

MagicNaming: Consistent Identity Generation by Finding a “Name Space” in T2I Diffusion Models

Jing Zhao¹, Heliang Zheng³, Chaoyue Wang⁴, Long Lan¹, Wanrong Huang¹, Yuhua Tang^{2*}

¹Department of Intelligent Data Science, College of Computer Science National University of Defense Technology, China

² HPCL, College of Computer Science, National University of Defense Technology, China

³ University of Science and Technology of China, ⁴ University of Sydney

{zhaojing,long.lan,huangwanrong12}@nudt.edu.cn,zhenghl@mail.ustc.edu.cn,
chaoyue.wang@outlook.com,yhtang62@163.com

Abstract

Large-scale text-to-image diffusion models, (e.g., DALL-E, SDXL) are capable of generating famous persons by simply referring to their names. *Is it possible to make such models generate generic identities as simple as the famous ones, e.g., just use a name?* In this paper, we explore the existence of a “Name Space”, where any point in the space corresponds to a specific identity. Fortunately, we find some clues in the feature space spanned by text embedding of celebrities’ names. Specifically, we first extract the embeddings of celebrities’ names in the Laion5B dataset with the text encoder of diffusion models. Such embeddings are used as supervision to learn an encoder that can predict the name (actually an embedding) of a given face image. We experimentally find that such name embeddings work well in promising the generated image with good identity consistency. Note that like the names of celebrities, our predicted name embeddings are disentangled from the semantics of text inputs, making the original generation capability of text-to-image models well-preserved. Moreover, by simply plugging such name embeddings, all variants (e.g., from Civitai) derived from the same base model (i.e., SDXL) readily become identity-aware text-to-image models.

Homepage — <https://magicfusion.github.io/MagicNaming/>

Introduction

Large-scale text-to-image diffusion models such as DALL-E (Ramesh et al. 2022) and Stable Diffusion XL (SDXL) (Rombach et al. 2022) have revolutionized the field of generative models by enabling the synthesis of images from textual descriptions (Arar et al. 2023; Chen et al. 2023; Gal et al. 2023b; Hinz, Heinrich, and Wermter 2022; Li et al. 2022; Liu and Chilton 2022; Liu et al. 2023a; Prabhudesai et al. 2023; Shan et al. 2023; Xie et al. 2023). These models have the impressive capability of generating consistent identity for celebrities simply based on input of textual names. This leads to an intriguing question: Is it possible to extend this capability to generate images of non-known or entirely fictional identities with the same ease, using just a name?

*Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: Identity and textual semantics are disentangled.

In this paper, we investigate the existence of a “Name Space” (\mathcal{N} Space), aiming to predict a “name” for each generic identity and guide the generative model to produce consistent identity generation in a manner similar to celebrity name formats. Specifically, we first investigate how celebrity names enable the generative model to produce consistent IDs. Experiments reveal that name embeddings and textual semantics in a text-to-image model are disentangled; e.g., altering name embeddings or their positions within the model does not disrupt the generation of the original semantics. Likewise, maintaining name embeddings while changing semantics does not affect the consistency of the generated IDs. As shown in Figure 1.

The key to extending the generative model’s capability for consistent ID generation from celebrities to generic identities lies in finding corresponding name embeddings for the latter. To this end, we collect a subset from the laion5B dataset that contains celebrity name and image pairs. We perform text encoding on the names to obtain their respective name embeddings, thereby constructing a large-scale celebrity dataset with their name embeddings. Using this data set, we train an image encoder capable of predicting “name” for any reference individual. To enhance alignment with the textual embedding space, our image encoder employs the same CLIP model as used for text embeddings. Building upon this, we refine the output through a fully connected trilayer network to yield the final name embedding prediction. In the inference generation process, a technique analogous to classifier-free guidance was employed, involv-

ing the recalibration of name embeddings through the utilization of their mean values. Experiments indicate that our proposed image encoder identifies a \mathcal{N} Space, wherein each point corresponds to an individual identity. By embedding this point (i.e., name embedding) into text guidance, we can achieve consistent ID generation for generic identities.

