K \cdot \tau\theta(y)$  and  $V_g = W_g V \cdot \tau\theta(y)$ . Benefiting from the hierarchical semantics learned by different layers in the CLIP image encoder, the feature from the deepest layer (i.e., layer 25) possesses the highest comprehension of the image. And the word embedding associated



with the deepest feature is learned to describe the primary concept (i.e., the object), while keeping great editability. In contrast, word embeddings from shallower features are learned to describe the irrelevant disturbances, such as background (see Sec. 4.2 for more details). Note that, we use only the word embedding of the deepest feature during local training (Sec. 3.3) and image generation stages for better editing capacity.

### 3.3. Local Mapping

Usually, a single word embedding is not sufficient to faithfully describe the details of a given concept, while multiple word embeddings suffer from degraded editing capacity. To improve the consistency between the given concept and synthesized image without sacrificing editability, we further propose a local mapping network. As shown in Fig. 3(b), the local mapping network  $MI(\cdot)$  encodes the multi-layer CLIP features into the textual feature space (i.e., the output space of text encoder),  $e = MI \circ \psi \theta(x * m)$ . (5) where  $m$  is the object mask to ease the redundant details of background.  $e \in \mathbb{R}^{p \times p \times d}$  keeps the spatial structure and  $p$  is the feature size. Each pixel of  $e$  contains the details information of a local patch in given images. To inject the local Figure 4. Visual comparisons of concept generation. For generating images, A photo of a  $S^*$  is used for Textual Inversion [13] and ours, while A photo of a  $S^*$  [category] is used for Dreambooth [29] and Custom Diffusion [15]. Our ELITE is comparable with competing methods.

Figure 6. Visual comparisons of different variants. Left:  $[v]$  denotes the generation results with full word embeddings  $v$ , while  $[w]$  denotes the generation results with the primary word embedding  $w$ . Either learning single word embedding or multiple word embeddings from the single layer (i.e. the deepest layer) fails to achieve reliable editing capacity. In contrast, our proposed method of learning multiple words from multiple layers successfully learns an editable primary word embedding. Right: Our proposed local mapping improves the consistency of details between the input image and the generated image significantly, while maintaining editability.


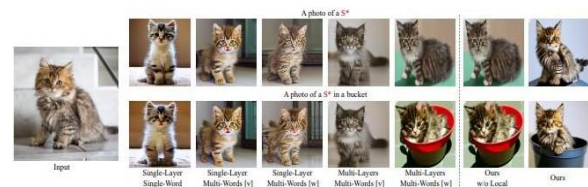


Figure 5. Visualization of learned word embeddings. The word associated deepest feature (i.e.,  $w_0$  from layer 25) describes the primary concept (i.e., corgi, teddybear), while other words describe irrelevant disturbances. information of  $e$  into generation, additional cross attention is introduced. Specifically, we add two additional projection layers  $W_I K$  and  $W_I V$  into cross attention module, and the local attended feature is obtained by  $Attention(Q, K_I, V_I)$  where  $K_I = W_I K \cdot (e * m)$  and  $V_I = W_I V \cdot (e * m)$ . Then the local attended features are then fused with global attended features through,  $Out = Attention(Q, K_g, V_g) + \lambda Attention(Q, K_I, V_I)$ , (6) where  $\lambda$  is a hyperparameter and set as 1 during training. To focus on the object region, the obtained attention map  $Q K_I T$  is reweighted by  $Q K_g T_i$ , where  $i$  is the index of  $w_0$  in the text prompt. To train the local mapping network, we also adopt the Eqn. (1), while regularizing the local values  $V_I$ ,  $L_{local} = L_{LDM} + \lambda_{local} ||V_I||_k$ , (7) where  $\lambda_{local}$  is a trade-off hyperparameter.






Figure 7. Visual comparisons on the effect of value of  $\lambda$ . As  $\lambda$  increases, the consistency of details between the generated image and input image improves, but at the cost of reduced editability.

Table 1. Ablation study. [v] denotes the testing results are generated with full word embeddings  $v$ , while [w] denotes testing results are generated with the primary word embedding  $w$ .

