

CharaConsist: Fine-Grained Consistent Character Generation

Mengyu Wang^{1,3}, Henghui Ding^{2†}, Jianing Peng^{1,3}, Yao Zhao^{1,3}, Yunpeng Chen⁴, Yunchao Wei^{1,3†}

¹Institute of Information Science, Beijing Jiaotong University, Beijing, China

²Institute of Big Data, Fudan University, Shanghai, China

³Visual Intelligence + X International Joint Laboratory of the Ministry of Education

⁴Alkaid Pte. Ltd., Singapore

mengyu.wang@bjtu.edu.cn, henghui.ding@gmail.com, wychao1987@gmail.com

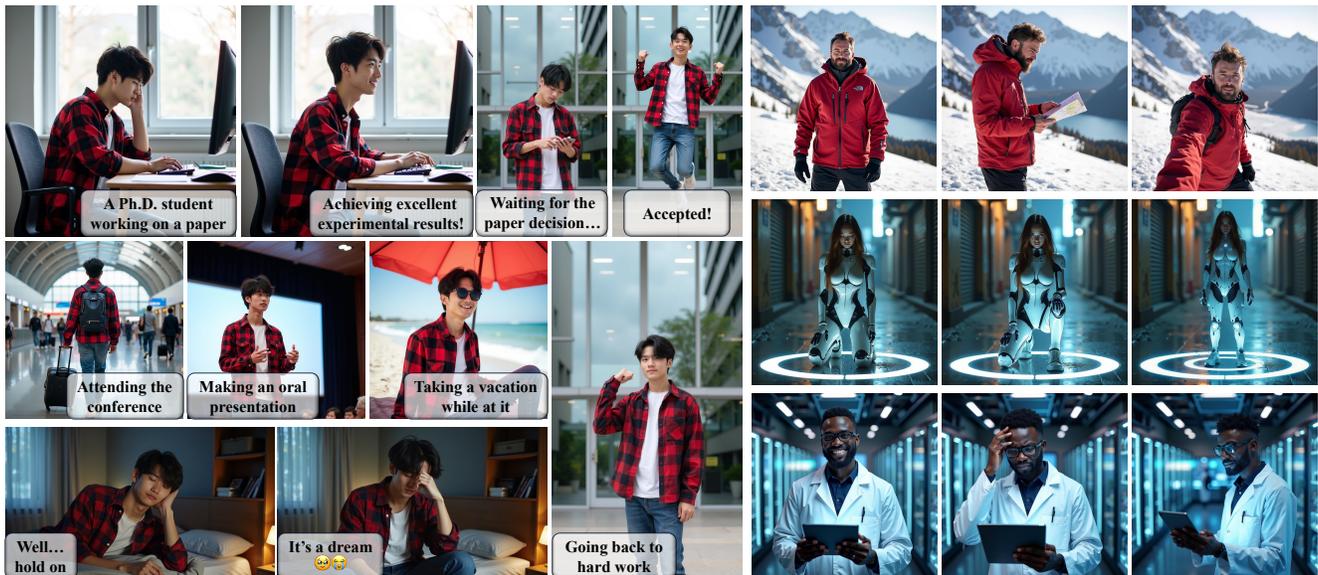


Figure 1. *CharaConsist* achieves fine-grained consistency maintaining. The left part shows that, *CharaConsist* enables more flexible storytelling, supports controllable scene transitions, and achieves character consistency within a fixed scene, across different scenes, and across resolutions. The right part highlights *CharaConsist*'s ability to preserve extensive background details with complete consistency.

Abstract

In text-to-image generation, producing a series of consistent contents that preserve the same identity is highly valuable for real-world applications. Although a few works have explored training-free methods to enhance the consistency of generated subjects, we observe that they suffer from the following problems. First, they fail to maintain consistent background details, which limits their applicability. Furthermore, when the foreground character undergoes large motion variations, inconsistencies in identity and clothing details become evident. To address these problems, we propose *CharaConsist*, which employs point-tracking attention and adaptive token merge along with decoupled control of the foreground and background. *CharaConsist*

enables fine-grained consistency for both foreground and background, supporting the generation of one character in continuous shots within a fixed scene or in discrete shots across different scenes. Moreover, *CharaConsist* is the first consistent generation method tailored for text-to-image DiT model. Its ability to maintain fine-grained consistency, combined with the larger capacity of latest base model, enables it to produce high-quality visual outputs, broadening its applicability to a wider range of real-world scenarios. The source code has been released at [CharaConsist.git](https://github.com/CharaConsist/CharaConsist).