Note that existing research efforts (Wang et al. 2024; Yan et al. 2023; Xiao et al. 2023; Li et al. 2023; Ye et al. 2023) also show the ability to generate consistent identities; however, these approaches predominantly rely on fine-tuning generative models (lora or cross attention layers) on specific datasets, resulting in a bias towards close-up head images. This significantly diminishes the model’s capacity, e.g., for scene construction, stylization, emotional editing, and the action control of humans within generated images. In contrast to these approaches, our method does not introduce additional image conditions for text-to-image models, effectively maintaining the original generative capabilities of the diffusion model. Moreover, our approach does not prescribe a fixed generative model; instead, the \mathcal{N} Space we have identified can be integrated with any variant model based on Stable Diffusion, facilitating the rapid generation of consistent identities. Our contributions are as follows:

- We investigated the relationship between prompts and name embeddings, discovering that textual semantics and name embeddings are disentangled. Consistent ID generation for generic identities can be achieved by finding its name embedding.
- We extracted celebrity images from the Laion5B dataset and obtained name embedding by encoding names with text encoder. Through multi-level data filtering and refinement, we constructed the LaionCele dataset, which comprises 42,000 celebrities and approximately 810,000 authentic images.
- We designed and trained an image encoder to accomplish the mapping of any reference portrait to the \mathcal{N} Space to get the corresponding name embedding. By simply plugging such name embedding into the text embedding, consistent ID generation centered around a specific identity can be achieved. In addition, ID interpolation can also generate fictional characters.

Related works

Text-to-Image Generation

Recent years have witnessed a surge of interest in text-to-image generation (Arar et al. 2023; Chen et al. 2023; Gal et al. 2023b; Hinz, Heinrich, and Wermter 2022; Li et al. 2022; Liu and Chilton 2022; Liu et al. 2023a; Prabhudesai et al. 2023; Shan et al. 2023; Xie et al. 2023; Ho and Salimans 2022) within the field of computer vision, with a plethora of research endeavors propelling rapid advancements in this domain. Stable Diffusion (Rombach et al. 2022) introduced the latent diffusion model, significantly enhancing the generative performance and inference speed of diffusion models for high-resolution images. Building upon this, a myriad of derivative visual tasks have been proposed and explored, including subject-driven text-to-image

generation (Gal et al. 2023a; Ruiz et al. 2023; Wei et al. 2023; Hua et al. 2023; Alaluf et al. 2023; Ku et al. 2023; Ye et al. 2023), text-guided image editing (Avrahami, Lischinski, and Fried 2022; Hertz et al. 2023; Tumanyan et al. 2023; Elarabawy, Kamath, and Denton 2022), model or subject integration (Liu et al. 2022; Kumari et al. 2023; Liu et al. 2023b) and identity-preserved personalization (Yan et al. 2023; Xiao et al. 2023; Li et al. 2023). Overall, text-to-image generation is catalyzing a significant wave in the era of AI-generated content.

Identity-preserved Personalization

The task of Identity-preserved personalization image generation (Yan et al. 2023; Xiao et al. 2023; Li et al. 2023; Peng et al. 2023; Wang et al. 2024; Zhang et al. 2023; Pikoulis, Filntisis, and Maragos 2023; Wang et al. 2024) demands the generation of personalized images centered around a specific individual, with the consistency of the person’s identity being a critical metric of success. IP-Adapter (Ye et al. 2023) is an AI painting tool that generates style-transferred images according to prompts without training Lora, and supports multiple reference images and multiple feature extractions. PhotoMaker (Li et al. 2023) encodes an arbitrary number of input ID images into a stack ID embedding for preserving ID information. Based on the proposed celebrity foundation, the new identity in INDM (Yuan et al. 2023) customized model exhibits better conceptual combination capabilities than previous personalized methods. However, its image quality and id consistency are not ideal, and it still needs fine-tuning. In contrast, our method does not need to fine-tune U-net model, effectively preserving the model’s original generative prowess. Moreover, the introduction of \mathcal{N} Space is compatible with any variant of the SDXL model.

Methods

Decoupling Name Embedding

It is well-established that diffusion models like SDXL, trained on large-scale image-text datasets, inherently possess the capability to generate images with ID consistency of celebrities, where the ID information is specified and controlled through the inclusion of names in the prompts. Our experiments reveal that ID information is localized to the text embeddings associated with name positions in the prompts, i.e., name embedding. As shown in Figure 1, extracting and integrating these ID-specific embeddings into other semantic embedding allows for the consistency of ID in the generated images. This phenomenon shows that ID within text embeddings are disentangled from textual semantics, suggesting that the technique of ID consistency through name embedding in image generation is viable.

Integrating Name Embedding

Since the name embedding contains complete personal ID information, it is crucial to devise a method for integrating name embedding into semantic representations in a way that preserves the integrity of the semantic content while ensuring the generated images maintain consistent ID attributes.

