
OneActor: Consistent Subject Generation via Cluster-Conditioned Guidance

Jiahao Wang^{1,2}, Caixia Yan^{1,*}, Haonan Lin¹, Weizhan Zhang^{1,*}, Mengmeng Wang^{3,4},
Tieliang Gong¹, Guang Dai⁴, Hao Sun⁵

¹School of Computer Science and Technology, MOEKLINNS, Xi'an Jiaotong University

²State Key Laboratory of Communication Content Cognition

³College of Computer Science and Technology, Zhejiang University of Technology

⁴SGIT AI Lab, State Grid Corporation of China

⁵China Telecom Artificial Intelligence Technology Co.Ltd

uguisu@stu.xjtu.edu.cn {yancaixia, zhangwzh}@xjtu.edu.cn

Abstract

Text-to-image diffusion models benefit artists with high-quality image generation. Yet their stochastic nature hinders artists from creating consistent images of the same subject. Existing methods try to tackle this challenge and generate consistent content in various ways. However, they either depend on external restricted data or require expensive tuning of the diffusion model. For this issue, we propose a novel one-shot tuning paradigm, termed OneActor. It efficiently performs consistent subject generation solely driven by prompts via a learned semantic guidance to bypass the laborious backbone tuning. We lead the way to formalize the objective of consistent subject generation from a clustering perspective, and thus design a cluster-conditioned model. To mitigate the overfitting challenge shared by one-shot tuning pipelines, we augment the tuning with auxiliary samples and devise two inference strategies: semantic interpolation and cluster guidance. These techniques are later verified to significantly improve the generation quality. Comprehensive experiments show that our method outperforms a variety of baselines with satisfactory subject consistency, superior prompt conformity as well as high image quality. Our method is capable of multi-subject generation and compatible with popular diffusion extensions. Besides, we achieve a $4\times$ faster tuning speed than tuning-based baselines and, if desired, avoid increasing the inference time. Furthermore, our method can be naturally utilized to pre-train a consistent subject generation network from scratch, which will implement this research task into more practical applications. Project page: <https://johnniewang.github.io/OneActor-webpage/>.

1 Introduction

Diffusion probabilistic models [13, 31, 34] have shown great success in image generation. As a prominent subset of them, text-to-image (T2I) diffusion models [14, 28] significantly improve artistic productivity. Designers can simply describe a subject (e.g. a character, an object, an art style) that they desire and then obtain high-quality images of the subject. However, relying on the random sampling, diffusion models fail to maintain a consistent appearance of the subject across the generated images. Taking Fig. 1(a) as an example, ordinary diffusion models denoise a random noise from the noise space to a clean latent code guided by the given prompt, where the clean code corresponds to a

* Corresponding authors.

This work was completed during the internship at SGIT AI Lab, State Grid Corporation of China.

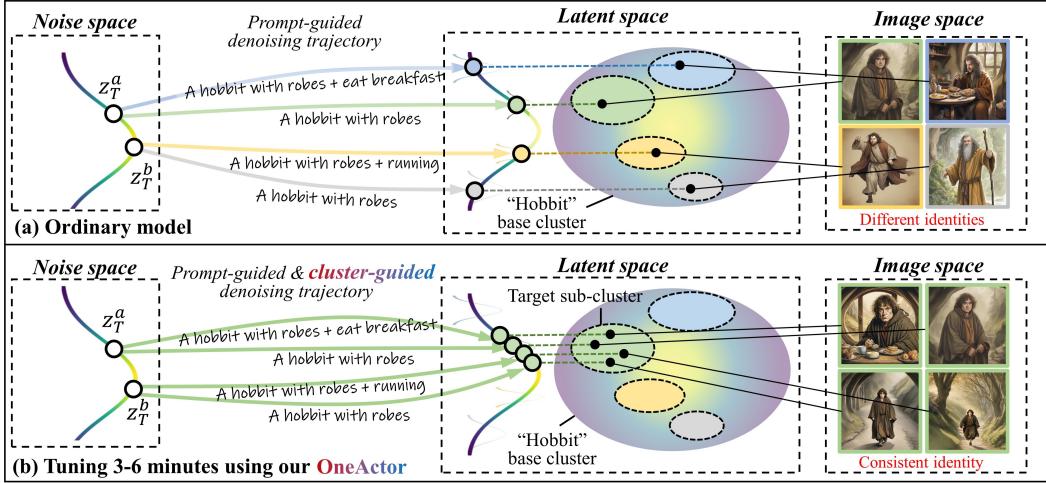


Figure 1: For every subject in the latent space, there are identity sub-clusters within the subject base cluster. (a) Given different prompts and initial noises, ordinary diffusion models generate inconsistent images from different identity sub-clusters of the "hobbit" base cluster. (b) While our OneActor, after a quick tuning, provides an extra cluster guidance and thus generates images from the same target sub-cluster that show a consistent identity. Different colors denote different identity sub-clusters.

salient image in the image space. As shown, if given 4 different prompts of the same subject "hobbit" and random noises, ordinary models will generate 4 hobbits with different identities. In other words, the generation of the ordinary models lacks *subject consistency*.

The subject consistency is a necessity in practical scenarios such as designing an animation character, advertising a product and drawing a storybook protagonist. As diffusion models prevail, many works try to harness the diffusion models to generate consistent content through the following paths. *Personalization* [8, 33, 42, 39] learns to represent a new subject from several given images and thus generates images of that. *Storybook visualization* [23, 30] manages to generate consistent characters throughout a story. However, they require external data to function, either a given image set or a specific storybook dataset. Such a requirement not only complicates practical usage, but also limits the ability to illustrating imaginary, fictional or novel subjects. More recently, *consistent subject generation* [4] is first proposed to generate consistent images of the same subject only driven by prompts. Thus, artists can easily describe the subject and start their creation, which is a more intuitive manner. Despite the pioneering work, their repetitive tuning process of the backbone model leads to expensive computation cost and possible quality degradation. Later, a tuning-free approach [37] equips the backbone model with handicraft modules to eliminate the tuning process. Yet, the extra modules double the inference time and exhibit inefficiency when generating a large number of images.

In fact, diffusion models can generate two consistent images albeit with low probability, demonstrating their inherent potential for consistent subject generation. Motivated by this, we believe all the ordinary model needs is a learned semantic guidance to tame its internal potential. To this end, we propose a novel one-shot tuning paradigm, termed as OneActor. We start from the insight that in the latent space of a pre-trained diffusion model, samples of different subjects form different clusters [4, 41], i.e. *base cluster*. In a specific base cluster, samples that share common appearance gather into the same sub-cluster, i.e. *identity sub-cluster*. In the latent space of Fig. 1, for example, a "hobbit" base cluster contains four different identity sub-clusters. Ordinary generations spread out into four different sub-clusters, causing subject inconsistency. While in our paradigm of Fig. 1(b), users first choose one satisfactory image from the generated proposals as the target. Then after a quick tuning, our OneActor quickly learns to find the denoising trajectories towards the target sub-cluster, generating images of the same subject. Throughout the process, we expect to learn in the semantic space and avoid harming the inner capacity of the latent space.

To achieve high-quality and efficient generation, we pioneer the cluster-guided formalization of the consistent subject generation task and derive a cluster-based score function. This score function underpins our novel cluster-conditioned paradigm which leverages posterior samples as cluster repre-

sentations. To systematically address overfitting, we enhance the tuning with auxiliary samples to robustly align the representation space to the semantic space. We further develop the semantic interpolation and cluster guidance scale strategies for more nuanced and controlled inference. Extensive experiments demonstrate that with superior subject consistency and prompt conformity, our method forms a new Pareto front over the baselines. Our method requires only 3-6 minutes for tuning, which is at least $4\times$ faster than tuning-based pipelines. This efficiency gain is achieved without necessarily increasing the inference time, making our method highly suitable for large-scale image generation tasks. Additionally, our method’s inherent flexibility allows consistent multi-subject generation and seamless integration into various workflows like ControlNet [43].

Our main contributions are: (1) We revolutionize the consistent subject generation by pioneering the cluster-guided task formalization, which supersedes the laborious backbone tuning with a learned semantic guidance to perform efficient generation. (2) We introduce a novel cluster-conditioned generation paradigm, termed as OneActor, to pursue more generalized, nuanced and controlled consistent subject generation. It highlights an auxiliary augmentation during tuning and features the semantic interpolation and cluster guidance scale strategies. (3) Extensive experiments demonstrate our method’s superior subject consistency, excellent prompt conformity and the $4\times$ faster tuning speed without inference time increase. (4) We first establish the semantic-latent guidance equivalence of T2I models, offering a promising tool for precise generation control.

2 Preliminaries

Before introducing our method, we first give a brief review of ordinary text-to-image diffusion models. To start with, Gaussian diffusion models [34, 13] assume a forward Markov process that gradually adds noise to normal image \mathbf{x}_0 :

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad (1)$$

where $t \in [0, T]$, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\bar{\alpha}_t$ are a set of constants. Meanwhile, a denoising network ϵ_θ , usually a U-Net [32], is trained to reverse the forward process by estimating the noise given a corrupted image:

$$\mathcal{L}(\theta) = \mathbb{E}_{t \in [1, T], \mathbf{x}_0, \epsilon_t} [\|\epsilon_t - \epsilon_\theta(\mathbf{x}_t, t)\|^2]. \quad (2)$$

Once trained, we can sample images \mathbf{x}_0 from a Gaussian noise \mathbf{x}_t by gradually removing the noise step by step with ϵ_θ . To introduce conditional control, classifier-free guidance [14] trains ϵ_θ in both unconditional and conditional manners: $\epsilon_\theta(\mathbf{x}_t, t, \mathbf{c}_\emptyset)$ and $\epsilon_\theta(\mathbf{x}_t, t, \mathbf{c})$, where \mathbf{c} is the given condition and \mathbf{c}_\emptyset indicates no condition. Images are then sampled by the combination of two types of outputs:

$$\epsilon_\theta(\mathbf{x}_t, t, \mathbf{c}_\emptyset) + s \cdot (\epsilon_\theta(\mathbf{x}_t, t, \mathbf{c}) - \epsilon_\theta(\mathbf{x}_t, t, \mathbf{c}_\emptyset)), \quad (3)$$

where s is the guidance scale. For text control, conditions \mathbf{c} are generated by a text encoder E_t , which projects text prompts \mathbf{p} into semantic embeddings: $\mathbf{c} = E_t(\mathbf{p})$. To accelerate the pipeline, latent diffusion models [31] pre-train an autoencoder to compress images into latent codes: $\mathbf{z} = E_a(\mathbf{x})$, $\mathbf{x} = D_a(\mathbf{z})$. Thus, the whole diffusion process can be carried out in the latent space instead of the salient image space.

3 Method

In our task, given a user-defined description prompt \mathbf{p}^{tar} (e.g. *a hobbit with robes*), a user-preferred image \mathbf{x}^{tar} is generated by an ordinary diffusion model ϵ_θ and chosen as the target subject. Our goal is to equip the original ϵ_θ with a supportive network ϕ , formulating $\epsilon_{\theta, \phi}$. After a quick one-shot tuning of ϕ , our model can generate consistent images of the same subject with any other subject-centric prompt $\mathbf{p}^{\text{sub}} = \{\mathbf{p}^{\text{tar}}, \mathbf{p}\}$ (e.g. *a hobbit with robes + walking on the street*). To accomplish this task, We first give mathematical analysis in Sec. 3.1. Then we construct a cluster-conditioned model and tune it in a generalized manner in Sec. 3.2. During inference, we generate diverse consistent images with the semantic interpolation and cluster guidance in Sec. 3.3.

