

XVerse: Consistent Multi-Subject Control of Identity and Semantic Attributes via DiT Modulation

Bowen Chen*, Mengyi Zhao*, Haomiao Sun*, Li Chen*,[†] Xu Wang,
Kang Du, Xinglong Wu
Intelligent Creation Team, ByteDance
{chenbowen.cbw, zhaomengyi.pl, sunhaomiao, chenli.phd, wangxu.ailab,
dukang.daniel, wuxinglong}@bytedance.com



Figure 1: XVerse enables single/multi-subject personalization and the additional control of semantic attributes such as pose, style, and lighting. Input conditions are highlighted with red dots.

Abstract

Achieving fine-grained control over subject identity and semantic attributes (pose, style, lighting) in text-to-image generation, particularly for multiple subjects, often undermines the editability and coherence of Diffusion Transformers (DiTs).

*Equal contribution.

[†]Corresponding author, project lead.

Many approaches introduce artifacts or suffer from attribute entanglement. To overcome these challenges, we propose a novel multi-subject controlled generation model XVerse. By transforming reference images into offsets for token-specific text-stream modulation, XVerse allows for precise and independent control for specific subject without disrupting image latents or features. Consequently, XVerse offers high-fidelity, editable multi-subject image synthesis with robust control over individual subject characteristics and semantic attributes. This advancement significantly improves personalized and complex scene generation capabilities. Project Page: <https://bytedance.github.io/XVerse>; Github Link: <https://github.com/bytedance/XVerse>.

1 Introduction

The field of text-to-image generation has advanced remarkably [1; 2; 3; 4; 5; 6], enabling the creation of highly realistic and diverse images from textual descriptions. Initial breakthroughs in personalization focused on single subjects [7; 8; 9; 10]. These methods demonstrated strong control over individual subject appearances and showed good editability. However, the growing demand for more complex visual narratives and personalized content has spurred interest in extending these capabilities to scenarios involving multiple subjects within a single, coherent scene. This transition to multi-subject personalization presents substantial challenges, particularly in preserving individual identity fidelity and alleviating attribute entanglement.

Many current state-of-the-art multi-subject methods [11; 12; 13; 14] tried to leverage the attention mechanism in Diffusion Transformers (DiTs) [15] for injecting information from reference images. But this direct injection or strong reliance on image features can substantially impact the generation quality of the base model. This often leads to artifacts, distortions, attribute entanglement, and can compromise the overall structural integrity and coherence of the generated image. These limitations highlight a critical gap, underscoring the need for novel techniques that offer fine-grained, independent control over multiple subjects while preserving image quality, editability.

To address these limitations, this paper introduces XVerse, a novel method for consistent multi-subject control of identity and semantic attributes. We identify that the inherent modulation vectors within DiT blocks [15], which are typically employed for general conditioning, represent an underexplored yet highly promising pathway to achieve nuanced, subject-specific control. Building upon this insight, XVerse pioneers an approach centered on learning offsets within the text-stream modulation mechanism of DiTs. With the reference images provided, XVerse utilizes an adapter to transform them into share offsets and per-block offsets for token-specific text-stream modulation. This technique allows for condition injection from diverse reference images while preserving the image's underlying structural integrity. To enhance fine-grained details, we incorporate VAE-encoded image features into the single-stream block of FLUX [6]. Instead of being the main conditioning factor, the VAE-derived features play a supporting role in enhancing details for the backbone network. This strategy successfully minimizes the occurrence of artifacts and distortions, enabling XVerse to achieve exceptional multi-subject controlled generation results (as shown in Figure 1). Extensive testing conducted on our benchmark validates XVerse's exceptional performance in terms of both flexibility of editing and maintaining the appearance of the subject.

Our contributions are summarized as follows:

- We propose XVerse, a novel framework for fine-grained multi-subject controllable generation. Our approach integrates the reference images into the text-stream modulation offsets. Additionally, we leverage VAE-encoded image features to enhance fine-grained details which are difficult to express in the semantic space. This allows XVerse to achieve a high degree of consistency with reference images, while preserving the original diffusion model's editability by maintaining the image composition.
- Through training on a high-quality datasets constructed by our data pipeline, XVerse achieves outstanding performance in generation tasks under multi-subject control. Moreover, due to the flexibility of text-stream modulation, XVerse also demonstrates strong generalization performance in tasks such as maintaining semantic information such as posture, lighting, and background.

- We present a comprehensive benchmark XVerseBench, which evaluates both single-subject and multi-subject controlled image generation. This benchmark provides a rigorous methodology for assessing a model’s ability to support flexible editing, maintain subject characteristics, and preserve distinct identities.