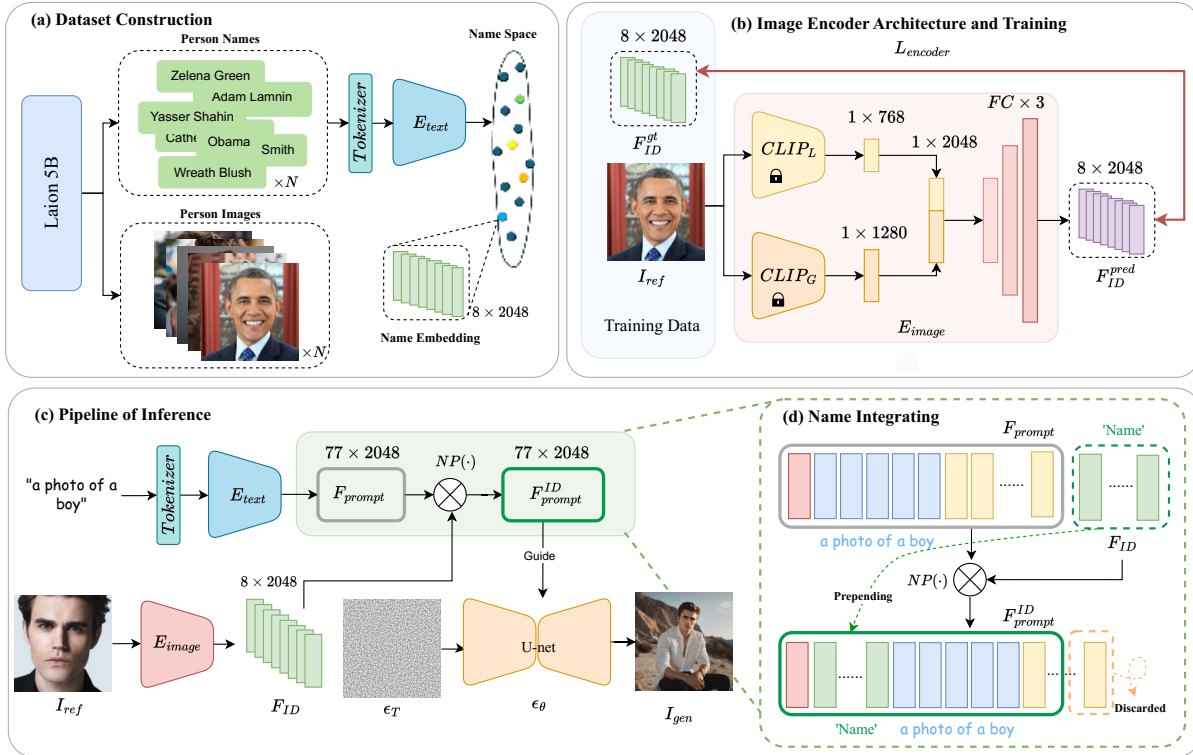


Figure 2: Overview of the Proposed Method. (a) Dataset Construction. Celebrity images and their corresponding names were extracted from the Laion5b dataset, with name embeddings generated through a text encoder E_{text} based on the names. (b) Image Encoder Architecture and Training. Features were extracted from input images using two CLIP image encoders, followed by a three-layer fully connected network to produce the “name” prediction. Mean Squared Error (MSE) loss was computed against the ground truth name embedding F_{ID}^{gt} . (c) Pipeline Inference. The name embedding predicted by the image encoders E_{image} were combined with the original text embeddings by “name” Prepending ($NP(\cdot)$) to obtain final embeddings F_{prompt}^{ID} . These embeddings were then used to guide a U-net model for denoising. (d) “Name” Integrating. There is no need to set a specific placeholder or identify its position, simply inserting the name embedding between the start token (red block) and the first semantic token (blue block) is sufficient to achieve consistent identity generation. The padding tokens (yellow block) exceeding the length of 77 will be discarded.

“Name” substitution is an intuitive method that begins by inserting placeholders like “a man” or “a person” into the original sentence, and then replacing these placeholders with the target ID information in the corresponding name embedding. However, a limitation of this approach is the inconsistency in placeholder positions across various prompts, which makes identifying the location and length of the placeholder within the text embedding challenging.

In this paper, we employ an alternative approach, namely “Name” Prepending (NP) Method. As illustrated in Figure. 2(d), the target “name” F_{ID} is inserted between the start token (the red block) and the first semantic embedding (the blue blocks) of the original text embedding F_{prompt} , getting rid of considering the placeholder’s position or length. Since the text encoder performs non-semantic padding (the yellow block) for texts with fewer than 77 tokens, discarding the excess portion can resolve the embedding length overflow caused by name insertion. The original text embedding $F_{prompt} = E_{text}(T(p))$, where E_{text} denotes the text en-

coder in the diffusion model and T represents the tokenizer which converts a prompt p into tokens. The text embedding resulting from “name” prepping is denoted as F_{prompt}^{ID} , which can be calculated as follows, where $NP(\cdot)$ represents the operation of “name” prepping.

$$F_{prompt}^{ID} = NP(F_{ID}, F_{prompt}) \quad (1)$$

Encoder for Image to \mathcal{N} Space

SDXL inherently possesses the capability to generate ID-consistency images of celebrities; the challenge lies in extending this ability to generic identities. We extract a vast array of celebrity images from the Laion5B dataset and utilize the text encoder to obtain name embeddings corresponding to the names, thereby constructing an extensive ID-name dataset. Employing this dataset, we trained an image encoder E_{image} to achieve the mapping from individual portraits to a \mathcal{N} Space, which contains a variety of name embeddings. The architecture of the image encoder is shown in Figure. 2(b).