3.1 Derivation of Cluster-Guided Score Function

To start with, let’s revisit the ordinary generation process. Given N initial noises and the same subject prompt, generations of ϵ_θ fail to reach one specific sub-cluster, but spread to a base region $\mathcal{S}^{\text{base}}$ of

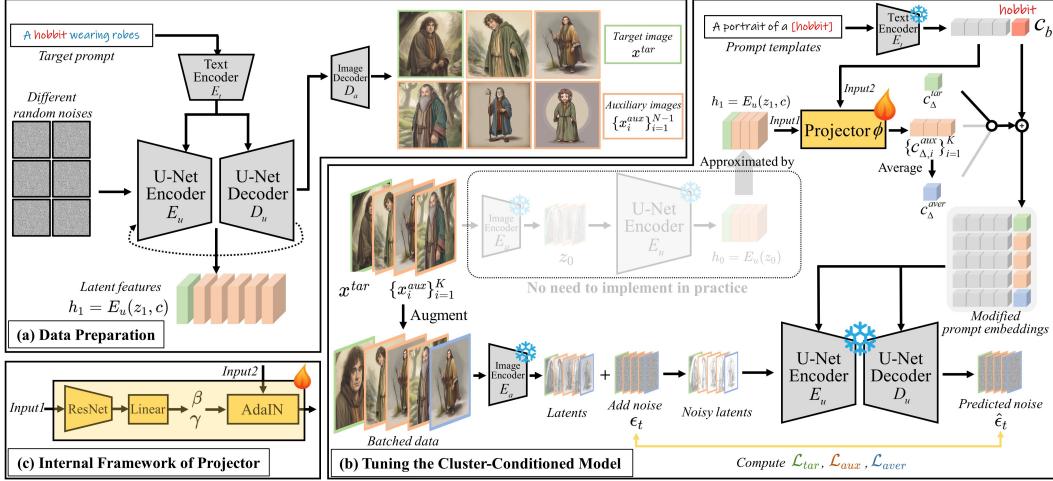


Figure 2: The overall architecture of our method. (a) We first generate base images and construct the target and auxiliary set. (b) We design a cluster-conditioned model and tune the projector with batched data. (c) The projector consists of a ResNet network, linear and AdaIN layers. Tuning and freezing weights are denoted by fire and snowflake marks. The items used to compute different objectives are outlined in different colors. The unimplemented theoretical models are semi-transparent.

different sub-clusters. If we choose one sub-cluster as the target \mathcal{S}^{tar} and denote the rest as auxiliary sub-clusters \mathcal{S}_i^{aux} , then $\mathcal{S}^{base} = \mathcal{S}^{tar} \cup \{\mathcal{S}_i^{aux}\}_{i=1}^{N-1}$. The key to consistent subject generation is to guide the denoising trajectories of ϵ_{θ} towards the expected target sub-cluster \mathcal{S}^{tar} . From a result-oriented perspective, we expect to increase the probability of generating images of the target sub-cluster \mathcal{S}^{tar} and reduce that of the auxiliary sub-clusters \mathcal{S}_i^{aux} . Thus, if we consider the original diffusion process as a prior distribution $p(\mathbf{x})$, our expected distribution can be denoted as:

$$p(\mathbf{x}) \cdot \frac{p(\mathcal{S}^{tar} | \mathbf{x})}{\prod_{i=1}^{N-1} p(\mathcal{S}_i^{aux} | \mathbf{x})}. \quad (4)$$

We take the negative gradient of the log likelihood to derive:

$$-\nabla_{\mathbf{x}} \log p(\mathbf{x}) - \nabla_{\mathbf{x}} \log p(\mathcal{S}^{tar} | \mathbf{x}) + \sum_{i=1}^{N-1} \nabla_{\mathbf{x}} \log p(\mathcal{S}_i^{aux} | \mathbf{x}). \quad (5)$$

With the reparameterization trick of [13], we can further express the scores as the predictions of the denoising network θ from a latent diffusion model [31] in a classifier-free manner [14]:

$$\epsilon_{\theta}(z_t, t) + \eta_1 \cdot \underbrace{[\epsilon_{\theta}(z_t, t, \mathcal{S}^{tar}) - \epsilon_{\theta}(z_t, t)]}_{\text{target attraction score}} - \eta_2 \cdot \sum_{i=1}^{N-1} \underbrace{[\epsilon_{\theta}(z_t, t, \mathcal{S}_i^{aux}) - \epsilon_{\theta}(z_t, t)]}_{\text{auxiliary exclusion score}}, \quad (6)$$

where η_1, η_2 are guidance control factors. If we introduce the concept of score function [35], Eq. (6) can be regarded as a combination of *target attraction score* and *auxiliary exclusion score*. This formula, termed as *cluster-guided score function*, is the core of our method. We will manage to realize it in our subsequent parts. The detailed derivation is shown in Appendix F.

3.2 Generalized Tuning of Cluster-Conditioned Model

In accordance with theoretical analysis of Eq. (6), we propose a novel tuning pipeline as shown in Fig. 2. In general, we incorporate the cluster representations to construct a cluster-conditioned model $\epsilon_{\theta}(z_t, t, \mathcal{S})$ with a supportive network ϕ . We then prepare the data and tune the model with the help of auxiliary samples.

Cluster-Conditioned Model. In this model, the supportive network ϕ is designed to process the posterior samples into cluster representations of the semantic space. To specify, for each sample

\mathbf{x}, ϕ transforms its latent code \mathbf{z} and prompt embedding \mathbf{c} into a subject-specific vector c_Δ , i.e. $c_\Delta = \phi(\mathbf{z}, \mathbf{c})$. For the framework of ϕ , a naive way is to use a feature extractor and a space projector. Since the original U-Net encoder E_u is already well-trained to extract features from the latent codes, we use it directly as the extractor. However, the U-Net extractor may cause extra computational burden. To bypass the extractor, we approximate the features of \mathbf{z}_0 with that of \mathbf{z}_1 :

$$\mathbf{h} = E_u(\mathbf{z}_1, \mathbf{c}) \approx E_u(\mathbf{z}_0). \quad (7)$$

Thus, we save the intermediates \mathbf{h} of the U-Net encoder during data generation and then directly feed them into the projector for the semantic output: $c_\Delta = \phi(\mathbf{h}, \mathbf{c}) = \phi(\mathbf{h}, E_t(\mathbf{p}))$. This approximation condenses ϕ to only a projector and lowers the computational cost by 30%. The output vector c_Δ represents the semantic direction of the sample's sub-cluster. We then split \mathbf{c} into word-wise embeddings c_i to locate the base word embedding c_b and offset it by:

$$c'_b = c_b + c_\Delta. \quad (8)$$

The modified embedding c'_b later replaces c_b to form \mathbf{c}' and guides the noise predictions. By now, all the factors (i.e. \mathbf{p}, \mathbf{h}) that could determine a sub-cluster are involved in the cluster-conditioned model, so we can express the cluster-related terms in Eq. (6) as:

$$\epsilon_{\theta}(\mathbf{z}_t, t, \mathcal{S}) = \epsilon_{\theta, \phi}(\mathbf{z}_t, t, \mathbf{c}, c_\Delta) = \epsilon_{\theta, \phi}(\mathbf{z}_t, t, E_t(\mathbf{p}), \phi(\mathbf{h}, E_t(\mathbf{p}))). \quad (9)$$

Generalized Tuning with Auxiliary Samples. The major challenge of one-shot tuning is overfitting because insufficiency of data may lead to severe bias and limited diversity of generated images. To overcome this challenge, we tune the model with not only target samples, but also auxiliary samples. To elaborate, as shown in Fig. 2(a), given \mathbf{p}^{tar} (e.g. *a hobbit with robes*) which contains the base word p_i (e.g. **hobbit**), we first input it into the ordinary diffusion model for N base images and the corresponding intermediates to form a base set: $\mathcal{X}^{base} = \{\mathbf{x}_i^{base}, \mathbf{h}_i^{base}\}_{i=1}^N$. We randomly choose one image as the target sample $\mathbf{x}^{tar}, \mathbf{h}^{tar}$ and gather the others to form an auxiliary set: $\mathcal{X}^{aux} = \{\mathbf{x}_i^{aux}, \mathbf{h}_i^{aux}\}_{i=1}^{N-1}$. We apply face crop and image flip to target image for an augmented set: $\mathcal{X}^{tar} = \{\mathbf{x}_i^{tar}, \mathbf{h}_i^{tar}\}_{i=1}^M$. Then in every tuning step, we randomly select 1 target and K auxiliary samples to form a batch of data: $\mathcal{B} = \{\mathbf{x}^{tar}, \mathbf{h}^{tar}\} \cup \{\mathbf{x}_i^{aux}, \mathbf{h}_i^{aux}\}_{i=1}^K$. Thus, more data contributes to the optimization and the batch normalization can be applied to the projector, resulting in a more generalized projection. To carry out the tuning, in Fig. 2(b), we first get the latent codes by: $\mathbf{z} = E_a(\mathbf{x})$, $\mathbf{x} \in \mathcal{B}$ and add noise ϵ_t to them. Then we input the noisy latent \mathbf{z}_t , prompt \mathbf{p} and feature \mathbf{h} to the cluster-conditioned model with the projector ϕ . The prompt \mathbf{p} is a random template filled with the base word (e.g. *a portrait of a hobbit*). We apply the standard denoising loss for both target and auxiliary samples:

$$\mathcal{L}_{tar}(\phi) = \mathbb{E}_{t \in [1, T], \mathbf{z}_0^{tar}, \epsilon_t} [\|\epsilon_t - \epsilon_{\theta, \phi}(\mathbf{z}_0^{tar}, t, E_t(\mathbf{p}), \phi(\mathbf{h}^{tar}, \mathbf{c}))\|^2], \quad (10)$$

$$\mathcal{L}_{aux}(\phi) = \mathbb{E}_{t \in [1, T], \mathbf{z}_0^{aux}, \epsilon_t} [\|\epsilon_t - \epsilon_{\theta, \phi}(\mathbf{z}_0^{aux}, t, E_t(\mathbf{p}), \phi(\mathbf{h}^{aux}, \mathbf{c}))\|^2]. \quad (11)$$

Simplifying into Average Condition. For the auxiliary exclusion score in Eq. (6), it's laborious to calculate the noise predictions of $N - 1$ auxiliary conditions in every inference step. To simplify, we substitute with an average condition c_Δ^{aver} , which is derived by averaging the semantic vectors of the auxiliary instances:

$$c_\Delta^{aver} = \frac{1}{K} \sum_{i=1}^K c_{\Delta, i}^{aux} = \frac{1}{K} \sum_{i=1}^K \phi(\mathbf{h}_i^{aux}, \mathbf{c}). \quad (12)$$

Intuitively, this average condition indicates the center of all the auxiliary sub-clusters and repels the denoising trajectories away. The denoising loss is also used in this condition:

$$\mathcal{L}_{aver}(\phi) = \mathbb{E}_{t \in [1, T], \mathbf{z}_0^{aver}, \epsilon_t} [\|\epsilon_t - \epsilon_{\theta, \phi}(\mathbf{z}_0^{aver}, t, E_t(\mathbf{p}), c_\Delta^{aver})\|^2]. \quad (13)$$

Note that the data \mathbf{z}^{aver} is randomly chosen from target and auxiliary set. This average condition will act as the empty condition of classifier-free guidance [14] later in Sec. 3.3. During tuning, we only fire the projector from scratch and freeze all other components. As shown in Fig. 2(c), the light-weight projector comprises a ResNet [11] network, linear and AdaIN [16] layers. The complete tuning objective consists of the above 3 losses weighted by hyper-parameters λ_1, λ_2 :

$$\mathcal{L}(\phi) = \mathcal{L}^{tar}(\phi) + \lambda_1 \mathcal{L}^{aux}(\phi) + \lambda_2 \mathcal{L}^{aver}(\phi). \quad (14)$$

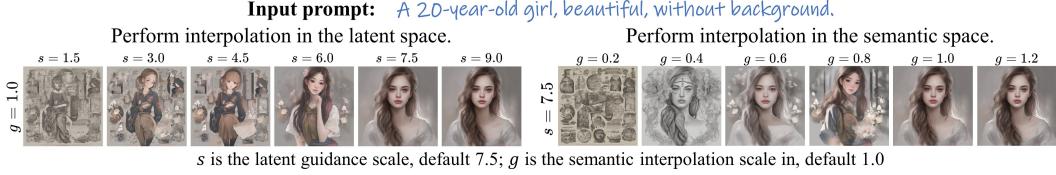


Figure 3: The observation of semantic-latent guidance equivalence property. We vary the latent guidance scale on the left side and the semantic interpolation scale on the right, respectively. The semantic and latent manipulations show the same effect, which proves our argument.