2 Related Work

2.1 Subject-driven Generation

Subject-driven generation tasks aimed at synthesizing user-specific content, and have made significant progress in recent years. Early endeavors predominantly concentrated on single-subject personalization. These approaches can be broadly classified into two main categories: (1) Fine-tuning-based methods, such as Textual Inversion [7] and DreamBooth [8], which adapt pre-trained models to embed novel concepts from a few exemplar images of a single subject. (2) Tuning-free methods, including IP-Adapter [9] and Photoverse [10], leverage powerful vision encoders to inject subject identity directly into the generation process without test-time fine-tuning, thereby offering enhanced flexibility. However, extending personalization to accommodate multiple subjects within a single, cohesive image presents substantial challenges, particularly in preserving individual identity fidelity and mitigating attribute entanglement. Recent progress leveraging DiTs [15] has begun to address these complexities. For instance, OmniControl [11] demonstrated versatile control by conditioning DiT inputs. Concurrently, a line of research has focused on unified frameworks for complex multi-subject image customization, exemplified by UniReal [12], UNO [13], and DreamO [14]. While these methods significantly advance multi-subject generation, their primary strategies involve conditioning the input token sequence, imposing explicit constraints, or modifying attention mechanisms to guide the generative process. XVerse introduces a distinct perspective, operating as a tuning-free method for multi-subject personalization. With the utilization of the text-stream modulation mechanism in DiTs, XVerse can enable precise, subject-specific conditioning while preserving the structural integrity of the generated image

2.2 Modulation in Generative Models

Modulation mechanisms have been instrumental in enhancing the controllability and expressiveness of generative models. This concept involves dynamically adjusting a model’s internal activations or parameters based on conditioning information. StyleGAN [16] pioneered the use of modulation by introducing adaptive instance normalization (AdaIN), which modulates features using a style vector. A key discovery associated with this approach was that even small modifications to these modulation parameters could induce smooth and semantically meaningful perturbations in the generated image. This insight spurred a wave of research focused on leveraging modulation layers, leading to significant advancements and successful applications in tasks such as image editing and manipulation [17; 18; 19; 20; 21].

Transformer-based models, such as DiTs, commonly employ modulation mechanisms like adaptive layer normalization (AdaLN) [22]. These mechanisms allow conditioning information (e.g., text embeddings, timesteps) to guide image generation by modulating normalization layers within transformer blocks. Tokenverse [23] explored modulation for subject preservation, they typically rely on fine-tuning strategies. Such methods necessitate extensive training on predefined datasets to learn subject-specific modulations, thus limiting their adaptability to novel, unseen subjects without retraining or inference-time optimization. In contrast, XVerse is designed to inject rich identity and semantic information for arbitrary subjects without such subject-specific fine-tuning.

3 Method

3.1 Preliminaries

DiTs have introduced significant advancements in the quality, efficiency, and scalability of image synthesis, establishing DiTs as the foundational architecture for most state-of-the-art models. The Attention Block in DiTs process text and image tokens concurrently within their transformer layers. This attention mechanism provides a pathway for injecting control signals, such as features from control images, into the token representations. However, while injecting control signals through




Figure 2: Overview of the XVerse framework. The reference images are processed by a T-Mod Resampler and subsequently injected into the per-token modulation adapter. Additionally, to supplement image details, the VAE-encoded features of the reference image are also utilized as input to the single-stream block of DiTs.

attention can achieve good similarity improvements, this approach can also cause the model's sampling trajectory to deviate, leading to a reduction in image generation quality. In this work, We try to further explore the potential of the modulation mechanism in DiTs to achieve more precise subject control editing.

Modulation Mechanism in DiTs. The modulation mechanism in DiTs, adapted from techniques popularized by StyleGAN[16], refines neural network activations by applying learned scale factors and bias terms. For instance, models like Stable Diffusion 3 [15] and FLUX [6] utilize a Multi-Layer Perceptron (MLP) to process both the diffusion timestep t and a CLIP [24] embedding of the text prompt (e.g., $f_p = \text{CLIP}(p)$), thereby deriving a conditioning vector y :

$$y = \text{MLP}(t, f_p). \quad (1)$$

This vector y is then further processed, typically through a linear layer, to generate a set of modulation parameters. In the DiT architecture, this yields twelve distinct parameters for each block. Six of these parameters modulate text features, and the remaining six modulate image features. The integration of these modulations is commonly achieved using techniques such as Adaptive Layer Normalization (AdaLN) and residual connections. AdaLN is implemented prior to the Attention Layer and Feed-Forward Layer, and can be expressed as:

$$\text{AdaLN}(x, \alpha, \beta) = \alpha \cdot \frac{x - \mu(x)}{\sigma(x)} + \beta. \quad (2)$$

where x is the input feature, $\mu(x)$ and $\sigma(x)$ are its mean and standard deviation respectively, and α (scale) and β (bias) are modulation parameters derived from the conditioning vector y . Residual connections are added after the Attention layer and the Feed-Forward layer. They help stabilize training and enable deeper networks, typically taking the form:

$$x_{out} = \gamma \mathcal{F}(x_{in}) + x_{in}. \quad (3)$$

where x_{in} is the input to a layer or block, $\mathcal{F}(x_{in})$ is the learned residual mapping, x_{out} is the output, and γ is a scaling parameter applied to the residual term, also derived from the conditioning vector y . These mechanisms facilitate fine-grained control over feature representations at various stages within the network. Such adaptive modulation significantly enhances the model's capacity to align image generation with the conditioning inputs.