To better approximate the text embedding space, our image encoder incorporates two pre-trained CLIP image encoders $CLIP_L$ and $CLIP_G$, which correspond to the text encoder used in SDXL. For any given reference image I_{ref} , features with dimensions 768 and 1280 are extracted using the two CLIP image encoders, respectively. These features are then concatenated and subsequently processed through a series of fully connected layers, resulting in a feature with a dimension of 16384. A reshape operation is finally applied to transform the predicted feature into a name embedding F_{ID}^{pred} with a shape of 8×2048 . F_{ID}^{pred} can be denoted as follows,

$$F_{ID}^{pred} = E_{image}(I_{ref}). \quad (2)$$

The loss function in the training process is defined as follows,

$$\mathcal{L}_{encoder} = \mathbb{E}_{\varepsilon} \|E_{image}(I_{ref}), F_{ID}^{gt}\|_2^2, \quad (3)$$

where F_{ID}^{gt} denotes the ground-truth name embedding, ε denotes the parameters of E_{image} .

Consistent Identity Generation

we adopt classifier-free guidance in our sampling process, which extrapolates model predictions toward text guidance and away from null-text guidance. Moreover, we design a variant of classifier-free guidance in our name embedding processing by leveraging the mean of name embeddings F_{ID}^{mean} . The specific computation rule is as follows,

$$F'_{ID} = F_{ID}^{mean} + \delta(F_{ID} - F_{ID}^{mean}), \quad (4)$$

where F_{ID}^{mean} denotes the average of 42000 name embeddings in the proposed training dataset LaionCele and δ is used to control the strength of the classifier-free guidance. We further introduce a hyper-parameter η to adjust the numeric scale of name embeddings, with the specific computation formula as follows,

$$F''_{ID} = \eta \frac{F'_{ID}}{Norm(F'_{ID})}, \quad (5)$$

where $Norm(\cdot)$ computes the vector norm of a given tensor in PyTorch. F''_{ID} will be the final name embedding.

To obtain the final generated image, we follow previous work (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2021) to conduct an iterative denoising process. Figure. 2(c) illustrates the process of consistent identity generation. Keeping the predicted name embedding F_{ID} unchanged, variations in the input prompts do not affect the individual identity in the generated images. A concise summary of the consistent identity generation is presented in Algorithm 1 in the supplementary material.

Experiments

Dataset Construction

In our study, we filtered celebrity images from the Laion5B dataset based on the associated image captions and matched them with corresponding names. Utilizing the prior knowledge of diffusion models, we encoded the names using a text

encoder to obtain the respective name embedding. Considering the text encoder’s completion mechanism for short texts, we selected the first 8 dimensions of the embedding to serve as the “name” for the individuals. After a rigorous selection and data cleansing process, we retained sample pairs with a minimum of five images per “name”, resulting in a identity-name dataset comprising 42,000 IDs and 810,000 corresponding images, which we have named LaionCele. The detailed construction process of the dataset involves numerous procedural intricacies, which are elaborated in the supplementary material.

Results and Comparison

Evaluation Setup To demonstrate the effectiveness of our method, we conducted both quantitative and qualitative assessments, comparing and analyzing them against relevant works. e.g., IP-Adapter (Ye et al. 2023) and PhotoMaker (Li et al. 2023). For the **testing prompt**, we classified them into four tasks: (a)scene construction, (b)action control, (c)stylization and (d)emotional editing. These prompts present greater challenges compared to simple directives such as changing clothes or wearing hats. It is noteworthy that (a)scene construction and (b)action control pursue the generation of images that contain complete human figures and rich scene semantics rather than merely headshots. For the **evaluation images**, we meticulously selected person images that SDXL could not generate with consistent ID based solely on name text, thereby mitigating the influence of the generative model’s inherent person recognition capabilities on ID consistency of generated images. Ultimately, for each evaluation task, we produced 1,000 test images for each method. The detailed configuration of the **evaluation dataset** (prompts and images) and the specific **implementation details** of the methodology are elaborated in the supplementary material.

Visualization and Comparison The visual outcomes of the aforementioned tests are partially illustrated in Figure. 3, with additional results presented in the supplementary material. SDXL (Rombach et al. 2022), serving as the baseline model, demonstrates its inherent complex semantic generative capabilities for the respective tasks. Given that SDXL lacks identity insertion functionality, we facilitate person image generation by inserting the corresponding personal names into the prompt. The outputs indicate that SDXL inherently lacks the capability to generate consistent IDs for generic identities. Subsequent sections will provide a comparative analysis for each task individually.