3.3 Inference Strategies for Versatile Generation

In the inference stage, traditional tuning-based pipelines exhibit the imbalance between consistency and diversity and fail to extend to multiple subjects generation. To address these challenges, we propose a semantic interpolation strategy and a latent guidance scale strategy to achieve more nuanced and controlled inference. We further develop two variants to perform multiple subjects generation.

Proof and Implementation of Semantic Interpolation. Classifier-free guidance [14] proves that the conditional interpolation and extrapolation in the latent space, controlled by a guidance scale, reflect the same linear effect in the image space. We argue that the semantic space of the prompt embeddings also shares this property. To prove it, we carry out a toy experiment on SDXL [28]. Since the empty condition c_\emptyset is also a semantic embedding obtained by encoding an empty string, we perform interpolation in the semantic space: $c' = c_\emptyset + g \cdot (c - c_\emptyset)$. Then we replace c with c' in the inference process to generate the images in Fig. 3. As shown, for the latent interpolation (left), when we increase the guidance scale s in Eq. (3) within a proper range, the generated image will be more and more compliant with the prompt. While for the semantic interpolation (right), with g increasing, we can observe the same liner reflection in the image space. This proves that the semantic space also possesses the ability of guidance interpolation. We believe that this is because the semantic and latent space are entangled by the denoising network and thus share some properties in common. With this key insight, we can precisely control the offset guidance effect of c_Δ by replacing Eq. (8) with:

$$c'_b = c_b + v \cdot c_\Delta, \quad (15)$$

where v is a controllable semantic scale. We further investigate its effect in Appendix A.4. In short, a moderate v leads to an optimal balance of consistency and diversity.

Sample with Cluster-Guided Score. Given a subject-centric prompt p^{sub} in the inference stage, we first input its semantic embedding c^{sub} to the tuned projector ϕ^* for two representations:

$$c_\Delta^{tar} = \phi^*(\mathbf{h}^{tar}, c^{sub}), \quad (16)$$

$$c_\Delta^{aver} = \frac{1}{N-1} \sum_{i=1}^{N-1} \phi^*(\mathbf{h}_i^{aux}, c^{sub}). \quad (17)$$

Then we approximate the N-1 auxiliary predictions in Eq. (6) with one prediction under the average condition:

$$\sum_{i=1}^{N-1} [\epsilon_{\theta, \phi^*}(\mathbf{z}_t, t, c^{sub}, c_{\Delta, i}^{aux}) - \epsilon_{\theta}(\mathbf{z}_t, t, c_\emptyset)] \approx \epsilon_{\theta, \phi^*}(\mathbf{z}_t, t, c^{sub}, c_\Delta^{aver}) - \epsilon_{\theta}(\mathbf{z}_t, t, c_\emptyset), \quad (18)$$

where the cluster-conditioned terms are written in the form of Eq. (9). With this approximation, Eq. (6) can be further transformed into:

$$\begin{aligned} \epsilon_{\theta}(\mathbf{z}_t, t, c_\emptyset) + \eta_1 \cdot [\epsilon_{\theta, \phi^*}(\mathbf{z}_t, t, c^{sub}, c_\Delta^{tar}) - \epsilon_{\theta}(\mathbf{z}_t, t, c_\emptyset)] \\ - \eta_2 \cdot [\epsilon_{\theta, \phi^*}(\mathbf{z}_t, t, c^{sub}, c_\Delta^{aver}) - \epsilon_{\theta}(\mathbf{z}_t, t, c_\emptyset)]. \end{aligned} \quad (19)$$

Now we are ready to sample consistent images with the cluster-guided score of Eq. (19). Since η_1 and η_2 play the same role as s in Eq. (3), we can manipulate them for an optimal performance. For large-quantity image generation, we can set $\eta_2 = 0$ to avoid inference time increase at the expense of a slight quality decrease. Besides, since our method acts by guiding the latent code to a specific

cluster step by step, we can loosen the restrictions by applying cluster guidance to certain steps instead of all steps. The effects of these inference strategies are also discussed in Appendix A.4.

Generating Multiple Subjects via Two Variants. To extend the generation of single subject to multiple subjects, we develop two different variants. In the first variant, we treat the multi-subject generation as a joint distribution problem and assume that an image of two subjects belongs to the intersection sub-cluster $\mathcal{S}^{\text{tar}} = \mathcal{S}_1^{\text{tar}} \cap \mathcal{S}_2^{\text{tar}}$. Given multiple target prompts $\mathcal{P} = \{p_j^{\text{tar}}\}_{j=1}^L$ (e.g. *{a hobbit with robes, a white dog}*), we combine the prompts together to one prompt p^{tar} (e.g. *a hobbit with robes and a white dog*) containing multiple base words. We generate multi-subject images with p^{tar} and modify the projector to output multiple semantic vectors $\mathcal{C} = \{c_{\Delta,j}\}_{j=1}^L$. Thus, we can run the tuning and inference process above to generate consistent multi-subject images. However, with all subjects fixed from the beginning, this variant is unable to perform continual subject addition.

In the second variant, for each p_j^{tar} , we treat it as an individual and perform single-subject tuning for multiple projectors ϕ_j^* . Then given multi-subject-centric prompt during inference, we offset the embeddings of multiple base words with the semantic outputs $\mathcal{C} = \{c_{\Delta,j}\}_{j=1}^L$ of the projectors. To avoid intersecting leakage of semantic information, we constrain the effect of every $c_{\Delta,j}$ to the subject’s corresponding spatial mask in the denoising process. Since all subjects are independent in this variant, we can continually add new subjects without affecting existing ones.

4 Experiments

4.1 Baselines

To evaluate the performance of OneActor, we compare it with a wide variety of baselines: (1) tuning-based personalization pipelines, Textual Inversion (TI) [8] and DreamBooth (DB) [33] in a LoRA [15] manner; (2) encoder-based personalization pipelines, IP-Adapter (IP) [42], BLIP-Diffusion (BL) [20] and ELITE (EL) [39]; (3) consistent subject generation pipelines, TheChosenOne (TCO) [4] and ConsiStory (CS) [37]. We will denote them by the abbreviations for simplification. For personalization pipeline, we first generate one target image with the target prompt. Then we use the target image as the input to perform the personalization. Our method and the tuning-based pipeline are implemented on SDXL [28]. Note that TCO and CS have not offered official open-source codes, so we compare to the results and illustrations from their written materials for fairness. We also construct 3 ablation models: original SDXL, our model excluding $\mathcal{L}_{\text{aver}}$ and our model excluding \mathcal{L}_{aux} , $\mathcal{L}_{\text{aver}}$. More implementation details and ablation study are presented in Appendix A.3 and H.

4.2 Qualitative Illustration

Single Subject Generation. We illustrate the visual results of personalization baselines and our OneActor in Fig. 4. As shown, TI generates high-quality and diverse images but fails to maintain the consistency with the target image. With global tuning, DB shows, albeit a certain degree of consistency, limited prompt conformity. IP maintains consistency in some cases, but is unable to adhere closely to the prompt. BL suffers from inferior quality during generation. In contrast, benefiting from the intricate cluster guidance and the undamaged denoising network, our method demonstrates a balance of superior subject consistency, content diversity as well as prompt conformity. We specifically compare our OneActor with other consistent subject generation pipelines. Fig. 5 shows that all pipelines manage to generate consistent and diverse images. Yet our method preserves more consistent details (e.g. the gray shirt of the man, the green coat of the boy). Additionally, we illustrate more examples in Fig. 6, which reaffirms the excellent performance of our method.

Multiple Subjects Generation. We show double subjects generation comparison between baselines and the two variants of our OneActor in Fig. 7. Most pipelines, represented by DB and TCO, fail to perform multiple subjects generation. CS extends generation of single subject to multiple subjects. However, CS shows mediocre layout diversity, probably due to their unstable SDSA operations [37]. In comparison, both variants of our method excellently maintain the appearances of two subjects and display superior image diversity. In general, our method is able to generate different types of subjects including humans, animals and objects (or their combinations). Also, our method is versatile for arbitrary styles and expressions ranging from realistic to imaginary. More qualitative results including triple subjects generation and amusing applications are illustrated in Appendix B and C.



Figure 4: The qualitative comparison between personalization pipelines and our OneActor. TI lacks consistency, while DB and IP exhibit limited prompt conformity and diversity. BL suffers from poor quality in certain cases. In contrast, our method shows superior consistency, diversity as well as stability. Target prompts and base words are marked blue and red, respectively.



Figure 5: The qualitative comparison between consistent subject generation methods. Though all methods generate consistent images given different prompts, our OneActor refines more details such as the characters' clothes.

4.3 Quantitative Evaluation

Metrics. To provide an objective assessment, we carry out comprehensive quantitative experiments. We incorporate multi-objective optimization evaluation and encompass two objectives, identity consistency and prompt similarity, which are popular in personalization tasks [8, 33]. We instruct ChatGPT [25] to generate subject target prompts and templates. For a fair and comprehensive comparison, we adopt two metric settings from the two consistent subject generation pipelines [4, 37]. For the first setting of [4], the identity consistency is CLIP-I, the normalized cosine similarity between the CLIP [29] image embeddings of generated images and that of target image, while the prompt similarity is CLIP-T, the similarity between the CLIP text embeddings of the inference prompts and the CLIP image embeddings of the corresponding generated images. For the second setting of [37], the identity consistency is the DreamSim score [7] which focuses on the foreground subject and the prompt similarity is the CLIP-score [12].



Figure 6: More qualitative results of OneActor. Our method demonstrates excellent subject consistency across characters, animals and objects.

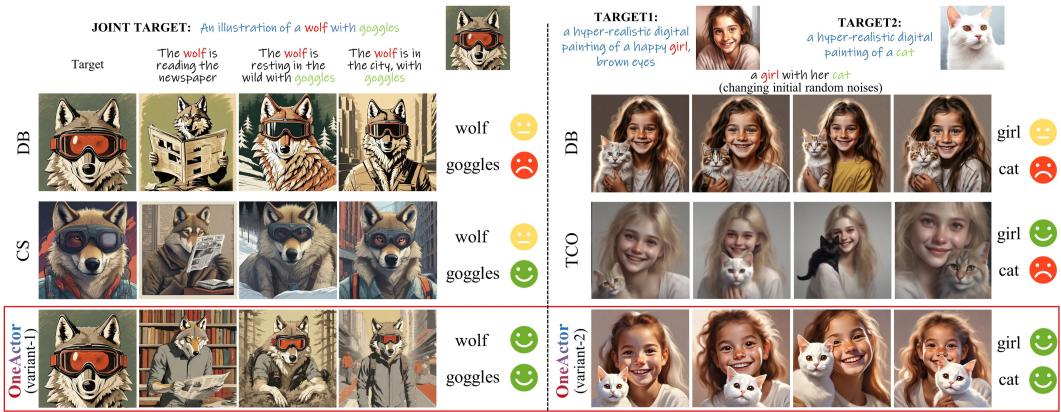


Figure 7: Double subjects generation comparison between baselines and the two variants of our OneActor. Only ConsiStory and our method are able to maintain consistency of multiple subjects.