An essential aspect of this modulation mechanism is its ability to operate separately from the main data flow of the attention mechanism. This separation allows for the precise integration of visual feature representation into specific words without interfering with the denoising process. This has the potential to minimize errors and enhance the overall quality of the generated image. Furthermore, text-based features often offer clearer semantic directionality compared to image-based features. This indicates that manipulating text-stream modulation signals can offer a more straightforward and comprehensible method to control the generative process. Consequently, XVerse focuses on leveraging text-side feature modulation to attain precise control over image synthesis.

3.2 Multi-subject Controlled Generation

The framework of XVerse is shown in Figure 2. We first leverage the modulation mechanism for its robust preservation of key attributes from the reference images. Subsequently, we inject fine-grained details with the VAE-encoded image features in single-stream blocks of FLUX model. The subsequent sections will detail these enhancements.

Enhanced Text-Stream Modulation with Image Feature Control. To enhance fine-grained control over the generation process, we augment the existing text-stream modulation framework by introducing image features as an additional control signal. Specifically, given a conditioning image I_c , we first extract its deep features using the CLIP model, represented as $f_c = \text{CLIP}(I_c)$. Subsequently, we combine these image features with the CLIP-encoded features f_p of the text prompt, using a perceiver resampler [25] as the text-stream modulation adapter (T-Mod Adapter). Here, the text features f_p serve as the query vector, and the resampler outputs an offset Δ_{cross} :

$$\Delta_{\text{cross}} = \text{Resampler}(f_p, f_c). \quad (4)$$

This output, Δ_{cross} , encapsulates synergistic information from both the textual prompt and the conditioning image, acting as a corrective or refining signal. For example, if the text prompt p is about a “handbag” and the accompanying image I_c shows a “brown leather handbag”, Δ_{cross} helps the model adjust its representation from a generic handbag to the specific “brown leather handbag” with a series of visual characteristics (such as material, color, and style).

Crucially, Δ_{cross} is formulated as an offset vector targeting specific tokens. This offset is added to the corresponding token embeddings injected into the model, enabling precise control over the semantics of each token while preserving the structure of the text-to-image result. The adjusted conditioning signal y^* is thereby derived:

$$y^* = \text{MLP}(t, f_p) + \Delta_{\text{cross}}. \quad (5)$$

This adjusted conditioning signal y^* is then utilized to adjust the original modulation parameters (e.g., scale and shift parameters) applied to the network’s activations, thereby refining the conditioning influence. Since our method primarily performs injection in the textual space, it can naturally generalize to the control of high-level semantic attributes such as pose, lighting, and style, without requiring additional differentiation for these conditions.

To further enhance the level of control over the final image output, we draw inspiration from StyleGAN’s expansion of the $\mathcal{W}$ latent space to $\mathcal{W}^+$, enabling each Transformer block i in our model to receive customized conditioning. This is achieved by decomposing the Δ_{cross} signal. Instead of a single offset, we compute a shared component, Δ_{shared} (applied across all DiT blocks), and individual components for each block i , denoted as $\Delta_{\text{per-block}}^i$ (which can be computed on a per-block or per-stage basis):

$$y_i^* = \text{MLP}(t, f_p) + \Delta_{\text{shared}} + \Delta_{\text{per-block}}^i. \quad (6)$$

This structured decomposition of the image-derived text offset facilitates more precise and adaptive control over the influence of text conditioning at various levels of the generation process.

Refined Attention Module with Controlled VAE Integration. While using only text-stream modulation can achieve good editability and generation results, its ability to preserve detailed information is still limited. Inspired by OmniControl [11], we introduce VAE-encoded image features as an auxiliary module to further enhance the capability of our approach for maintaining consistency of fine-grained features. To avoid potential negative impacts of directly injecting image features (such as artifacts or degradation of image quality), we constrain the role of VAE features, making them primarily an auxiliary module for supplementing image details, rather than the dominant mechanism for feature injection. Specifically, we restrict the injection of VAE features only to single-stream blocks within the FLUX model. To effectively distinguish different image patch regions within the latent space of the conditioning image, the position index for the latent is changed from (i, j) into individual index as UNO[13].

3.3 Regularizations

In modulation space $\mathcal{M}^+$, feature vectors derived from highly similar subjects often become entangled, leading to subject confusion and unintended fusion in generated results. To address this, we




Figure 3: Training Data Construction Pipeline.

introduce two critical regularization techniques: region preservation loss and text-image attention loss, which facilitate multi-subject disentanglement.