(a)Scene Construction. Applications of scene construction with close-up head are highly constrained. Effective scene content and character portrayal are crucial for applications like print advertising and film production. As shown in Figure. 3(a), both IP-Adapter and PhotoMaker primarily generate portraits from the chest or neck up, with scene content corresponding to the text being incorporated only by adding relevant elements to the head’s background. These approaches fall short in comprehensive scene generation. Conversely, our approach consistently produces complete character representations and well-composed scenes in vari-




Figure 3: Visualization and Comparison. The results concludes four tasks, i.e., scene construction, stylization, action control and emotional editing. Please devote attention to semantic consistency and visual aesthetics of images with the same level of concern as given to ID consistency. The results demonstrate that our approach maintains ID consistency while perfectly preserving the original semantic performance (complex semantic consistency) of the generator(SDXL), a feat not achieved by other works.

ous contexts, with high-resolution images, semantically rich environments and visual appeal, demonstrating substantial creative potential in role-based scene construction.

(b) Action Control. The inherent capability of the SDXL to generate human action scenes is notably aesthetically pleasing and impressive, e.g., “a woman is riding a horse”. However, after fine-tuning, most models show a significant drop in this ability. As observed in Figure. 3(b), the generated images by PhotoMaker and IP-Adapter remain constrained to upper-body of individuals, resulting in incomplete scenes lacking in visual appeal. In contrast, our approach not only maintains identity consistency but also delivers stunning visual effects, with complete human actions and a semantically rich scene representation.

(c) Stylization. One of SDXL’s key strength lies in its ability to synthesize diverse-styled imagery from textual descriptions alone. We conducted tests across several classic styles, with the results displayed in Figure. 3(c). The results

show that the IP-adapter has negligible stylization capability. PhotoMaker, while integrating stylistic features, diverges from intended styles and is plagued by poor image fidelity and content ambiguity. Even more, instances of failed identity consistency have been observed in PhotoMaker. Conversely, our method maintains SDXL’s inherent stylization proficiency, producing stylized images with superior quality, clear content and consistent identity.

(d) Emotional Editing. Our method maintain the intrinsic generative power of the generator model (e.g., SDXL), which enables us to leverage sophisticated semantic manipulation features, including the ability to edit emotional expressions. As illustrated in Figure 3(d), our method successfully facilitates the generation of diverse emotional states, while ensuring the synthesized identities remain consistency. In contrast, the facial expressions of characters generated by IP-Adapter and PhotoMaker exhibit less pronounced emotional variations.

	CLIP-TI score \uparrow					FID score with SDXL \downarrow					ID Cons. \uparrow
	Style	Scene	Emotion	Action	All	Style	Scene	Emotion	Action	All	
SDXL	40.28	40.48	27.40	37.01	36.29	-	-	-	-	-	0.3365
IP-Adapter	29.73	33.91	24.86	32.93	30.36	111.72	127.55	108.40	95.28	110.73	0.5345
PhotoMaker	40.37	40.72	29.11	36.93	36.78	76.11	106.68	109.58	88.27	95.16	0.5466
Ours	40.57	43.45	29.31	39.63	38.24	49.52	57.98	73.25	46.25	56.75	0.5624

Table 1: The main quantitative evaluation results. Our approach demonstrates superior performance across all metrics.



Figure 4: Interpolating creates fictional ID with consistent ID generation.



Figure 5: The comparison of training an image encoder using the proposed LaionCele dataset and the generated dataset. It is evident that LaionCele is indispensable, and holds substantial value and contribution to the field.

Quantitative Evaluation In our research, we concentrate more on maintaining complex semantic consistency in images produced by different generative models, which are driven by challenging textual prompts. Additionally, we place a significant emphasis on the visual appeal of the generated images, striving to achieve a high level of aesthetic quality. In this study, we performed grouped evaluations of the aforementioned four specified tasks utilizing a pair of evaluation metrics: the CLIP Text-Image Consistency (CLIP-TI) score and the Fréchet Inception Distance (FID). Given the widespread acclaim of industry and user for the image generation results of SDXL (Rombach et al. 2022), we evaluated the similarity of each model’s outputs to those of SDXL by calculating their FID score. A lower FID score denotes a higher fidelity to SDXL’s synthesis quality

and visual appeal. Table. 1 presents the quantitative results of our evaluation. The outcomes indicate that our approach outperforms other works across all tasks(style, Scene, Emotion, Action) and overall(All) in terms of the CLIP-TI metric. Two points merit particular attention. Firstly, our method achieves higher CLIP-TI scores than SDXL due to the latter’s inability to recognize name-associated identity attributes, resulting in the omission of character within the scene. The introduction of name embedding in our technique reinforces the concept of person, thereby mitigating the issue of character loss and enhancing semantic consistency. Secondly, although PhotoMaker’s CLIP-TI scores are comparable to those of SDXL, qualitative assessments reveal that PhotoMaker’s close-up head images are noticeably inferior in semantic consistency to the well-composed scenes of SDXL. This suggests that CLIP-TI scores may not effectively discriminate between models in terms of scene construction and visual aesthetic appeal.