Results. In Fig. 8(a), we can observe that the baselines originally form a Pareto front in the first setting. To specify, DB exhibits highest identity consistency at the cost of lowest prompt similarity and diversity. Quite the contrary, TI and BL sacrifice identity consistency for prompt similarity. TCO and IP form a moderate region with balanced identity consistency and prompt similarity. In contrast, since our method makes full use of the capacity of the original diffusion model instead of harming it, it achieves a significant boost of prompt similarity as well as satisfactory identity consistency. Note that our method Pareto-dominates IP, TI and BL. In Fig. 8(b), the original Pareto front is determined by IP and three variants of CS in the second setting. Our method pushes the front forward and shows dominance over EL, TI, DB and two variants of CS. In summary, whether in the first or second setting, our OneActor establishes a new Pareto front with a significant margin. These results correspond to the qualitative illustrations above. More experiments are shown in Appendix A.

5 Related Work

Consistent Text-to-Image Generation. A variety of works aim to overcome the randomness challenge of the diffusion models and generate consistent images of the same subject. To start with, *personalization*, given several user-provided images of one specific subject, aims to generate consistent images of that subject. Some works [8, 38] optimize a semantic token to represent the subject, while others [33, 3] tune the entire diffusion model to learn the distribution of the given subject. Later, [1, 10, 18, 36] find that tuning partial parameters of the diffusion model is effective enough. Apart from the tuning-based methods above, encoder-based methods [9, 5, 2, 42, 39] carefully encode the given subject into a representation and manage to avoid user-tuning. While *multi-concept composition* [19, 6] goes a step further to display multiple custom subjects on the same generated image via specially designed attention mechanisms. However, depending on external given images, they fail to generate imaginary or novel subjects. Recently, a new *consistent subject generation* [4] task is proposed, which aims to generate images of one subject given only its descriptive prompts.

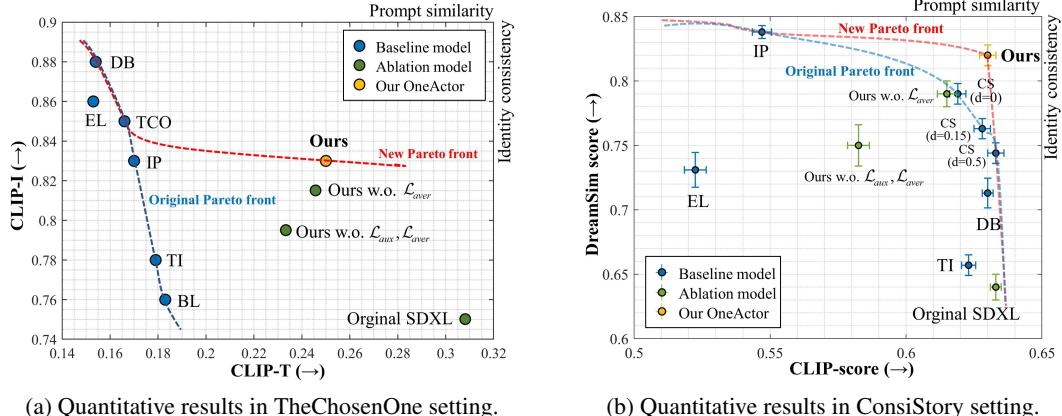


Figure 8: The quantitative comparison between baselines and our OneActor in (a) TheChosenOne setting and (b) ConsiStory setting. Either way, our method establishes a new Pareto front with superior subject consistency and prompt conformity.

Though, its tuning of the whole diffusion model is expensive and may degrade the generation quality. Later, a new approach [37] designs handicraft modules to avoid the tuning process and extends consistent generation from single subject to multiple subjects. Yet the extra modules significantly increase the inference time of every image. Besides, based on the localization of the subject, it is incapable of abstract subject generation like artistic style. For these issues, we design a new paradigm from a clustering perspective. Learning only in the semantic space, it requires shorter tuning without necessarily increasing inference time. Compared to the training-free pipelines, our pipeline can be naturally utilized to pre-train a consistent subject generation network from scratch.

Semantic Control of Text-to-Image Generation Models. The semantic control starts from classifier-free guidance [14], which first proposes a text-conditioned denoising network to perform text-to-image generation. It entangles the semantic and latent space so the prompt can guide the denoising trajectory towards the expected destination. Since then, many works manage to manipulate the semantic space to accomplish various tasks. For personalization, they either learn an extra semantic token to guide the latent to the subject cluster [8] or push the token embedding to its core distribution to alleviate the overfitting [45]. For image editing, they calculate a residual semantic embedding to indicate the editing direction [27, 24]. Besides, previous works [40, 21] utilize the semantic interpolation of two conditional embeddings for a mixed visual effect. These works repeatedly confirm that solely manipulating the semantic space is an effective way to harness diffusion models for various goals. Hence, in our novel cluster-guided paradigm, we transform the cluster information into a semantic representation. This representation will later guide the denoising trajectories to the corresponding cluster.

6 Conclusion

This paper proposes a novel one-shot tuning paradigm, OneActor, for consistent subject generation. Leveraging the derived cluster-based score function, we design a cluster-conditioned pipeline that performs a lightweight semantic search without compromising the denoising backbone. During both tuning and inference, we devise several strategies for better performance and efficiency. Extensive and comprehensive experiments demonstrate that intricate semantic guidance is sufficient to maintain superior subject consistency, as achieved by our method. In addition to excellent image quality, prompt conformity and extensibility, our method significantly improves efficiency, with an average tuning time of just 5 minutes and avoidable inference time increase. Notably, the semantic-latent guidance equivalence property proven in this paper is a potential tool for fine generation control. Furthermore, our method is also feasible to pre-train a consistent subject generation network from scratch, which implements this research task into more practical applications.

Acknowledgements

This work was supported by the National Natural Science Foundation of China under Grant 62192781, Grant 62172326, Grant 62137002 and Grant 62302384, China Postdoctoral Science Foundation under Grant 2023M742790, Research Project Funded by the State Key Laboratory of Communication Content Cognition under Grant No. A202403, and by the Project of China Knowledge Centre for Engineering Science and Technology.

References

- [1] Yuval Alaluf, Elad Richardson, Gal Metzer, and Daniel Cohen-Or. A neural space-time representation for text-to-image personalization. *ACM Transactions on Graphics*, 2023.
- [2] Moab Arar, Rinon Gal, Yuval Atzmon, Gal Chechik, Daniel Cohen-Or, Ariel Shamir, and Amit H. Bermano. Domain-agnostic tuning-encoder for fast personalization of text-to-image models. In *ACM SIGGRAPH Asia*, 2023.
- [3] Omri Avrahami, Kfir Aberman, Ohad Fried, Daniel Cohen-Or, and Dani Lischinski. Break-a-scene: Extracting multiple concepts from a single image. In *ACM SIGGRAPH Asia*, 2023.
- [4] Omri Avrahami, Amir Hertz, Yael Vinker, Moab Arar, Shlomi Fruchter, Ohad Fried, Daniel Cohen-Or, and Dani Lischinski. The chosen one: Consistent characters in text-to-image diffusion models. In *ACM SIGGRAPH*, 2024.
- [5] Wenhui Chen, Hexiang Hu, Yandong Li, Nataniel Ruiz, Xuhui Jia, Ming-Wei Chang, and William W. Cohen. Subject-driven text-to-image generation via apprenticeship learning. In *Advances in Neural Information Processing Systems*, 2023.
- [6] Ganggui Ding, Canyu Zhao, Wen Wang, Zhen Yang, Zide Liu, Hao Chen, and Chunhua Shen. Freecustom: Tuning-free customized image generation for multi-concept composition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [7] Stephanie Fu, Netanel Tamir, Shobhita Sundaram, Lucy Chai, Richard Zhang, Tali Dekel, and Phillip Isola. Dreamsim: Learning new dimensions of human visual similarity using synthetic data. In *Advances in Neural Information Processing Systems*, 2023.
- [8] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit Haim Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. In *International Conference on Learning Representations*, 2023.
- [9] Rinon Gal, Moab Arar, Yuval Atzmon, Amit H. Bermano, Gal Chechik, and Daniel Cohen-Or. Encoder-based domain tuning for fast personalization of text-to-image models. *ACM Transactions on Graphics*, 2023.
- [10] Ligong Han, Yinxiao Li, Han Zhang, Peyman Milanfar, Dimitris N. Metaxas, and Feng Yang. Svdiff: Compact parameter space for diffusion fine-tuning. In *IEEE/CVF International Conference on Computer Vision*, 2023.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [12] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. In *Conference on Empirical Methods in Natural Language Processing*, 2021.
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020.
- [14] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *CoRR*, 2022.
- [15] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022.
- [16] Xun Huang and Serge J. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *IEEE International Conference on Computer Vision*, 2017.
- [17] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloé Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross B. Girshick. Segment anything. In *IEEE/CVF International Conference on Computer Vision*, 2023.

- [18] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [19] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [20] Dongxu Li, Junnan Li, and Steven C. H. Hoi. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. In *Advances in Neural Information Processing Systems*, 2023.
- [21] Zhen Li, Mingdeng Cao, Xintao Wang, Zhongang Qi, Ming-Ming Cheng, and Ying Shan. Photomaker: Customizing realistic human photos via stacked ID embedding. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [22] Haonan Lin, Mengmeng Wang, Jiahao Wang, Wenbin An, Yan Chen, Liu Yong, Feng Tian, Guang Dai, Jingdong Wang, and Qianying Wang. Schedule your edit: A simple yet effective diffusion noise schedule for image editing. *CoRR*, 2024.
- [23] Adyasha Maharana, Darryl Hannan, and Mohit Bansal. Storydall-e: Adapting pretrained text-to-image transformers for story continuation. In *European Conference on Computer Vision*, 2022.
- [24] Thao Nguyen, Yuheng Li, Utkarsh Ojha, and Yong Jae Lee. Visual instruction inversion: Image editing via image prompting. In *Advances in Neural Information Processing Systems*, 2023.
- [25] OpenAI. Chatgpt. <https://chat.openai.com/>, Accessed: 2023-10-15.
- [26] Maxime Oquab, Timothée Darcret, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael G. Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jégou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision. *CoRR*, 2023.
- [27] Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, and Jun-Yan Zhu. Zero-shot image-to-image translation. In *ACM SIGGRAPH*, 2023.
- [28] Dustin Podell, Zion English, Kyle Lace, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: improving latent diffusion models for high-resolution image synthesis. *CoRR*, 2023.
- [29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, Proceedings of Machine Learning Research, 2021.
- [30] Tanzila Rahman, Hsin-Ying Lee, Jian Ren, Sergey Tulyakov, Shweta Mahajan, and Leonid Sigal. Make-a-story: Visual memory conditioned consistent story generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [31] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, 2015.
- [33] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [34] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, 2015.