1. Introduction

Text-to-image (T2I) generation has advanced significantly in recent years [18, 21], especially with the widespread emergence of Diffusion Transformers (DiTs) [3, 12, 17]. Their successful scaling-up has led to significant improve-

[†]Corresponding authors: Henghui Ding and Yunchao Wei.

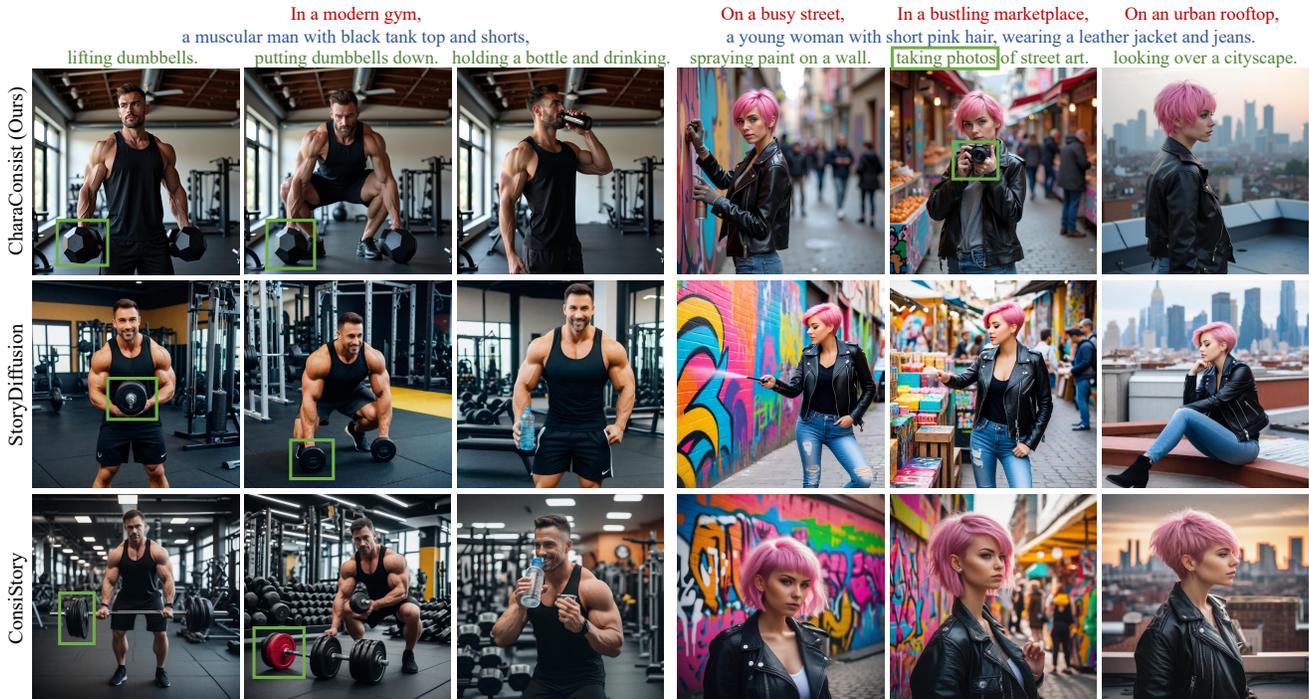


Figure 2. **Qualitative comparison with StoryDiffusion [35] and ConsiStory [30].** For continuous shots within a fixed scene (the left part), previous methods cannot maintain consistent background and foreground details. For discrete shots across different scenes (the right part), previous methods either exhibit a copy-paste effect or suffer from inconsistency for the characters. In contrast, our *CharaConsist* maintains fine-grained consistency for both characters and environments, while also providing greater flexibility in character actions.

ments in both image quality and text alignment, unlocking new possibilities for many real-world applications. A key challenge is generating multiple images while ensuring subject consistency, a crucial requirement that has gained increasing attention [16, 30, 35]. Consistency is essential for applications such as visual storytelling, virtual character design, and long video generation, where continuity across images enhances usability and realism.