Superior similarity with SDXL by FID score indicates exceptional capabilities in scene construction, stylization, emotional editing, action control, and the visual aesthetics of the imagery. The results presented in Table. 1 indicate that our method exhibits a significant advantage in terms of the FID metric. The underlying reason is that our approach does not alter the generator (SDXL) in any way, and the “name” prepending operation during the inference process does not induce a shift in the text embedding space. Consequently, this maintains the consistency of identity generation while perfectly preserving the original performance of the generative model. Table. 1 also presents the quantitative assessment results for generated identity consistency. To mitigate the visual bias introduced by effective stylization, as well as the impact of reduced facial features due to the diminution of faces in full-figure presentations within scene construction and action control tasks, we employed the emotional editing task that showcases facial details for the evaluation of identity consistency. The results indicate that our approach also outperforms other methods in terms of generating consistent identities.

Create Fictional Identities

In the \mathcal{N} space, any given point corresponds to a person identity. For real individuals, we can obtain their mapping in the \mathcal{N} space using the image encoder proposed and trained in this paper, thereby achieving consistent identity generation. Furthermore, by interpolating between any two points (i.e., name embedding) within the name space, we can cre-

		CLIP-TI \uparrow	Ref Simi. \uparrow	ID Cons. \uparrow
Ours(w/o ft)		43.45	0.4008	0.4407
LaionCele	LoRA	37.54	0.5004	0.6236
	U-net	32.55	0.4475	0.6631
Ins700	LoRA	40.66	0.4816	0.5742
	U-net	36.51	0.4502	0.6356

Table 2: Quantitative Evaluation of Fine-Tuning.



Figure 6: Application on StyleUnet: By integrating a “name” into any variant of a stable diffusion, consistent ID generation can be achieved.

ate new fictional characters. As illustrated in Figure 4, the appearance of the fictional characters can be flexibly controlled by adjusting the interpolation ratio. Notably, these fictional characters also support consistent identity generation. This will provide significant practical value for film and game production.

Ablation Study

The Necessity of Constructing LaionCele Datasets. An intuitive approach to construct dataset is leveraging the prior knowledge inherent in SDXL for generation. However, we eschewed this method in favor of an arduous process of filtering and cleaning celebrity image data from the LAION-5B dataset, ultimately yielding a large-scale identity-name dataset, LaionCele. To assess the disparity between the two datasets (LaionCele and Generated dataset), we also trained an image encoder, E_{gen} , with the generated images, and compared its performance with the our image encoder E_{cele} trained with LaionCele. Figure. 5 illustrates the comparative results. E_{cele} is capable of producing high-quality images with ID and semantic consistency. In contrast, images generated by E_{gen} exhibit numerous artifacts and lack identity consistency. The results indicate that, on one hand, even well-performing generative models like SDXL exhibit a significant gap between their image outputs and real images. On the other hand, the gap between real test images and generated training images also impacts inference performance. Therefore, it is evident that the dataset constructed in this study, LaionCele, holds substantial value and contribution to the field.

How Does Fine-tuning Impair Generative Capacity?

The primary distinction between our approach and other works lies in our decision not to fine-tune the generative

model, thereby preserving its original performance intact. Consequently, one might inquire: how does fine-tuning impair model’s Generative Capacity? we explored the way of fine-tuning, entire U-net or LoRA on two dataset, i.e., the proposed LaionCele dataset and Ins700(comprising images of 700 individuals scraped from the Instagram website). The scene construction task is employed for evaluation and the results are presented in Table. 2. The implementation details of fine-tuning are provided in the supplementary material. Experimental results indicate that: 1) Fine-tuning U-net and LoRA both incur a trade-off, enhancing identity consistency(ID Cons.) and reference similarity(Ref Simi.) at the expense of generative performance(CLIP-TI). 2)Fine-tuning U-net results in a more pronounced degradation of generative capability, yet it yields greater improvements in identity consistency. 3) Fine-tuning LoRA is more beneficial for augmenting reference similarity. 4) Compared to the larger LaionCele dataset, smaller datasets (Ins700) exert less influence on all metrics measured.