- [35] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- [36] Yoad Tewel, Rinon Gal, Gal Chechik, and Yuval Atzmon. Key-locked rank one editing for text-to-image personalization. In *ACM SIGGRAPH*, 2023.
- [37] Yoad Tewel, Omri Kaduri, Rinon Gal, Yoni Kasten, Lior Wolf, Gal Chechik, and Yuval Atzmon. Training-free consistent text-to-image generation. *ACM Transactions on Graphics*, 2024.
- [38] Yael Vinker, Andrey Voynov, Daniel Cohen-Or, and Ariel Shamir. Concept decomposition for visual exploration and inspiration. *ACM Transactions on Graphics*, 2023.
- [39] Yuxiang Wei, Yabo Zhang, Zhilong Ji, Jinfeng Bai, Lei Zhang, and Wangmeng Zuo. ELITE: encoding visual concepts into textual embeddings for customized text-to-image generation. In *IEEE/CVF International Conference on Computer Vision*, 2023.
- [40] Yuxuan Yan, Chi Zhang, Rui Wang, Yichao Zhou, Gege Zhang, Pei Cheng, Gang Yu, and Bin Fu. Facestudio: Put your face everywhere in seconds. *CoRR*, 2023.
- [41] Mingzhao Yang, Shangchao Su, Bin Li, and Xiangyang Xue. Exploring one-shot semi-supervised federated learning with pre-trained diffusion models. In *AAAI Conference on Artificial Intelligence*, 2024.
- [42] Hu Ye, Jun Zhang, Sibo Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *CoRR*, 2023.
- [43] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *IEEE/CVF International Conference on Computer Vision*, 2023.
- [44] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [45] Xulu Zhang, Xiao-Yong Wei, Jinlin Wu, Tianyi Zhang, Zhaoxiang Zhang, Zhen Lei, and Qing Li. Compositional inversion for stable diffusion models. In *AAAI Conference on Artificial Intelligence*, 2024.

Appendix

Table of Contents

A Extra Experiments	15
A.1 Motivation Verification	15
A.2 Efficiency Analysis	15
A.3 Ablation Study	16
A.4 Parameter Analysis	17
A.5 More Quantitative Evaluation	18
B Applications	19
B.1 Style Transfer	19
B.2 Storybook Generation	19
B.3 Integration with Diffusion Extensions	19
C More Qualitative Results	22
D Limitations	25
E Societal Impacts	25
F Detailed Derivation of Cluster-Conditioned Score Function	25
G Pseudo-Code	26
H Experiment Details	27
H.1 Implement Details	27
H.2 Packages License	27
H.3 Prompts for Evaluation	27

A Extra Experiments

A.1 Motivation Verification

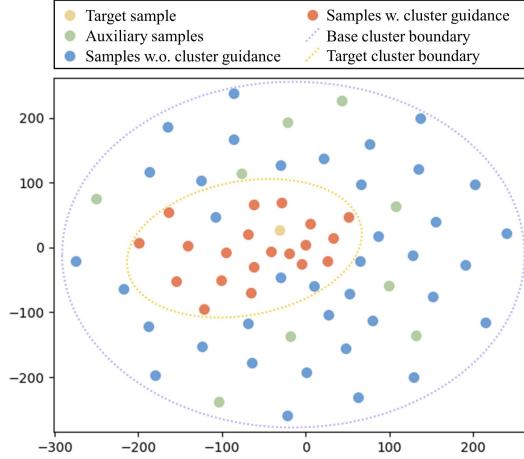


Figure 9: T-SNE illustration of the target sample, auxiliary samples, generated samples with and without our method.

to generate consistent images. In general, the result confirms our motivation and proves an expected function of our cluster guidance.

A.2 Efficiency Analysis

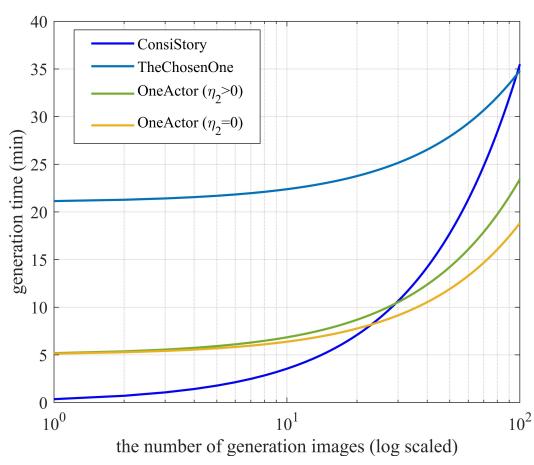


Figure 10: The latency comparison between consistent subject generation baselines and OneActor.

advantages. For example, when generating 100 images, ConsiStory takes about 35 minutes while our method only needs about 23 minutes. If we set $\eta_2 = 0$, with an affordable compromise of generation details, our method will reduce the total time to about 18 minutes. With the generation number increasing, our advantages will continue to grow.

Since our method is built on the clusters of the latent space, we conduct an intuitive experiment to verify our motivation. To specify, We run the whole process to generate samples, choose a target and tune the projector. Then we execute the inference twice, with and without the tuned projector. We collect the latent codes of the target sample, auxiliary samples and two rounds of generated samples. We perform t-SNE analysis to them and illustrate the results in Fig. 9. It clearly demonstrates the cluster structure in the latent space. The target (yellow) and auxiliary (green) samples together form a base cluster (the purple dotted line). Without our cluster guidance, the generated samples (blue) spread out in the base cluster region. In contrast, with our cluster guidance, the generated samples (orange) gather into a sub-cluster (the golden dotted line) centered around the target sample. Notably, some samples without cluster guidance fall into the target sub-cluster, which verifies that the original model has the potential

To evaluate the efficiency of consistent subject generation pipelines, we illustrate the generation latency of each pipeline in Fig. 10. The generation latency consists of the average tuning time and inference time. On a single NVIDIA A100, TheChosenOne [4] needs 20 minutes of tuning. After that, it takes about 8.3 seconds to generate an consistent image. ConsiStory [37] requires no tuning time, but needs about 21.3 seconds to sample one image. While our OneActor needs averagely 5 minutes of tuning. With $\eta_2 = 0$, the inference time is 8.3 seconds, which is the same as the original model. When $\eta_2 > 0$ and applying cluster guidance to steps 1-20, the inference time will rise to about 11.0 seconds. From a user’s perspective, within 20 images, ConsiStory is the fastest, followed by our OneActor and then TheChosenOne. However, considering the application scenarios, the user is very likely to have a demand for large-quantity generation. In this case, our method demonstrates definite

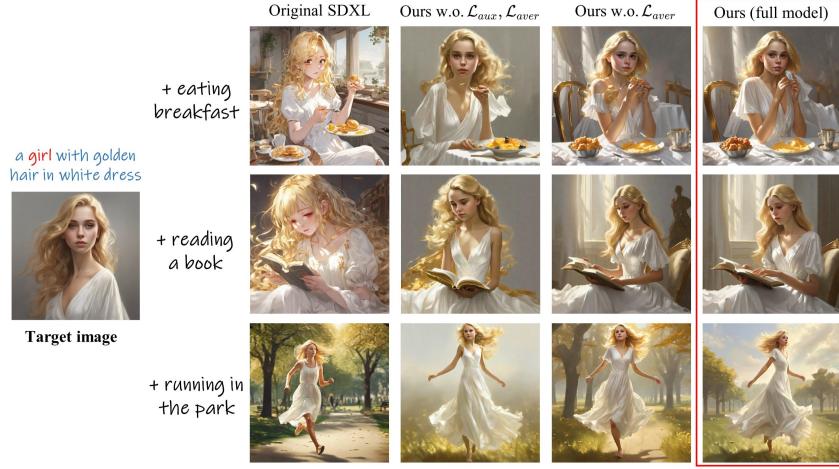


Figure 11: Illustration of the component ablation study. It shows that the proposed objective component significantly contribute to the enhanced character consistency and content diversity.



Figure 12: Illustration of the effect of semantic scale v . We find $v = 0.8$ is optimal because larger v damages the content diversity while smaller v degrades the character consistency.

A.3 Ablation Study

The key to our tuning method lies in the three objectives in Sec. 3.2. To evaluate the effect of every functional component, we carry out a series of ablation experiments. We also construct three ablation models: original SDXL, our model excluding \mathcal{L}_{aver} and our model excluding \mathcal{L}_{aux} , \mathcal{L}_{aver} . We run the process on these models to obtain character-centric images. For a more intuitive demonstration, we input the same initial noise and prompt to all the models and show the visual results in Fig. 11. We can observe that, first, without \mathcal{L}_{aux} and \mathcal{L}_{aver} , the model can maintain consistency to a certain extent. However, it exhibits serious bias issue and damages the diversity of the generated images. Next, including \mathcal{L}_{aux} helps stabilize the tuning and significantly improve the image quality and diversity.

Yet the identity consistency is not satisfying. Finally, in the full model, all components add up to function and generate consistent, diverse and high-quality images. The ablation experiments further validate our analysis in Sec. 3.2.

A.4 Parameter Analysis

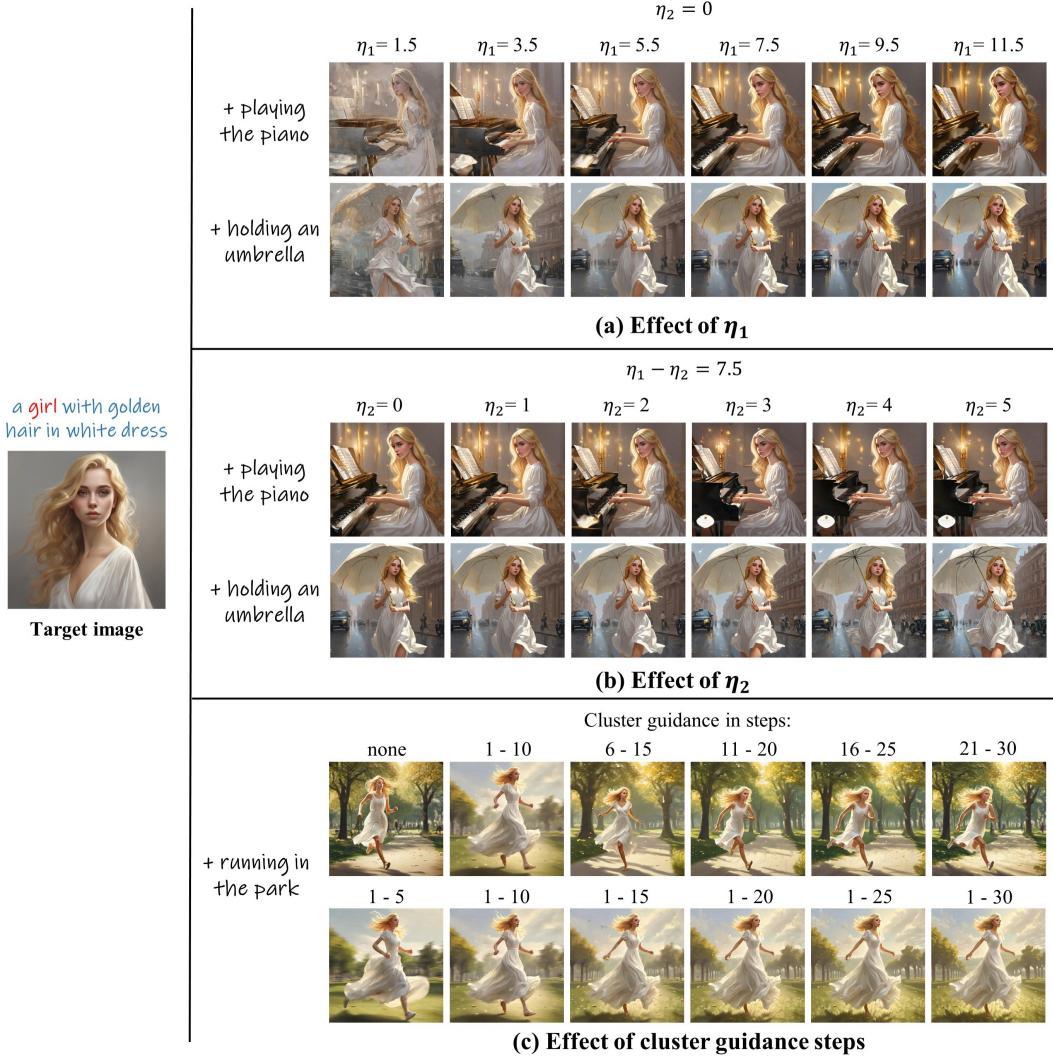


Figure 13: Illustration of parameter analysis of (a) guidance scale η_1 , (b) guidance scale η_2 and (c) cluster guidance steps.