Region Preservation Loss. We construct a new training sample by randomly selecting two existing samples and concatenating them in a side-by-side left-right configuration, with their captions combined into a unified overall caption. We randomly retain modulation injection from only one side (left or right), rather than both. For regions without modulation injection, we enforce consistency between the model’s output and the T2I branch’s output in those regions using L2 loss.

$$\mathcal{L}_{\text{region}} = \mathbb{E}_{x,y} \left[\|(1 - M_c) \odot (V_{\theta'}(z_t, t, y^*) - V_{\theta}(z_t, t, y))\|_2^2 \right] \quad (7)$$

where z_t denotes the noisy latent at timestep t , $V_{\theta}(z_t, t, y)$ denotes T2I branch’s output, $V_{\theta'}(z_t, t, y^*)$ denotes the XVerse’s output, and $1 - M_c$ denotes the non-modulated mask region. This approach not only regularizes subject-specific features but also serves as a data augmentation strategy for multi-subject datasets, enhancing the model’s ability to distinguish and preserve subject characteristics.

Text-Image Attention Loss. To maintain the compositional and editability properties of the T2I branch after modulation injection, we align the cross-attention dynamics between two branches. Specifically, we compute an L2 loss with normalization over the text-image cross-attention maps of the modulated model and the reference T2I branch. This encourages the modulated model to retain attention patterns that closely match those of the T2I branch, ensuring that semantic interactions between text and image regions remain consistent and editable.

3.4 Collection of Training data

High-quality multi-entity data remains scarce, and modulation injection necessitates knowledge of the corresponding text tokens in the target prompt for each conditional image. To address this, we introduce a high-quality general-purpose multi-entity data annotation pipeline.

Single-Image Data Construction. We curated a 1M-scale dataset of images with resolutions exceeding 512 pixels from LAION [26], constructing a universal multi-entity single-image dataset via the data workflow depicted in 3. Specifically, we first employ Florence2 [27] for joint image caption generation and phrase grounding, followed by large language model (LLM)-driven phrase filtering and classification to exclude non-entity terms (e.g., “sky”, “water surface”) that do not align with our definitional criteria. Subject segmentation is then performed using SAM2 [28], with additional face detection and extraction applied specifically for human entities. While open-source data includes diverse scenarios, challenges such as suboptimal aesthetic quality and scene complexity impede model learning. To mitigate this, we supplement the dataset with an additional 1M-scale corpus of high-aesthetic-quality images synthesized using FLUX [6].

Cross-Image Data Construction. For human-centric data, we harvested a 100K-scale single-person multi-view dataset from proprietary in-house collections, forming up to three image pairs per subject ID. For general objects, we leveraged the Subject200K [11] dataset. The construction workflow mirrors that of single-image data, with domain-specific refinements: human data incorporates ID similarity thresholds, while general object data employs DINOv2 [29] similarity filtering—both applied independently to enforce cross-view consistency.

4 Experiments

4.1 Experiments Setting

Implementation Details. In XVerse, we design a three-stage training pipeline to achieve precise multi-subject controlled image generation. We first train the text-stream modulation adapter to




Figure 4: Data distribution and samples for XVerseBench. XVerseBench includes evaluations of single-subject, dual-subject, and triple-subject controlled image generation. The figure also illustrates the number of test samples allocated to each category.

establish foundational semantic alignment. It serves as the basis for conditional feature injection which ensures high-level consistency between injected images and generated results. Building upon the first stage, we introduce VAE-encoded features to enhance detail preservation. While maintaining the global structure and key characteristics obtained from stage one, this phase focuses on injecting fine-grained visual details through hierarchical feature fusion. The first two stages employ hybrid training data containing both single-subject and multi-subject samples from our single-image dataset. However, we found that such a reconstruction paradigm tends to produce copy-paste effects during inference, resulting in reduced diversity and compromised structural fidelity. To address this issue, we propose a third training phase incorporating cross-image data where the subject appearances differ between injected and target images. This extension enhances the capability of our approach to establish robust appearance mapping under heterogeneous visual conditions. Throughout all stages, we maintain mixed training on both single-subject and multi-subject datasets.

Based on FLUX.1-dev [6] text-to-image generation model, we employ LoRA [30] with a rank of 128 to efficiently fine-tune the model while maintaining its generalization capabilities. We utilize two three-layer resamplers with an intermediate dimension of 3072 to generate the shared offsets and per-block offsets for text-stream modulation. The model was trained for 70K, 150K, and 10K iterations in the respective stages. Both the text-stream modulation adapter and LoRA layers were optimized with a learning rate of $5e-6$. The region preservation loss was assigned a weight of 10, while the text-image attention loss was weighted at 0.01. Training was conducted on 64 NVIDIA A800 GPUs (40GB each), taking approximately 7 days in total — distributed as 2 days for Stage 1, 4 days for Stage 2, and 1 day for Stage 3.

XVerseBench Details. Existing controlled image generation benchmarks often focus on either maintaining identity or object appearance consistency, rarely encompassing datasets that rigorously test both aspects. To comprehensively assess the models’ single-subject and multi-subject conditional generation and editing capabilities, we constructed a new benchmark by merging and curating data from DreamBench++ [31] and Unsplash50 [32].