Application: Making any U-net Identity-aware

The proposed \mathcal{N} Space for consistent ID generation is agnostic to the generative model and can be seamlessly integrated with any variant based on SDXL to achieve ID-consistent image generation. In the aforementioned experiments, our uniformly employed SDXL v1.0 as the generator. To demonstrate the broad applicability of the \mathcal{N} Space, we conducted additional experiments on a U-net variant (styleUnet) specialized in style generation. Figure. 6 present the consistent ID generation results with styleUnet. The experimental outcomes indicate that the “name” sampled from the \mathcal{N} Space can be easily utilized with alternative generative models to produce consistent ID generation without compromising the original specialized capabilities.

Conclusions

In this paper, we present a novel approach for consistent identity generation of generic identities by introducing a \mathcal{N} Space. Specifically, we construct a large-scale identity-name dataset, LaionCele, and train an image encoder that maps real images into this \mathcal{N} Space. This encoder operates independently of the generator, allowing for integration with any SDXL-based generative model to achieve consistent ID generation, offering broad practical applications. A key advantage of our method is that it preserves the original generative capabilities of the generator model, including scene construction, stylization, emotional editing, action control, and so on, thus enhancing the flexibility and creativity of consistent ID generation. Extensive experimental results demonstrated that our approach surpasses existing and concurrent works in terms of ID consistency, semantic consistency, image quality and visual aesthetics of the generated images.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Grant Nos. 62106278, 91948303-1, 611803375, 12002380, 62101575) and the National Key R&D Program of China (Grant No. 2021ZD0140301).