Semantic Interpolation Scale. We specially add semantic interpolation to enhance the flexibility and controllability during inference in Sec. 3.3. To investigate their influence on the generations, we conduct a parameter analysis of semantic scale v . As shown in Fig. 12, with v increasing, the generated images progressively become more and more consistent with the target image. However, too large v may distort the semantic embeddings and thus damage the content diversity. Hence, we set v to 0.8 for a moderate performance. It's worthy noting that the results correspond with the experiments in Sec. 3.3, which proves that the interpolation is feasible to the learned Δc .

Cluster Guidance Scale. Since η_1 and η_2 in Eq. (19) are two crucial guidance scales for consistent subject generation, we conduct experimental analysis on them. In Fig. 13a, we set η_2 to 0 for an isolated experiment of η_1 . We can observe that when η_1 increases, the image aligns more with prompt descriptions. Hence, η_1 acts like the conditional scale s of the original model in Eq. (3). We use $\eta_1 = 7.5$ as the standard setting. In Fig. 13(b), we keep $\eta_1 - \eta_2 = 7.5$ and change η_2 to explore its

effect. It turns out that the performance under $\eta_2 = 0$ is already satisfying. As η_2 grows, the images exhibit more consistent details such as the facial attributes. These results align with our motivation and designs. It also proves that η_1 and η_2 can be manipulated to achieve our intended effect.

Cluster Guidance Steps. During inference, our cluster guidance functions like a navigator. Since the inference is step-wise, we change the inference steps to which we apply the cluster guidance, to explore the influence. Fig. 13(c) shows one interesting fact of the denoising process that, applying guidance to early steps plays a more significant role. It indicates that the early steps control the primary direction of the trajectory and determine which sub-cluster the trajectory falls in. Its influence is so deterministic that guidance in later steps can hardly change it. Hence, if not specified, we apply cluster guidance to steps 1-20.

A.5 More Quantitative Evaluation

Method	Subject consistency		
	DINO-fg(\uparrow)	CLIP-I-fg(\uparrow)	LPIPS-fg(\uparrow)
Textual Inversion	0.452 \pm 0.130	0.644 \pm 0.082	0.322\pm0.061
DreamBooth-LoRA	<u>0.735\pm0.078</u>	0.833\pm0.091	0.296 \pm 0.064
BLIP-Diffusion	0.668 \pm 0.118	0.750 \pm 0.067	<u>0.308\pm0.035</u>
IP-Adapter	0.655 \pm 0.099	0.698 \pm 0.073	0.289 \pm 0.045
OneActor(ours)	0.756\pm0.073	0.821 \pm 0.089	0.299 \pm 0.057

Table 1: Subject consistency scores of the baselines and OneActor. The best and second-best results are marked bold and underlined, respectively.

Method	Prompt similarity		Background diversity		
	CLIP-T-score(\uparrow)	DINO-bg(\downarrow)	CLIP-I-bg(\downarrow)	LPIPS-bg(\downarrow)	
Textual Inversion	0.654\pm0.069	0.318\pm0.074	0.412\pm0.089	0.467 \pm 0.063	
DreamBooth-LoRA	0.594 \pm 0.058	0.441 \pm 0.051	0.438 \pm 0.041	0.482 \pm 0.097	
BLIP-Diffusion	0.547 \pm 0.074	0.362 \pm 0.063	0.446 \pm 0.064	0.423\pm0.077	
IP-Adapter	0.446 \pm 0.051	0.388 \pm 0.054	0.483 \pm 0.077	0.494 \pm 0.101	
OneActor(ours)	0.620 \pm 0.038	0.358 \pm 0.066	0.426 \pm 0.120	0.453 \pm 0.077	

Table 2: Prompt similarity and background diversity scores of the baselines and OneActor.

To further evaluate the performance, we devise automatic quantitative experiments on the baselines and our OneActor. The basic settings remain the same with Sec. 4. We segment the foregrounds (fg) and backgrounds (bg) of the images by SAM [17]. For the metrics, we generally report the cosine similarity between the visual embeddings of the target image and the generated images. These visual embeddings are extracted by DINO [26], CLIP [29] and LPIPS [44]. We focus on three dimensions of the metrics: (1) subject consistency: we calculate the similarity on the image foregrounds to obtain DINO-fg, CLIP-I-fg and LPIPS-fg; (2) prompt similarity: we calculate the CLIP-T-Score [12] on the whole images; (3) background diversity: we calculate the similarity on the image backgrounds to obtain DINO-bg, CLIP-I-bg and LPIPS-bg. We use the standard deviation as the error bar throughout this paper.

The results are presented in Tabs. 1 and 2. In summary, our method achieves the highest DINO-fg score and the second-highest CLIP-I-fg score, showcasing superior subject consistency. Additionally, our method ranks second across various metrics in prompt similarity and background diversity. These evaluation results highlight our method’s comprehensive performance in generating consistent images.

B Applications

Based on a mild guidance rather than a tuned backbone or attention manipulations, our OneActor can be naturally extended to various application scenarios.

B.1 Style Transfer

For spatial location-based pipelines like ConsiStory [37], the inherent flaw is that they can't deal with abstract subjects (e.g. art styles, brushstrokes, photo lighting). In contrast, relying on the semantic guidance, our OneActor is able to deal with any type of subject as long as it can be described. For example in Fig. 14, we perform style transfer using OneActor. To specify, we first generate an image with a target style with the prompt "*a picture in a style*" and choose *style* as the base word. Then after tuning, we generate styled images with style-centric prompts (e.g. *a fancy car in a style*).

B.2 Storybook Generation

Since OneActor is able to perform multiple subjects generation, we can extend the ability to story visualization. As shown in Fig. 15, we first generation three target characters with three descriptive prompts. Then we perform multiple subjects generation with the second variant of our method. After tuning, the pipeline is ready to generate consistent images of one specific character or the combination of the three characters. By describing the plots with the base words, we can obtain a consistent story revolving around the three characters. Note that different from traditional story visualization pipelines, the characters in our story are arbitrary generations by diffusion models with prompts.

B.3 Integration with Diffusion Extensions

Avoiding tuning of the denoising backbone, our method is naturally compatible with diffusion-based extensions. Here we show the integrations of StableDiffusion V1.5, StableDiffusion V2.1 and ControlNet XL in Fig. 16. The illustrations show that our method is robust across different versions of diffusion models. Besides, combined with ControlNet, our method is able to generate consistent images given different pose conditions, which highlights the compatibility of our method.

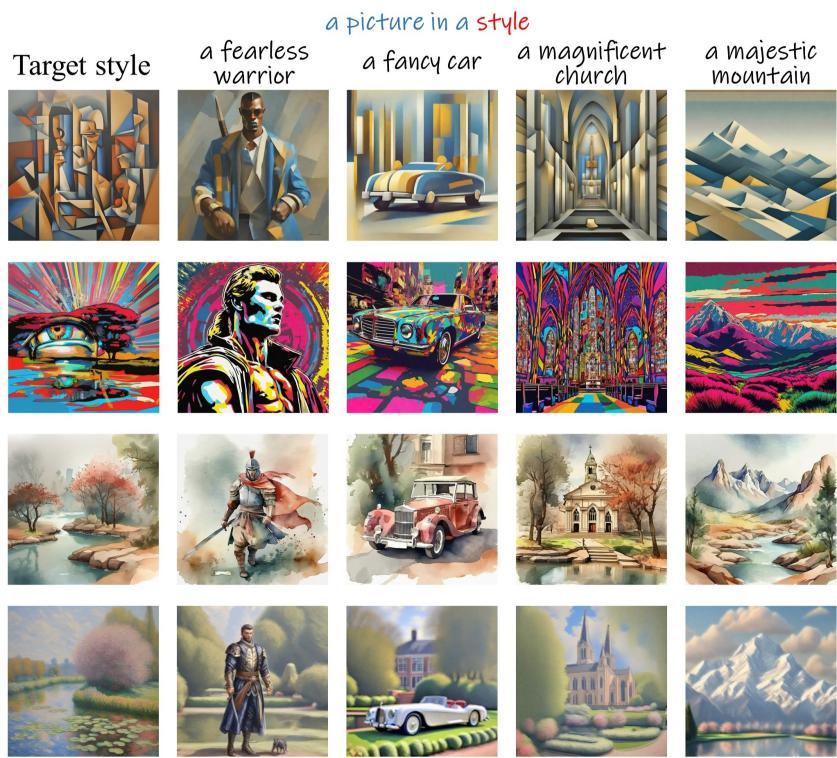


Figure 14: Style transfer using our OneActor. We generate the target image with a prompt and choose "style" as the base word. Then we perform our method to generate images in a consistent style.

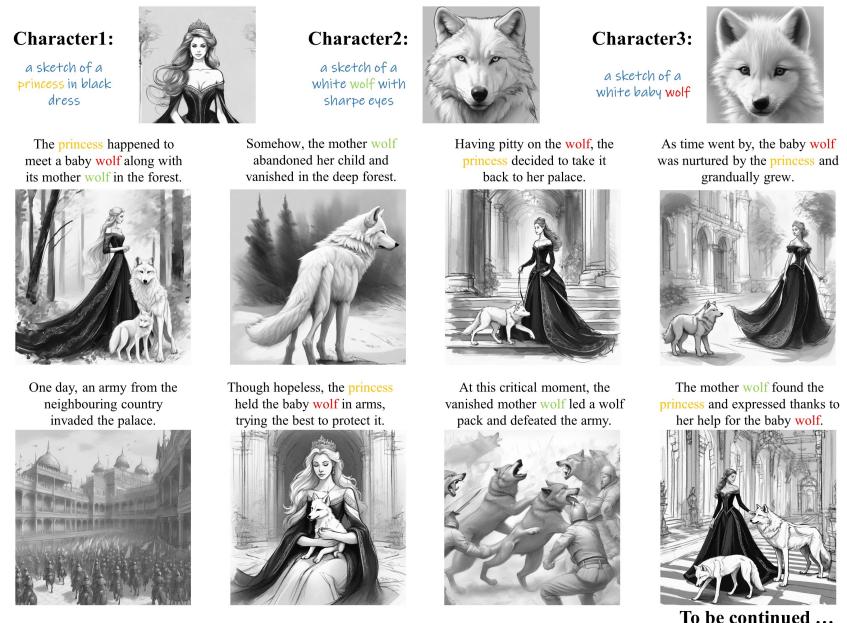


Figure 15: Story Visualization using our OneActor. We generate 3 target characters with prompts and perform the second variant of our method for multiple subjects generation. Different target characters are marked with different colors.