Our resulting benchmark XVerseBench comprises 20 distinct human identities, 65 unique objects, and 45 different animal species/individuals. To thoroughly evaluate model effectiveness in subject-driven generation tasks, we developed test sets specifically for single-subject, dual-subject, and triple-subject control scenarios. This benchmark includes 300 unique test prompts covering diverse combinations of humans, objects, and animals. Figure 4 shows more detail information and samples for each categories. For evaluation, we employ a suite of metrics to quantify different aspects of generation quality and control fidelity: including DPG score [33] to assess the model’s editing capability, Face ID similarity [34] and DINOv2 [29] similarity to assess the model’s preservation of human identity and objects, and Aesthetic Score [35] to measure to evaluate the aesthetics of the generated image. XVerseBench aims to provide a more challenging and holistic evaluation framework for state-of-the-art multi-subject controllable text-to-image generation models.

4.2 Comparisons with State-of-the-art Methods

We compare our proposed XVerse method with several leading multi-subject driven generation techniques, including MS-Diffusion [36], MIP-Adapter [37], OmniGen [38], UNO [13], and DreamO

Table 1: Quantitative results of single-subject and multi-subject driven generation on XVerseBench.

Method	Single-Subject					Multi-Subject					Overall
	DPG	ID-Sim	IP-Sim	AES	AVG	DPG	ID-Sim	IP-Sim	AES	AVG	
MS-Diffusion [36]	96.94	6.58	51.06	59.69	53.57	87.27	4.81	40.90	55.87	47.21	50.39
MIP-Adapter [37]	77.48	28.39	66.32	52.09	56.07	84.52	19.49	49.89	51.78	51.42	53.75
OmniGen [38]	85.19	<u>60.17</u>	70.73	51.89	67.00	81.71	42.18	52.11	51.35	56.84	61.92
UNO [13]	91.82	37.22	74.35	55.21	64.65	87.57	26.00	60.62	53.04	56.81	60.73
DreamO [14]	97.51	58.74	67.69	53.80	<u>69.44</u>	<u>89.75</u>	<u>44.21</u>	<u>60.87</u>	51.16	<u>61.50</u>	<u>65.47</u>
XVerse (Ours)	93.50	63.02	<u>71.35</u>	<u>56.63</u>	71.13	91.77	51.03	61.04	<u>53.68</u>	64.38	67.76




Figure 5: Qualitative comparison with different methods on XVerseBench.

[14]. The generation results are evaluated for both single-subject and multi-subject tasks, with quantitative findings presented in Table 1. Our XVerse method achieves the highest Overall score of 67.76, significantly outperforming all other compared methods. This clearly indicates a strong comprehensive advantage of our approach.

In the single-subject generation category, XVerse demonstrates exceptional performance, securing the top AVG score of 71.13. This underscores its robust capability in generating high-quality images focused on individual subjects. Notably, XVerse achieves the best identity similarity score (ID-Sim) of 63.02, suggesting superior preservation of subject identity. While DreamO leads in DPG with 97.51, XVerse’s strong average performance, bolstered by competitive scores in IP-Sim (Object Similarity) at 71.35 and AES (Aesthetic Score) at 56.63, highlights its well-rounded excellence.

Meanwhile, XVerse truly excels in the more challenging multi-subject generation tasks, achieving a leading AVG score of 64.38. This remarkable performance can be attributed to XVerse’s novel approach of learning offsets within the text-stream modulation mechanism of DiT. This allows for precise conditioning from diverse image types while crucially preserving the image’s structural integrity. Furthermore, the careful integration of VAE-derived features for detail refinement, rather than dominant conditioning, effectively mitigates artifacts and distortions, which is particularly vital for maintaining clarity and attribute disentanglement in multi-subject scenarios.

Figure 5 presents a qualitative comparison of our method with other state-of-the-art approaches. As illustrated, our model demonstrates a superior capability in maintaining the consistency and relevance

Table 2: Ablation Study of Joint Training of Text Modulation and VAE-Encoded Features

EXP ID	EXP Name	Init	Steps	Single-AVG	Multi-AVG	AVG
1	Text Modulation Only	Scratch	30K	57.55	52.73	55.14
			40K	58.30	53.54	55.92
2	Single-Stream VAE Only	Scratch	30K	60.71	53.32	57.02
			40K	59.60	53.37	56.49
3	Text Modulation + VAE	Exp 1, 30K	10K	62.62	55.16	58.89

Table 3: Ablation Study of Different Regulation Losses

EXP	Steps	Single-AVG	Multi-AVG	AVG
Baseline	30K	57.55	52.73	55.14
Baseline w/o RPL	30K	55.47	50.72	53.10
Baseline w/o TIAL	30K	53.42	44.51	48.97

between identities and associated objects within the generated images. This is a direct result of our refined modulation strategy. By carefully adjusting the modulation offsets, XVerse achieves enhanced text-image alignment, which is particularly evident in its accurate depiction of object quantities and the relationships between multiple subjects. Furthermore, when compared to existing methods, our model consistently produces images with a higher degree of naturalness and visual plausibility. This improved realism underscores the advantages of our approach in editing the text-stream modulation pathway, allowing for more faithful and aesthetically pleasing image synthesis.