References

- Alaluf, Y.; Richardson, E.; Metzger, G.; and Cohen-Or, D. 2023. A Neural Space-Time Representation for Text-to-Image Personalization. *ACM Trans. Graph.*, 42(6): 243:1–243:10.
- Arar, M.; Gal, R.; Atzmon, Y.; Chechik, G.; Cohen-Or, D.; Shamir, A.; and Bermano, A. H. 2023. Domain-Agnostic Tuning-Encoder for Fast Personalization of Text-to-Image Models. In *SIGGRAPH Asia*, 72:1–72:10. ACM.
- Avrahami, O.; Lischinski, D.; and Fried, O. 2022. Blended Diffusion for Text-driven Editing of Natural Images. In *CVPR*, 18187–18197. IEEE.
- Chen, H.; Zhang, Y.; Wang, X.; Duan, X.; Zhou, Y.; and Zhu, W. 2023. DisenBooth: Identity-Preserving Disentangled Tuning for Subject-Driven Text-to-Image Generation. *CoRR*, abs/2305.03374.
- Elarabawy, A.; Kamath, H.; and Denton, S. 2022. Direct Inversion: Optimization-Free Text-Driven Real Image Editing with Diffusion Models. *CoRR*, abs/2211.07825.
- Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2023a. An Image is Worth One Word: Personalizing Text-to-Image Generation using Textual Inversion. In *ICLR*. OpenReview.net.
- Gal, R.; Arar, M.; Atzmon, Y.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2023b. Designing an Encoder for Fast Personalization of Text-to-Image Models. *CoRR*, abs/2302.12228.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2023. Prompt-to-Prompt Image Editing with Cross-Attention Control. In *ICLR*. OpenReview.net.
- Hinz, T.; Heinrich, S.; and Wermter, S. 2022. Semantic Object Accuracy for Generative Text-to-Image Synthesis. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(3): 1552–1565.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *NeurIPS*.
- Ho, J.; and Salimans, T. 2022. Classifier-Free Diffusion Guidance. *CoRR*, abs/2207.12598.
- Hua, M.; Liu, J.; Ding, F.; Liu, W.; Wu, J.; and He, Q. 2023. DreamTuner: Single Image is Enough for Subject-Driven Generation. *CoRR*, abs/2312.13691.
- Ku, M.; Li, T.; Zhang, K.; Lu, Y.; Fu, X.; Zhuang, W.; and Chen, W. 2023. ImagenHub: Standardizing the evaluation of conditional image generation models. *CoRR*, abs/2310.01596.
- Kumari, N.; Zhang, B.; Zhang, R.; Shechtman, E.; and Zhu, J. 2023. Multi-Concept Customization of Text-to-Image Diffusion. In *CVPR*, 1931–1941. IEEE.
- Li, W.; Wen, S.; Shi, K.; Yang, Y.; and Huang, T. 2022. Neural Architecture Search With a Lightweight Transformer for Text-to-Image Synthesis. *IEEE Trans. Netw. Sci. Eng.*, 9(3): 1567–1576.
- Li, Z.; Cao, M.; Wang, X.; Qi, Z.; Cheng, M.; and Shan, Y. 2023. PhotoMaker: Customizing Realistic Human Photos via Stacked ID Embedding. *CoRR*, abs/2312.04461.
- Liu, N.; Li, S.; Du, Y.; Torralba, A.; and Tenenbaum, J. B. 2022. Compositional Visual Generation with Composable Diffusion Models. In *ECCV (17)*, volume 13677 of *Lecture Notes in Computer Science*, 423–439. Springer.
- Liu, V.; and Chilton, L. B. 2022. Design Guidelines for Prompt Engineering Text-to-Image Generative Models. In *CHI*, 384:1–384:23. ACM.
- Liu, X.; Zhang, X.; Ma, J.; Peng, J.; and Liu, Q. 2023a. InstaFlow: One Step is Enough for High-Quality Diffusion-Based Text-to-Image Generation. *CoRR*, abs/2309.06380.
- Liu, Z.; Feng, R.; Zhu, K.; Zhang, Y.; Zheng, K.; Liu, Y.; Zhao, D.; Zhou, J.; and Cao, Y. 2023b. Cones: Concept Neurons in Diffusion Models for Customized Generation. In *ICML*, volume 202 of *Proceedings of Machine Learning Research*, 21548–21566. PMLR.
- Peng, X.; Zhu, J.; Jiang, B.; Tai, Y.; Luo, D.; Zhang, J.; Lin, W.; Jin, T.; Wang, C.; and Ji, R. 2023. PortraitBooth: A Versatile Portrait Model for Fast Identity-preserved Personalization. *CoRR*, abs/2312.06354.
- Pikoulis, I.; Filntisis, P. P.; and Maragos, P. 2023. Photorealistic and Identity-Preserving Image-Based Emotion Manipulation with Latent Diffusion Models. *CoRR*, abs/2308.03183.
- Prabhudesai, M.; Goyal, A.; Pathak, D.; and Fragkiadaki, K. 2023. Aligning Text-to-Image Diffusion Models with Reward Backpropagation. *CoRR*, abs/2310.03739.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. *CoRR*, abs/2204.06125.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *CVPR*, 10674–10685. IEEE.
- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. DreamBooth: Fine Tuning Text-to-Image Diffusion Models for Subject-Driven Generation. In *CVPR*, 22500–22510. IEEE.
- Shan, S.; Ding, W.; Passananti, J.; Zheng, H.; and Zhao, B. Y. 2023. Prompt-Specific Poisoning Attacks on Text-to-Image Generative Models. *CoRR*, abs/2310.13828.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising Diffusion Implicit Models. In *ICLR*.
- Tumanyan, N.; Geyer, M.; Bagon, S.; and Dekel, T. 2023. Plug-and-Play Diffusion Features for Text-Driven Image-to-Image Translation. In *CVPR*, 1921–1930. IEEE.
- Wang, Q.; Bai, X.; Wang, H.; Qin, Z.; and Chen, A. 2024. InstantID: Zero-shot Identity-Preserving Generation in Seconds. *CoRR*, abs/2401.07519.
- Wei, Y.; Zhang, Y.; Ji, Z.; Bai, J.; Zhang, L.; and Zuo, W. 2023. ELITE: Encoding Visual Concepts into Textual Embeddings for Customized Text-to-Image Generation. In *ICCV*, 15897–15907. IEEE.
- Xiao, G.; Yin, T.; Freeman, W. T.; Durand, F.; and Han, S. 2023. FastComposer: Tuning-Free Multi-Subject Image Generation with Localized Attention. *CoRR*, abs/2305.10431.

Xie, J.; Li, Y.; Huang, Y.; Liu, H.; Zhang, W.; Zheng, Y.; and Shou, M. Z. 2023. BoxDiff: Text-to-Image Synthesis with Training-Free Box-Constrained Diffusion. In *ICCV*, 7418–7427. IEEE.

Yan, Y.; Zhang, C.; Wang, R.; Zhou, Y.; Zhang, G.; Cheng, P.; Yu, G.; and Fu, B. 2023. FaceStudio: Put Your Face Everywhere in Seconds. *CoRR*, abs/2312.02663.

Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. IP-Adapter: Text Compatible Image Prompt Adapter for Text-to-Image Diffusion Models. *CoRR*, abs/2308.06721.

Yuan, G.; Cun, X.; Zhang, Y.; Li, M.; Qi, C.; Wang, X.; Shan, Y.; and Zheng, H. 2023. Inserting Anybody in Diffusion Models via Celeb Basis. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Zhang, B.; Qi, C.; Zhang, P.; Zhang, B.; Wu, H.; Chen, D.; Chen, Q.; Wang, Y.; and Wen, F. 2023. MetaPortrait: Identity-Preserving Talking Head Generation with Fast Personalized Adaptation. In *CVPR*, 22096–22105. IEEE.