Method	Text-alignment (↑)	Image-alignment (↑)	KID (↓)
Single-Layer Single-Word	0.194	0.802	6.62
Single-Layer Multi-Words [v]	0.211	0.741	27.98
Single-Layer Multi-Words [w]	0.228	0.758	37.51
Multi-Layers Multi-Words [v]	0.198	<b>0.850</b>	6.54
Multi-Layer Multi-Words [w]	<b>0.266</b>	0.733	13.68
Ours w/o Local	0.266	0.733	13.68
Ours	<b>0.266</b>	0.804	<b>5.27</b>

## 4. EXPERIMENTS

### 4.1. Experimental Settings

**Datasets.** To train our local and global mapping networks, we use the testset of OpenImages [16] as our training dataset. It contains 125k images with 600 object classes. During training, we crop and resize the object image to 512x512 according to the bounding box annotations. While for local mapping training, mask annotations are also used to extract foreground objects. During inference, we adopt images from existing work [13, 15, 29] with various categories, such as dog, cat, and toy, etc.

**Evaluation metrics.** Following [15], we evaluate our method with the Image-alignment, Text-alignment, and KID. For Image-alignment, we calculate the CLIP visual similarity between the generated images and the target concept image. For Text-alignment, we calculate the CLIP text-image similarity between the generated images and given text prompts. The pseudo-word ( $S^*$ ) in the text prompt is replaced with the proper object category for extracting CLIP text feature. For KID, we select six evaluation datasets that each contain multiple images. For each dataset, we take only one image as the concept image, and calculate the KID [4] between generated images and all real images. Moreover, we also adopt the optimization time as a metric.

Table 2. Quantitative comparisons with existing methods.

Method	Text-alignment (↑)	Image-alignment (↑)	KID (↓)	Time (↓)
Textual Inversion [13]	0.183	0.689	51.22	50min
DreamBooth [29]	0.249	0.827	4.48	30min
Custom Diffusion [15]	0.231	0.868	14.32	6min
Ours	0.266	0.804	5.27	0.05 s

Table 3. User study. The numbers indicate the percentage (%) of volunteers who favor the results of our method over those of the competing methods based on the given question. Our method is preferred over competing methods.

Metric	Ours vs.		
	Textual Inversion	Dreambooth	Custom Diffusion
Text-alignment	90.90	74.54	80.36
Image-alignment	82.54	57.63	48.48
Editing-alignment	96.00	64.00	78.54

to evaluate the efficiency of each method. **Implementation Details.** We use the V1-4 version of Stable Diffusion in our experiments, and the mapping network is implemented with three-layer MLP (for both global mapping network and local mapping network). To extract multilayer CLIP features, features from five layers are selected, and the layer indexes are {25,4,8,12,16} in order. To train global mapping network, we use the batch size of 16 and  $\lambda_{\text{global}} = 0.01$ . The learning rate is set to 1e-6. To train the local mapping network, we adopt the batch size of 8 and  $\lambda_{\text{local}} = 0.0001$ . The learning rate is set to 1e-5. All experiments are conducted on 4xV100 GPUs. During image generation, unless additional stated, Unless mentioned otherwise, we use 100 steps of PLMS sampler with a scale 5. For concept generation (e.g., a  $S^*$ ), we use  $\lambda = 0.8$ , while for concept editing (e.g., a  $S^*$  is swimming), we use  $\lambda = 0.6$ .

### 4.2. Ablation Study

We conduct ablation studies to evaluate the effects of various components in our method, including the multilayer features in the global mapping network, local mapping network, and the value of  $\lambda$ . **Effect of Multi-layer Features.** Fig. 5 first gives the visualization of words learned by multi-layer features in global mapping network. For each word visualization, we use the text A photo of a [wi]. One can see that, the word embedding of the deepest feature (i.e.,  $w_0$ ) describes the primary concept (i.e., corgi, teddybear), while other words describe some finer details. Meanwhile, the obtained  $w_0$  maintains superior editability. To further demonstrate this, we have conducted experiments with several variants: i) Single-layer Single-word: learning a single word embedding from the deepest feature. ii) Single-