OneActor \times Stable Diffusion v1.5



OneActor \times Stable Diffusion v2.1



OneActor \times ControlNet-SDXL

Input different poses

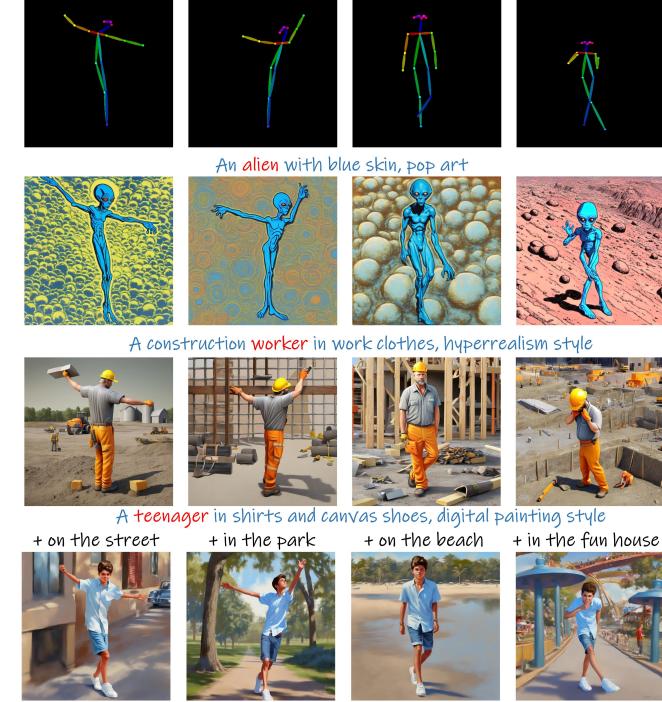


Figure 16: Integration results of our OneActor and several diffusion-based extensions. The results demonstrate the robust adaptability and compatibility of our method.

C More Qualitative Results

In order to comprehensively showcase the performance of our method, we provide more generation samples covering single subject, double subjects and triple subjects in Figs. 17, 18 and 19.

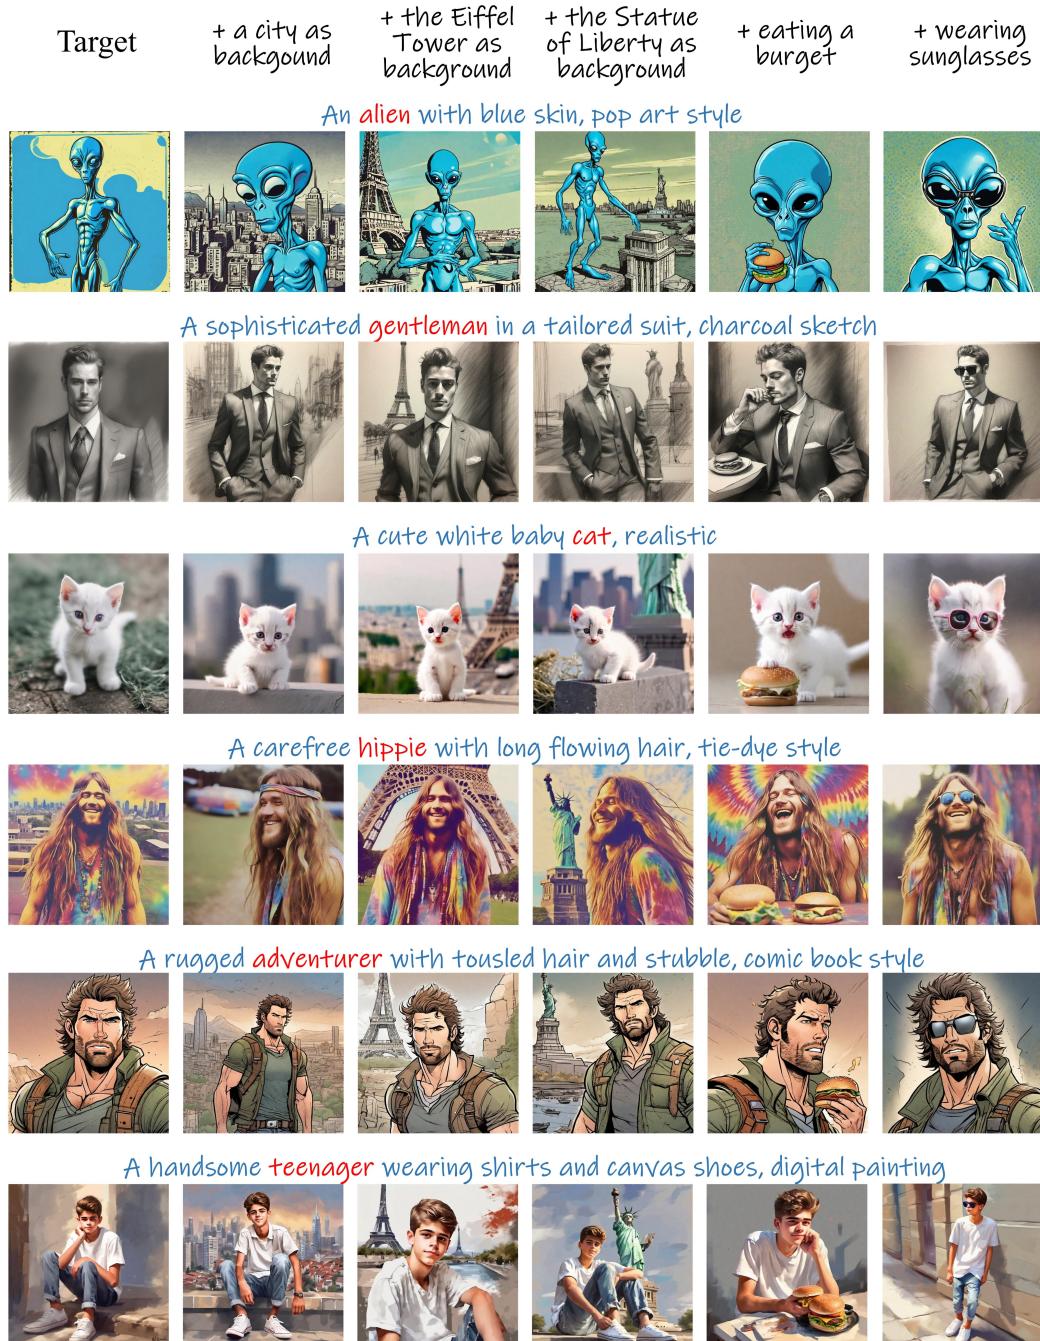


Figure 17: Illustrations of single subject generation.

boy & dog: Neonpunk style of a boy with black spiky hair, with his dog

+ sitting in the house

+ walking in the park

+ playing with a ball

+ walking in the snow



dog & hat:

+ eating food with a hat

+ sitting on the mat, with a hat

+ jumping over a puddle

+ walking on the street with a hat



wizard & warrior:

An old wizard and a young boy warrior, digital painting

+ walking in the snow

+ climbing on top of the mountain

+ resting inside the desert

+ cautiously get through a forest



explorer & lion:

+ getting off from his car

An watercolor style of an wild explorer with blue coat

+ drinking water near a lake

+ meeting a big lion

+ giving water to the lion



Figure 18: Illustrations of double subjects generation.

A digital painting of a happy girl, with a black dog and a white cat

+ holding together



+ sitting on the street



+ playing in the garden



+ walking on the street



A realistic photo of a man in black suit

+ drinking beer in a villa



+ walking on the beach with sunglasses



+ running in rain with sunglasses and an umbrella



+ sitting in a café with the umbrealla



A watercolor image of a mother with her son and daughter

+ playing in the garden



+ eating dinner in the house



+ camping besides a lake



+ walking on the grassland



A manga page of a warrior

+ fighting with a monster barehanded



+ holding a sword under the sun



+ walking on the road with the sword and a shield



+ resting while holding the shield



Figure 19: Illustrations of triple subjects generation.

D Limitations

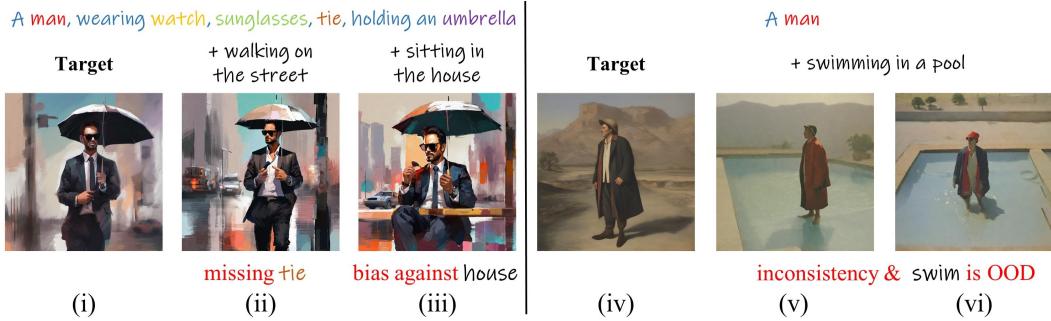


Figure 20: Limitations of OneActor. Our method struggles with generating numerous subjects on the same image and is restricted by the latent properties of diffusion models, causing bias and OOD.

Despite the above mentioned strengths, our method exhibits several limitations. One the one hand, though able to generate multiple subjects, our method struggles to simultaneously generate numerous subjects (usually > 4) on one image. As shown in Fig. 20(ii), in the 5 subjects generation case, the subject *tie* is neglected. On the other hand, based on a guidance manner, our method is affected by the latent properties of the original diffusion models. First, the latent clusters are entangled, which leads to inevitable bias. For the example in Fig. 20(iii), the *umbrella* causes bias against indoor environment, resulting in a mismatch with prompts. Second, as illustrated in Figs. 20(v) and (vi), if the target prompt is *a man*, its base cluster region will expand so further that it's difficult to precisely navigate to the target sub-cluster, causing inconsistency of some details. Yet, this problem can be solved by using more detailed prompts. Third, a sub-cluster distribution is rigid and incomplete, causing out of distribution (OOD) problem. In Figs. 20(vi) and (vii), *swim* is not within the scope of the target sub-cluster, which leads to mismatches with prompts. Addressing these challenging issues will be the focus of our future work.

E Societal Impacts

The advancements in consistent subject generation present significant societal impacts across various domains. Our work, which enhances subject consistency and efficiency in text-to-image generation, democratizes artistic creation, enabling artists and designers of all skill levels to produce consistent and high-quality visual content. It streamlines creative processes in animation, advertising and publishing, leading to substantial time and cost savings. The ability to maintain visual consistency in media and entertainment may revolutionize animation and storybook illustration.

Yet, the rapid development of creative diffusion models may lead to job displacement for artists and designers who rely on traditional methods, creating economic challenges and necessitating reskilling. Besides, ethical concerns arise around the misuse of generated content, such as the creation of deepfakes or misleading images, which can harm public trust and spread misinformation. Additionally, ensuring that generative models are trained on unbiased datasets is crucial to prevent reinforcing harmful stereotypes or biases. For these issues, since our model is based on the Stable Diffusion, we will inherit the safeguards such as the NSFW detection and the usage guideline.