Considering computational resources and time consumption, recent studies have explored training-free methods [16, 30, 35] to achieve consistent T2I generation. These methods introduce inter-image information sharing during generating a batch of images to enhance consistency and focus primarily on achieving the visual effect of **one character transitioning across different backgrounds**. However, this task formulation remains overly coarse, whereas many real-world applications demand **finer-grained consistency**. For example, as shown in Fig. 1, visual storytelling often requires generating a sequence of coherent story shots within the same scene. Similarly, in video-related applications, ensuring fully consistent keyframes is critical for guiding long video synthesis. Failure to do so can lead to flickering artifacts, compromising the final output quality.

Existing methods [16, 30, 35] face significant challenges when addressing such fine-grained requirements. First, their results exhibit **inconsistency of the static background**.

When generating a series of continuous shots, *e.g.*, the gym scene in Fig. 2, these methods failed to maintain consistency of neither the background instruments, nor the foreground dumbbell, which could lead to flickering artifacts when the generated results serving as video keyframes. Second, they suffer from a **trade-off between foreground consistency and action variation**. As shown in the right part of Fig. 2, the generated characters either exhibit a nearly copy-paste effect, deviating from the unique action description of each image, or suffer from inconsistency in details like clothing and hairstyle. Furthermore, these methods are inherently constrained by the **underpowered base model**. Most existing approaches are built on SDXL [18], an earlier T2I model that lags behind state-of-the-art Diffusion Transformers (DiTs) [3, 12, 17] in generating high-quality characters and complex scenes. Additionally, their UNet-specific implementation strategies cannot be directly adapted to newer DiT architectures.

To address these challenges, we propose *CharaConsist*, a training-free method built on FLUX.1 [12] to achieve the fine-grained consistency. First, we introduce point-tracking attention. This module first establishes positional correspondences across images and subsequently re-encode the positional embeddings during the inter-image information sharing process. Point-tracking attention enables the automatic tracking of critical features across images,

therefore not only enhancing consistency but also providing greater flexibility in character actions. Second, while re-encoding the positional embeddings mitigates the impact of positional shifts, it may inadvertently disrupt the original local geometric structure among tokens. To address this, we further introduce an adaptive token merge strategy, which refines the attention outputs by dynamically aggregating tokens based on their similarity and positional correspondences established in the first step. Third, to accommodate varying requirements for background consistency across different application scenarios, we introduce a foreground-background mask, which is derived from the attention differences between image tokens and text tokens corresponding to foreground and background elements. By leveraging this distinction, our approach enables decoupled control over foreground and background consistency, allowing for more flexible adaptation to diverse generation needs.

In experiments, we compare the proposed *CharaConsist* with previous training-free consistent T2I generation methods, including ConsiStory [30] and StoryDiffusion [35], as well as training-dependent identity-reference generation methods, such as IP-Adapter [34] and PhotoMaker [13]. The results demonstrate that *CharaConsist* significantly outperforms these methods in both foreground and background consistency, underscoring its effectiveness in fine-grained consistent generation. In summary, our key contributions are as follows:

- We highlight that existing consistent generation works adopt relatively coarse task formulations, primarily concentrating on transitions of similar characters across different scenes, while lacking the fine-grained consistency in character details and background environments.
- We propose *CharaConsist*, which leverages point tracking and mask extraction to automatically perceive critical features, thus achieving fine-grained consistency, making it well-suited for real-world applications.
- *CharaConsist* is the first training-free consistent T2I generation method built on a DiT model. It operates without additional training or extra modules, enabling the efficient utilization of large DiT models with billions of parameters. Moreover, the proposed DiT-based point tracking and mask extraction techniques can serve as effective tools for related tasks such as image editing.

2. Related Work

2.1. Diffusion-based T2I Generation

With the advent of DDPM [7], diffusion models [27] have gradually surpassed GANs [5] to become the dominant technology in visual generation. Earlier diffusion models were primarily based on the UNet [22] architecture, such as SD 1.x [21], SD 2.x, and SDXL [18]. These models process visual information through convolutional [6] and

self-attention [31] blocks, and capture textual prompt information through cross-attention with text tokens encoded by CLIP [19], enabling text-to-image generation.