4.3 Ablation Studies

To further evaluate the effectiveness of different modules in XVerse and analyze the impact of key components, we conduct a series of ablation studies. All quantitative comparisons are performed under the same configuration (e.g., hyperparameters, training schedule), unless otherwise specified. Due to computational resource limitations, all experiments are executed on 16 NVIDIA A800 (40GB) GPUs.

Joint Training of Text Modulation and VAE-Encoded Features We test three setups to examine text-VAE collaboration: (1) Text Modulation Only, (2) Single-Stream VAE Only, (3) Joint Training. The results are shown in Table 2. After 30K steps, Text Modulation Only scores 57.55 (single-subject) and 52.73 (multi-subject) AVG; extending to 40K steps (additional 10K) brings marginal gains (58.30/53.54). Single-Stream VAE Only outperforms text-only at 30K (60.71/53.32) but degrades at 40K, reflecting standalone VAE instability (attribute blending disrupts consistency). Notably, Experiment 3 (Joint Training, initialized from Exp.1’s 30K checkpoint, 10K extra steps) boosts scores to 62.62 (single) and 55.16 (multi), mitigating VAE injection instability and enabling more disentangled generation. Figure 6 shows the qualitative comparison in both dual-subject and triple-subject controlled generation tasks. The experimental results show that by converting the reference image to text-stream modulation offset vectors, XVerse achieves the personalization of the condition subjects while maintaining the composition identical to the original T2I outcomes. Additionally, the VAE features play a supportive role by helping the model add specific details without compromising the overall structure and semantic coherence established by the text-stream modulation. This collaborative approach allows XVerse to maintain subject consistency while achieving a high degree of editability.

Impacts of Different Regularization Losses. To verify the necessity of Region Preservation Loss (RPL) and Text-Image Attention Loss (TIAL), we compare the baseline (both losses) with RPL-ablated and TIAL-ablated models. Other configurations (e.g., training steps) are identical to isolate loss impacts. As shown in Table 3, the baseline achieves the highest AVG (single: 57.55, multi: 52.73). RPL ablation drops performance by 2.04% (as it helps to maintain category consistency), while TIAL ablation causes a 7.17% decline (as it ensures precise attribute localization, reducing irrelevant feature leakage). This confirms RPL and TIAL are complementary and indispensable.

Table 4: Ablation Study of Different VAE-encoded Feature Injection Strategies

EXP	Steps	Single-AVG	Multi-AVG	AVG
Single-stream, VAE Only	30K	60.71	53.32	57.02
All Blocks, VAE Only	30K	58.58	51.35	54.97



Figure 6: Effect of text-stream modulation resampler and VAE-encoded image features.




Figure 7: The control of semantic attributes, such as cloth, pose, lighting, and style.

VAE-encoded feature injection blocks. We evaluate two VAE injection strategies to identify the optimal configuration: “Single-stream, VAE Only” (injecting exclusively into single-stream blocks) and “All Blocks, VAE Only” (injecting into both double-stream and single-stream blocks). As shown in Table 4, the single-stream injection consistently outperform the all-blocks approach, achieving average scores of 60.71 versus 58.58 for single-subject evaluations, and 53.32 versus 51.35 for multi-subject tests. The inferior performance of the all-blocks injection is attributed to feature interference, where redundant signals disrupt the core generation process. Focusing injection on the core single-stream pathways results in better alignment and overall effectiveness compared to a more diffuse injection strategy.

4.4 Applications

Figure 7 demonstrates XVerse’s capability to control semantic attributes of the generated image. XVerse exhibits precise control over attributes such as lighting, subject posture, clothing, and artistic style. By injecting reference images toward targeted words, XVerse can manipulate these semantic attributes without the need for extensive training data on specific attribute categories. These results further underscore the model’s exceptional ability to generalize and edit effectively.

5 Conclusion

In this paper, we introduce XVerse, an innovative framework designed to excel in the complex task of precise multi-subject control within DiTs. XVerse injects reference image features through modulation offsets and control the token-specific representation within the DiT blocks. This approach facilitates adaptive, per-subject conditioning by precisely governing how textual embeddings are transformed and integrated throughout the diffusion process. Consequently, XVerse effectively mitigates common generation issues such as attribute entanglement and artifacts, demonstrating particular excellence in both single-subject and multi-subject controlled generation tasks.

Limitations. While XVerse showcases significant progress in fine-grained multi-subject control, there are certain limitations that need to be addressed and opportunities for future research. One major challenge is the lack of high-quality, large-scale cross-image multi-subject datasets, which are essential for training and evaluating models that can understand and generate complex interactions between subjects. Additionally, our current research has mainly focused on the text-stream modulation pathway, leaving room for exploration in utilizing image-modulation techniques for precise pixel-level or region-specific control.