F Detailed Derivation of Cluster-Conditioned Score Function

Given N initial noises and the same subject prompt, generations of ϵ_θ fail to reach one specific sub-cluster, but spread to a base region \mathcal{S}^{base} of different sub-clusters. If we choose one sub-cluster as the target \mathcal{S}^{tar} and denote the rest as auxiliary sub-clusters \mathcal{S}_i^{aux} , then $\mathcal{S}^{base} = \mathcal{S}^{tar} \cup \{\mathcal{S}_i^{aux}\}_{i=1}^{N-1}$. The key to consistent subject generation is to guide ϵ_θ towards the expected target sub-cluster \mathcal{S}^{tar} . From a result-oriented perspective, we expect to increase the probability of generating images of the target sub-cluster \mathcal{S}^{tar} and reduce that of the auxiliary sub-clusters \mathcal{S}_i^{aux} . Thus, if we consider the original

diffusion process as a prior distribution $p(\mathbf{x})$, our expected distribution can be denoted as:

$$p(\mathbf{x}) \cdot \frac{p(\mathcal{S}^{tar} | \mathbf{x})}{\prod_{i=1}^{N-1} p(\mathcal{S}_i^{aux} | \mathbf{x})}, \quad (20)$$

We take the negative gradient of the log likelihood to derive:

$$-\nabla_{\mathbf{x}} \log p(\mathbf{x}) - \nabla_{\mathbf{x}} \log p(\mathcal{S}^{tar} | \mathbf{x}) + \sum_{i=1}^{N-1} \nabla_{\mathbf{x}} \log p(\mathcal{S}_i^{aux} | \mathbf{x}). \quad (21)$$

From the score function-based perspective [35], Eq. (21) includes one unconditional score and one target conditional scores and $N - 1$ auxiliary scores. With the reparameterization trick of \mathbf{x} in terms of ϵ -prediction [13]: $\mathbf{x}_{\theta}(\mathbf{x}_t) = (\mathbf{x}_t - \sigma_t \epsilon_{\theta}(\mathbf{x}_t, t)) / \alpha_t$, where $t \in [0, T]$, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and α_t, σ_t are a set of constants, the unconditional can be estimated by:

$$\epsilon_{\theta}(\mathbf{x}_t, t) \approx -\sigma_t \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t). \quad (22)$$

We follow [14] to approximate the conditional scores in a classifier-free manner:

$$\epsilon_{\theta}(\mathbf{x}_t, t, \mathcal{S}) - \epsilon_{\theta}(\mathbf{x}_t, t) \approx -\eta \sigma_t \nabla_{\mathbf{x}_t} \log p(\mathcal{S} | \mathbf{x}_t), \quad (23)$$

where η is a guidance control factor. Then Eq. (21) can be transformed into:

$$\epsilon_{\theta}(\mathbf{x}_t, t) + \eta_1 \cdot [\epsilon_{\theta}(\mathbf{x}_t, t, \mathcal{S}^{tar}) - \epsilon_{\theta}(\mathbf{x}_t, t)] - \eta_2 \cdot \sum_{i=1}^{N-1} [\epsilon_{\theta}(\mathbf{x}_t, t, \mathcal{S}_i^{aux}) - \epsilon_{\theta}(\mathbf{x}_t, t)], \quad (24)$$

where η_1, η_2 are two scaled guidance control factors. Since we use a latent diffusion model in this paper, we transfer the above process from the image space to the latent space to derive:

$$\epsilon_{\theta}(\mathbf{z}_t, t) + \eta_1 \cdot [\epsilon_{\theta}(\mathbf{z}_t, t, \mathcal{S}^{tar}) - \epsilon_{\theta}(\mathbf{z}_t, t)] - \eta_2 \cdot \sum_{i=1}^{N-1} [\epsilon_{\theta}(\mathbf{z}_t, t, \mathcal{S}_i^{aux}) - \epsilon_{\theta}(\mathbf{z}_t, t)]. \quad (25)$$

G Pseudo-Code

Algorithm 1: Tuning Process of OneActor

Input :Original diffusion model θ , projector ϕ , target prompt \mathbf{p}^{tar} , base word w , prompt templates \mathcal{P} , max tuning step T

Output :Tuned projector ϕ^*

begin

```

Input  $\mathbf{p}^{tar}$  into  $\theta$  for dataset  $\mathcal{B} = \{\mathbf{z}^{tar}, \mathbf{h}^{tar}\} \cup \{\mathbf{z}_i^{aux}, \mathbf{h}_i^{aux}\}_{i=1}^K$ 
for  $s = 1$  to  $T$  do
    Randomly choose  $\mathbf{p} \in \mathcal{P}$  and fill it with  $w$ 
    Obtain the embedding  $\mathbf{c}$  of  $\mathbf{p}$  using the text encoder of  $\theta$ 
    for  $\mathbf{z}_0$  in  $\{\mathbf{z}^{tar}, \mathbf{z}_i^{aux}\}$  do
         $\mathbf{z}_t \leftarrow \sqrt{\bar{\alpha}_t} \mathbf{z}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$ , where  $t \in [0, T]$ ,  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  and  $\bar{\alpha}_t$  are constants
    end
     $c_{\Delta}^{tar} \leftarrow \phi(\mathbf{h}^{tar}, \mathbf{c})$ 
     $\{c_{\Delta,i}^{aux}\}_{i=1}^K \leftarrow \{\phi(\mathbf{h}^{tar}, \mathbf{c})\}_{i=1}^K$ 
     $c_{\Delta}^{aver} \leftarrow \frac{1}{K} \sum_{i=1}^K c_{\Delta,i}^{aux}$ 
    Calculate loss  $\mathcal{L}(\phi)$  using Eqs. (10)-(14)
     $\mathcal{L}(\phi).backward()$ 
    end
end

```

H Experiment Details

H.1 Implement Details

General Settings. We implement our method based on SDXL, which is a common choice by most of the related works. All experiments are finished on a single NVIDIA A100 GPU. All images are generated in 30 steps. During tuning, we generate $N = 11$ base images. We use $K = 3$ auxiliary images each batch. The projector consists of a 5-layer ResNet, 1 linear layer and 1 AdaIN layer. The weight hyper-parameters λ_1 and λ_2 are set to 0.5 and 0.2. We use the default AdamW optimizer with a learning rate of 0.0001, weight decay of 0.01. We tune with a convergence criterion, which takes 3-6 minutes in most cases. During inference, we set the semantic interpolation scale v to 0.8 if not specified. We set the cluster guidance scale η_1 and η_2 to 8.5 and 1.0. We apply cluster guidance to the first 20 inference steps and normal guidance to the last 10 steps.

Multi-Word Subject. Since we apply semantic offset to the base word of the target prompt, a multi-word subject (e.g. *pencil box*) may cause problem. For this case, our method automatically adapts the projector to output multiple offsets for the words, while other parts remain the same.

Subject Mask Extraction. When generating multiple subjects, our method perform mask extraction to the involved subjects. To specify, in each generation step, we collect all the cross-attention maps that correspond to the token of each subject. We apply Otsu’s method to the average of the maps for a binary mask for each subject.

H.2 Packages License

- SDXL [28] implementation at
<https://huggingface.co/stabilityai/stable-diffusion-xl-base-1.0>.
- DreamBooth-LoRA [33] and Textual Inversion [8] implementations on SDXL at
<https://huggingface.co/diffusers>.
- ELITE [39] implementation at
<https://github.com/csyxwei/ELITE>.
- BLIP-diffusion [20] implementation at
<https://huggingface.co/salesforce/blipdiffusion/tree/main>.
- IP-Adapater [42] implementation at
<https://github.com/tencent-ailab/IP-Adapter>.
- CLIP [29], DINO [26] and SAM [17] implementation at
<https://github.com/huggingface/transformers>.
- LPIPS [44] implementation at
<https://github.com/richzhang/PerceptualSimilarity>.

H.3 Prompts for Evaluation

Target prompts, where the base words are denoted bold:

- A sophisticated **gentleman** in a tailored suit, charcoal sketch.
- A graceful **ballerina** in a flowing tutu, watercolor.
- A carefree **hippie** with long flowing hair, tie-dye style.
- A wise **elder** with a wrinkled face and kind eyes, pastel portrait.
- A rugged **adventurer** with tousled hair and stubble, comic book style.
- A trendy **hipster** with thick-rimmed glasses and a beanie, street art.
- A spirited **athlete** in mid-stride, captured in a dynamic sculpture.
- A curious **child** with wide eyes and messy hair, captured in a whimsical caricature.
- A playful **panda** with black patches and fluffy ears, cartoon animation.
- A majestic **tiger** with orange stripes and piercing gaze, photorealistic illustration.

- A graceful **deer** with antlers and doe eyes, minimalist vector art.
- A curious **raccoon** with bandit mask and bushy tail, pastel-colored sticker.
- A wise **owl** with feathers and intense stare, ink brushstroke drawing.
- A sleek **dolphin** with gray skin and playful smile, digital pixel art.
- A mischievous **monkey** with long tail and banana, clay sculpture.
- A colorful **parrot** with vibrant feathers and beak, paper cut-out animation.
- A cuddly **koala** with fuzzy ears and button nose, felt fabric plushie.
- A majestic **eagle** with sharp talons and soaring wings, metallic sculpture.
- A noble **griffin** in a tapestry, its presence commanding attention with intricate detail and grandeur.
- A mischievous **imp** captured in charcoal, its sly grin dancing off the page.
- A radiant **unicorn** in stained glass, casting prisms of light with its ethereal beauty.
- An ominous **kraken** in sculpture, its tentacles frozen in time, threatening the bravest souls.
- A resplendent **phoenix** in metalwork, its fiery wings ablaze with eternal flames.
- An enigmatic **sphinx** in relief, guarding ancient secrets with stoic grace.
- A whimsical **faun** in woodcarving, its playful spirit captured in intricate detail.
- A formidable **werewolf** in digital art, its primal fury palpable yet contained.
- A mystical **mermaid** in underwater mural, her siren song echoing through ocean depths.
- A legendary **dragon** in castle tapestry, its scales shimmering with timeless majesty and power.
- A majestic **lion** in bronze sculpture, its mane flowing with regal power.
- A stealthy **ninja** in ink drawing, moving silently through the shadows.
- A serene **monk** in oil painting, meditating amidst tranquil surroundings.
- A futuristic **robot** in 3D rendering, with sleek lines and glowing eyes.

Subject-centric prompts, where [b] indicates the base word:

- [b], at the beach.
- [b], in the jungle.
- [b], in the snow.
- [b], in the street.
- [b], with a city in the background.
- [b], with a mountain in the background.
- [b], with the Eiffel Tower in the background.
- [b], near the Statue of Liberty.
- [b], near the Sydney Opera House.
- [b], floating on top of water.
- [b], eating a burger.
- [b], drinking a beer.
- [b], wearing a blue hat.
- [b], wearing sunglasses.
- [b], playing with a ball.
- [b], as a police officer.
- [b], on a rooftop overlooking a city skyline.
- [b], in a library surrounded by books.
- [b], on a farm with animals in the background.

- [b], in a futuristic cityscape.
- [b], on a skateboard performing tricks.
- [b], in a traditional Japanese garden.
- [b], in a medieval castle courtyard.
- [b], in a science laboratory conducting experiments.
- [b], in a yoga studio practicing poses.
- [b], at a carnival with colorful rides.
- [b], in a cozy coffee shop sipping a latte.
- [b], in a sports stadium cheering for a team.
- [b], in a space station observing Earth.
- [b], in a submarine exploring the depths of the ocean.
- [b], on a film set directing a movie scene.
- [b], in an art gallery admiring paintings.
- [b], in a theater rehearsing for a play.
- [b], at a protest holding a sign.
- [b], in a classroom teaching students.
- [b], in a space observatory studying the stars.
- [b], walking on the street.
- [b], eating breakfast in the house.
- [b], running in the park.
- [b], reading a book in the library.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: We accurately reflect our contributions and scope in the abstract and introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We discuss the limitations in Appendix [D](#).

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: We provide the theoretical analysis in Appendix F.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We submitted codes in the supplementary materials. We try to describe all the necessary details in Sec. 4 and Appendix H.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We submitted codes in the supplementary materials. We will release the codes officially online with detailed instructions after preparations.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We describe the details in Appendix H. We also submitted codes in the supplementary materials.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We report the standard deviation as the error bars in Sec. 4.3 and Appendix A.5. Partial results are not reported with error bars because some baseline works didn't offer codes or error bars in such setting.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: We provide information of the compute resources in Appendix [H](#).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: We try our best to conform with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[Yes\]](#)

Justification: We discuss the societal impacts in Appendix [E](#).

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [\[Yes\]](#)

Justification: The models will be protected by the safeguards of the Stable Diffusion, as described in Appendix [E](#).

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [\[Yes\]](#)

Justification: We try to credit all the assets in the reference and list the used packages in Appendix [H.2](#).

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](#) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We will provide documentation when we officially release the codes.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.