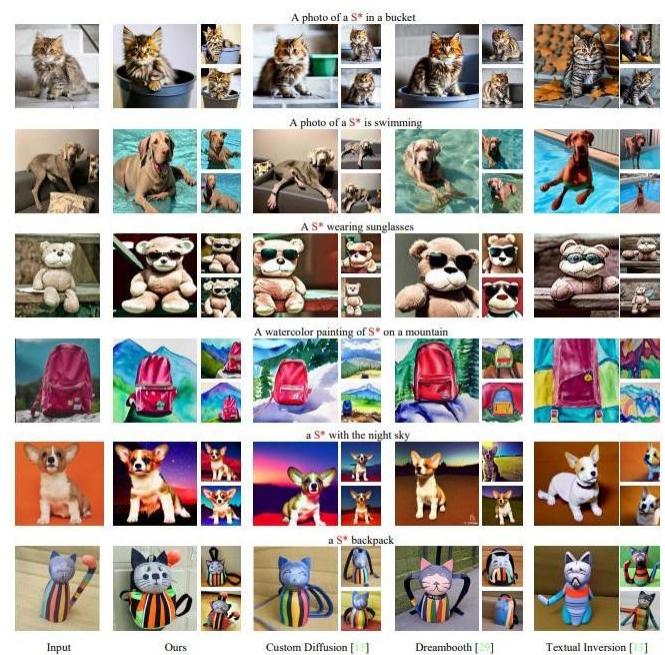
layer Multiwords: learning multiple word embeddings from the deepest feature separately. iii) Multi-layers Multi-words: our setting, learning multiple word embeddings from the multiple layer features separately. Fig. 6 illustrates the results of concept generation and editing for each variant. For multiple word settings, we show the results of the full embeddings (i.e.,  $[v]$ ) and the primary word embedding (denoted as  $[w]$ ), which is obtained by the above analysis. As shown in the figure, encoding concept image into one single word embedding leads to entangled embedding with poor editability. When learning multiple words from the deepest feature, the obtained full word embeddings  $v$  and the primary word embedding  $w$  are also not editable. In contrast, the primary word  $w$  learned by our multifeatures describes the object concept while maintaining superior editing capacity. That's why we only keep it during image generation. Though a single  $w$  is not sufficient to describe the details of the given concept faithfully, a local mapping network is further proposed to address this. Effect of Local Mapping. We further conduct the ablation to evaluate the effect of the proposed local mapping network. From Fig. 6, with the local mapping network, our ELITE generates image with higher consistency with the concept image. Meanwhile, from Table 1, the introduction of a local mapping network does not compromise the editable capability, which demonstrates its superiority over learning multiple words. Though its image alignment may not be the best, this is due to a trade-off between image alignment and text alignment. Effect of  $\lambda$ . From Eqn. 6,  $\lambda$  is introduced to control the fusion of information from the global mapping network and the local mapping network. To evaluate its effect, we vary its value from 0 to 1, and the generated results are shown in Fig. 7. From the figure, with the increasing of  $\lambda$ , the consistency between the synthesized image and concept image is improved. However, when the value of  $\lambda$  is too large, it may lead to degenerated editing results. Therefore, for a tradeoff between inversion and editability, we set  $\lambda = 0.6$  for editing prompts and  $\lambda = 0.8$  for generating prompts. We found these parameters work well for most cases. More ablations are provided in the Suppl.

Figure 8. Visual comparisons of concept editing. Our ELITE method demonstrates superior editability compared to Textual Inver.

### 4.3. Qualitative Results

To demonstrate the effectiveness of our ELITE, we compare it with existing optimization-based methods, including Textual Inversion [13], DreamBooth [29], and Custom Diffusion [15]. For a fair comparison, we trained all models using their official code 1 and default hyperparameters on a

single image.



<sup>1</sup>Since the official code of Dreambooth is not publicly available, we use the code implemented by <https://github.com/XavierXiao/> images generated with the text prompt A photo of a S\*. With only one concept image, Textual Inversion cannot learn a word emsion [13], Dreambooth [29], and Custom Diffusion [15].bedding that accurately describes the target concept. Although Dreambooth and Custom diffusion learn the [Dreambooth-Stable-Diffusion](#). For Textual Inversion, we use their stable diffusion version.

concept with detail consistency, their diversity may be limited. In comparison, our ELITE is capable of faithfully capturing the details of the target concept and generating diverse images. We also conduct experiments with editing prompts and compare our method with existing methods. As shown in Fig. 8, Dreambooth and Custom diffusion exhibit degraded editing ability, and in some cases, the editing prompts fail to produce the desired results (first row). In contrast, our method demonstrates superior editing performance. More qualitative results are provided in Suppl.

### 4.4. Quantitative Results

In addition to the qualitative comparisons, we further conduct the quantitative evaluation to validate the performance of our ELITE. From Table 2, our method achieves better text-alignment compared to the state-of-the-art methods,

demonstrating its superior editability. Moreover, our method achieves comparable detail consistency and image quality, indicating its capability to generate high-quality images. Furthermore, our method

provides a significant advantage in terms of computational efficiency. Unlike optimization-based methods that require several or tens of minutes to obtain the concept embedding, our method can obtain it in just 0.05s. This makes our method highly practical and efficient for realworld applications where speed is a critical factor. User Study. We then perform the user study to compare with competing methods. There are three questions: i) textalignment, the participants were given a text prompt and generated results from two different methods, and were required to select the image that has better consistency with the text. ii) image-alignment, the participants were given a concept image and generated results from two different methods, and were required to select the image that has better consistency with the given concept image. iii) editingalignment, the participants were given a text prompt, a concept image, and the corresponding generated results from two different methods, and were required to select the image that has better consistency with both the concept image and text. As shown in Table 3, our method is preferred over competing methods.