References

- [1] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems 27*, 2014.
- [2] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [3] Jascha Sohl-Dickstein, Eric A Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.
- [4] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851, 2020.
- [5] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.
- [6] Black Forest Labs. Flux: Official inference repository for flux.1 models, 2024. Accessed: 2024-11-12.
- [7] 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 *The Eleventh International Conference on Learning Representations*, 2023.
- [8] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023.
- [9] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023.
- [10] Li Chen, Mengyi Zhao, Yiheng Liu, Mingxu Ding, Yangyang Song, Shizun Wang, Xu Wang, Hao Yang, Jing Liu, Kang Du, et al. Photoverse: Tuning-free image customization with text-to-image diffusion models. *arXiv preprint arXiv:2309.05793*, 2023.
- [11] Zhenxiong Tan, Songhua Liu, Xingyi Yang, Qiaochu Xue, and Xinchao Wang. Ominicontrol: Minimal and universal control for diffusion transformer. *arXiv preprint arXiv:2411.15098*, 2024.
- [12] Xi Chen, Zhifei Zhang, He Zhang, Yuqian Zhou, Soo Ye Kim, Qing Liu, Yijun Li, Jianming Zhang, Nanxuan Zhao, Yilin Wang, et al. Unreal: Universal image generation and editing via learning real-world dynamics. *arXiv preprint arXiv:2412.07774*, 2024.
- [13] Shaojin Wu, Mengqi Huang, Wenxu Wu, Yufeng Cheng, Fei Ding, and Qian He. Less-to-more generalization: Unlocking more controllability by in-context generation. *arXiv preprint arXiv:2504.02160*, 2025.
- [14] Chong Mou, Yanze Wu, Wenxu Wu, Zinan Guo, Pengze Zhang, Yufeng Cheng, Yiming Luo, Fei Ding, Shiwen Zhang, Xinghui Li, et al. Dreamo: A unified framework for image customization. *arXiv preprint arXiv:2504.16915*, 2025.
- [15] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023.
- [16] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019.

- [17] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021.
- [18] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. In *Advances in neural information processing systems*, volume 33, pages 9841–9850, 2020.
- [19] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2287–2296, 2021.
- [20] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Transactions on graphics (TOG)*, 42(1):1–13, 2022.
- [21] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *European conference on computer vision*, pages 592–608, 2020.
- [22] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- [23] Daniel Garibi, Shahar Yadin, Roni Paiss, Omer Tov, Shiran Zada, Ariel Ephrat, Tomer Michaeli, Inbar Mosseri, and Tali Dekel. Tokenverse: Versatile multi-concept personalization in token modulation space. *arXiv preprint arXiv:2501.12224*, 2025.
- [24] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763, 2021.
- [25] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35: 23716–23736, 2022.
- [26] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021.
- [27] Bin Xiao, Haiping Wu, Weijian Xu, Xiyang Dai, Houdong Hu, Yumao Lu, Michael Zeng, Ce Liu, and Lu Yuan. Florence-2: Advancing a unified representation for a variety of vision tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4818–4829, 2024.
- [28] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.
- [29] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- [30] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.
- [31] Yuang Peng, Yuxin Cui, Haomiao Tang, Zekun Qi, Runpei Dong, Jing Bai, Chunrui Han, Zheng Ge, Xiangyu Zhang, and Shu-Tao Xia. Dreambench++: A human-aligned benchmark for personalized image generation. In *The Thirteenth International Conference on Learning Representations*, 2025.

- [32] Zinan Guo, Yanze Wu, Chen Zhuowei, Peng Zhang, Qian He, et al. Pulid: Pure and lightning id customization via contrastive alignment. In *Advances in neural information processing systems*, volume 37, pages 36777–36804, 2024.
- [33] Xiwei Hu, Rui Wang, Yixiao Fang, Bin Fu, Pei Cheng, and Gang Yu. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*, 2024.
- [34] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019.
- [35] discuss0434. Aesthetic predictor v2.5: Siglip-based aesthetic score predictor. <https://github.com/discuss0434/aesthetic-predictor-v2-5>, 2024. Accessed: 2024-12-08.
- [36] Xierui Wang, Siming Fu, Qihan Huang, Wanggui He, and Hao Jiang. Ms-diffusion: Multi-subject zero-shot image personalization with layout guidance. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [37] Qihan Huang, Siming Fu, Jinlong Liu, Hao Jiang, Yipeng Yu, and Jie Song. Resolving multi-condition confusion for finetuning-free personalized image generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 3707–3714, 2025.
- [38] Shitao Xiao, Yueze Wang, Junjie Zhou, Huaying Yuan, Xingrun Xing, Ruiran Yan, Shuting Wang, Tiejun Huang, and Zheng Liu. Omnigen: Unified image generation. *arXiv preprint arXiv:2409.11340*, 2024.

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 have presented the contributions of the paper in both 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 create a separate "Limitations" part in the Conclusion section.

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: [NA]

Justification: The paper does not include theoretical results.

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: [Yes]

Guidelines: The experimental details have been outlined in the main text, and we plan to make our code available at a later time.

- 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: [No]

Justification: We plan to make our code available at a later time.

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: In the beginning part of the experiments sections.

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: [No]

Justification: We are following the testing protocol of the previous method, where the results are represented by the mean performance of images generated by four random seeds to represent the model performance.

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: In the beginning part of the experiments sections.

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: The research conducted in the paper conform, in every respect, 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: In the supplementary materials.

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: We will have the detailed usage guidelines with our realised code.

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: The benchmark used in the paper are all from open-source datasets.

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 have the detailed usage guidelines with our realised code.

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

Justification: The data we use is come from those open-access dataset.

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.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

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

Justification: We use LLM for data labeling and model evaluation, and describe it in the paper.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Samples for Training Dataset




Figure 8: Examples of training data for multi-subject controlled generation.

Figure 8 presents examples of our training data for multi-subject controlled generation. As illustrated, the dataset covers a diverse range of scenarios, including human-object interactions, human-animal compositions, and complex multi-person scenes. For human-centric data, we intentionally randomly select facial images or full-body images as reference inputs. This strategy can further enhance the model's generalization performance. By utilizing this diverse and extensive dataset, which encompasses a wide range of scene variations and control types, our model is able to achieve impressive editing capabilities while maintaining high consistency with the reference images.

B Impact of Prompt Variation on the Generated Image




Figure 9: Impact of prompt variation on subject-controlled Image generation. The reference image and the initial output of our text-to-image generation model are shown on the left side. The right side illustrates the influence of different prompts on the generated output, with the prompt variances highlighted in red.

To evaluate the impact of prompt variation on subject-controlled image generation, we modified the injected words in the per-token text-modulation module while keeping the reference image constant. The results, shown in Figure 9, offer valuable insights. This visualization effectively illustrates

that a more detailed prompt description improves the preservation of the subject's identity in the generated image. Additionally, when the prompt closely matches the reference image, our model not only incorporates intricate image details but also maintains a high level of control over the subject's attributes, even allowing for successful changes such as gender. On the other hand, if there is a significant semantic mismatch between the injected prompt and the reference image (e.g., trying to generate a person's image from prompts like "a dog" or "a tree stump"), the injection process consistently fails. This highlights our model's ability to accurately target and incorporate reference image features into specific words, enabling precise control over the generated output.

C Comparison of the CLIP-T and DPG scores

When evaluating text-to-image generation models, the CLIP-T [24] score has been a prevalent metric in prior studies, assessing semantic consistency by leveraging CLIP's image-text embeddings. However, our research highlights the superior efficacy of the DPG (Dense Prompt Graph) score, particularly for intricate prompts. While CLIP-T offers a broad measure of semantic alignment, the DPG score is specifically designed to evaluate a model's capacity to interpret and execute detailed and complex textual instructions. It rigorously assesses editing abilities across multiple objects, diverse attributes, and intricate relationships, thereby capturing the nuanced and fine-grained semantic alignment crucial for advanced compositional generation. This provides a more comprehensive and robust evaluation for challenging scenarios.

D Illustration of Region Preservation Loss




Figure 10: Illustration of the region preservation loss.

Figure 10 shows the illustration of our region preservation loss. We form training samples by concatenating two existing samples, merging their captions, and randomly applying modulation to only one side. For the unmodulated regions, defined by M_c , we enforce consistency between our model's output ($V_{\theta'}(z_t, t, y^*)$) and the text-to-image branch's output ($V_{\theta}(z_t, t, y)$) via an L2 loss (Eq. 1). By using this regularization, XVerse can better inject the reference image into specific areas without affecting the generation of irrelevant areas, thereby achieving more precise generation control.

E Ablation Study for Text-Image Attention Loss




Figure 11: The qualitative comparison of Text-Image Attention Loss. This image shows the generated results and attention maps for "woman", "coffee cup", and "pink suit" for each method.

To validate the effectiveness of our text-image attention loss, we conducted an experiment where we excluded this regularization and examined the generated outputs along with their respective attention maps. The qualitative analysis presented in Figure 11 clearly demonstrates the significance of this method. It demonstrates our method's ability to maintain the structural and editable characteristics of the T2I branch following modulation injection. Through ensuring L2 consistency between the cross-attention maps of the modulated model and the reference T2I branch, our approach ensures the reliable preservation of text-image semantic interactions. This ultimately enables precise control over semantics, as visually evidenced by the generated results and attention maps for specific prompts like "woman," "coffee cup," and "pink suit."

F Broader Impacts

Our model, XVerse, marks a significant leap in multi-subject controllable text-to-image generation, leading to enhanced fidelity and editability. This breakthrough holds substantial positive societal impacts, particularly within the creative industries, where it can revolutionize the creation of personalized and complex visual content. Furthermore, XVerse can transform education and training by providing more engaging and tailored visual aids, and contribute to content inclusivity by enabling the representation of a wider range of individuals and scenarios.

However, this powerful technology also presents potential negative societal impacts. The improved generation capability could lead to misinformation and deepfakes, raise privacy concerns if used improperly, and potentially amplify biases present in training data. As foundational research, XVerse isn't directly tied to deployment. Yet, we believe it's crucial to acknowledge these risks. Future work will explore mitigation strategies like content detection and ethical guidelines, contributing to the responsible advancement of generative AI.