## 5. CONCLUSION

In this paper, we propose a novel learning-based encoder for fast and accurate customized text-to-image synthesis, termed ELITE. Compared with existing optimization-based methods, our ELITE directly encodes visual concepts into textual embeddings, significantly reducing the computational and memory burden of learning new concepts. Moreover, our method demonstrates superior flexibility in editing learned concepts into new scenes while preserving imagespecific details, making it a valuable tool for personalized text-to-image generation. In future work, We will explore to leverage multiple concept images for better inversion, and investigate effective methods for composing multiple concepts in ELITE.

## ACKNOWLEDGMENT

The authors are grateful to Jeppiaar Institute of Technology, Sriperumbudur, India for providing with the necessary facilities to carry out this piece of work.

## REFERENCES

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into thestylegan latent space? In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 4432–4441, 2019. 3
- [2] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan++: How to edit the embedded images? In Proceedings of the

IEEE/CVF conference on computer vision and pattern recognition, pages 8296–8305, 2020. 3

- [3] 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. 1, 2
- [4] Mikołaj Binkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. In International Conference on Learning Representations. 6
- [5] Piotr Bojanowski, Armand Joulin, David Lopez-Paz, and Arthur Szlam. Optimizing the latent space of generative networks. arXiv preprint arXiv:1707.05776, 2017. 3
- [6] Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. Muse: Textto-image generation via masked generative transformers. arXiv preprint arXiv:2301.00704, 2023. 2
- [7] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. In 2021 IEEE/CVF International Conference on Computer Vision (ICCV), pages 14347–14356. IEEE, 2021. 3
- [8] Antonia Creswell and Anil Anthony Bharath. Inverting the generator of a generative adversarial network. IEEE transactions on neural networks and learning systems, 30(7):1967– 1974, 2018. 3
- [9] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. Cogview: Mastering text-to-image generation via transformers. Advances in Neural Information Processing Systems, 34:19822–19835, 2021. 2
- [10] Ming Ding, Wendi Zheng, Wenyi Hong, and Jie Tang. Cogview2: Faster and better text-to-image generation via hierarchical transformers. arXiv preprint arXiv:2204.14217, 2022. 2



- [11] Ziyi Dong, Pengxu Wei, and Liang Lin. Dreamartist: Towards controllable one-shot text-to-image generation via contrastive prompt-tuning. arXiv preprint arXiv:2211.11337, 2022. 3
- [12] Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. Make-a-scene: Scenebased text-to-image generation with human priors. In Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XV, pages 89–106. Springer, 2022. 2
- [13] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. arXiv preprint arXiv:2208.01618, 2022. 1, 2, 3, 4, 5, 6, 7, 8
- [14] 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. 1
- [15] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. arXiv preprint arXiv:2212.04488, 2022. 1, 2, 3, 4, 5, 6, 7, 8
- [16] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision*, 128(7):1956–1981, 2020. 6
- [17] Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, and Philip Torr. Controllable text-to-image generation. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [18] Zachary C Lipton and Subarna Tripathi. Precise recovery of latent vectors from generative adversarial networks. arXiv preprint arXiv:1702.04782, 2017. 3
- [19] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. arXiv preprint arXiv:2211.09794, 2022. 1, 3
- [20] Minheng Ni, Zitong Huang, Kailai Feng, and Wangmeng Zuo. Imaginarynet: Learning object detectors without real images and annotations. arXiv preprint arXiv:2210.06886, 2022. 1
- [21] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741, 2021. 2
- [22] Xichen Pan, Pengda Qin, Yuhong Li, Hui Xue, and Wenhui Chen. Synthesizing coherent story with auto-regressive latent diffusion models. arXiv preprint arXiv:2211.10950, 2022. 1
- [23] Gaurav Parmar, Yijun Li, Jingwan Lu, Richard Zhang, Jun-Yan Zhu, and Krishna Kumar Singh. Spatially-adaptive multilayer selection for gan inversion and editing. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 11399–11409, 2022. 3
- [24] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In International conference on machine learning, pages 8748–8763. PMLR, 2021. 2, 4
- [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. 1, 2
- [26] 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
- [27] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 2287–2296, 2021. 2, 3
- [28] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Bjorn 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, 4