More recently, Diffusion Transformers (DiTs) [17] were introduced, where the authors replaced the commonly used UNet with a Transformer [31] and demonstrated its superior scalability. Over the past one year, numerous DiT-based T2I [3, 12, 14, 25] and T2V [10, 15, 26, 33] models have emerged and demonstrated a superior ability on producing visually stunning results. Among them, the latest T2I DiTs, such as SD3.x [3] and FLUX.1 [12], adopt similar multimodal-transformer architectures. These models use the T5 encoder [20] to encode text tokens and a vae [2, 8] to encode image tokens, and then concatenate these tokens into a single sequence to perform attention. Compared to UNet-based T2I models, these models no longer feature separate self-attention and cross-attention, making it challenging to directly adapt previous methods to DiTs, which often manipulate the specific layers in the UNet.

2.2. Training-free Consistent Generation

As a relatively new task, related research is primarily limited to three works: 1P1S [16], StoryDiffusion [35], and ConsiStory [30]. 1P1S is built on the consistency from text tokens, which concatenates the textual prompts of multiple images and encodes them together. However, this approach is limited by the maximum sequence length of the text encoder. As a result, it's challenging to employ it to scenarios requiring relatively long text prompts, such as the diverse characters and scenes generation in this work.

StoryDiffusion and ConsiStory share similar idea, introducing inter-image attention as core technique to enhance consistency. They randomly sample a subset of tokens from a batch of images and use them as shared keys and values in the attention operation for each image. However, we observe that such mechanism has low information retrieval efficiency. This results in two major issues: first, they require at least 2-4 images to be generated in parallel to enhance inter-image associations, leading to significant additional memory consumption. Second, the consistency of their results remains unsatisfactory. Therefore, our proposed *CharaConsist* introduces point-tracking-based attention and token merging, which can adaptively retrieves critical information to enhance consistency, without the need for any image parallelism.

Additionally, similar to our adaptive token merge module, ConsiStory also introduces similarity-based feature injection. However, their similarity matching relies on DIFT [29], a technique based on specific layers and steps of the UNet architecture. However, we found that this approach does not perform effectively on the DiT model.

2.3. Identity-reference Generation

Similar to consistent generation, identity-reference generation aims to align the identity of the generated results with the input reference face image. These works follow two main approaches: The first one involves introducing pre-trained image encoders and training additional projection module to embed the identity information into the generation process, such as IP-Adapter [34] and PhotoMaker [13]. The second approach, widely known as personalization, depends on per-subject training. These methods encodes a specific subject’s identity into a new textual embedding [4] or directly into the model’s parameters [11, 23], ultimately making the generated results overfit to a fixed identity.

However, these methods focus solely on facial identity preserving and cannot ensure consistency in other attributes such as clothing or scenes. Additionally, they require training on large-scale collected datasets, or performing per-subject training, which imposes high demands on computational resources or time consumption.

3. Method

In this section, we first provide a brief review and analysis of the core technique, inter-image attention, used to enhance consistency in previous works, along with some technical details of the base model FLUX.1 in Sec. 3.1. Then, we introduce our proposed *CharaConsist* in Sec. 3.2. Figure 4 presents the schematic of the overall framework and individual components of our method.

3.1. Preliminary

Inter-image Attention In previous works, inter-image attention is a core technique used to enhance consistency. During the generation of a batch of images, inter-image attention works by concatenating keys and values from neighboring images into each image’s self-attention process, enabling information sharing across images. This process helps maintain coherence in the generated sequence by promoting cross-image relationships.

However, we observe that although attention module has a global receptive field, directly applying inter-image attention struggles to capture the necessary correspondences for maintaining consistency. To intuitively demonstrate this, as shown in Fig. 3, we generate images with similar content but varying layouts and compute their averaged inter-image attention weights using both the UNet-based SDXL [18] and the DiT-based FLUX.1 [12]. Using the query coordinate as the origin, we divide each image into concentric rings of equal area and calculate the attention weight summation within each ring. The results show that as the distance from the query coordinate increases, the attention weight summation of the corresponding region decreases significantly. This phenomenon indicates that the model

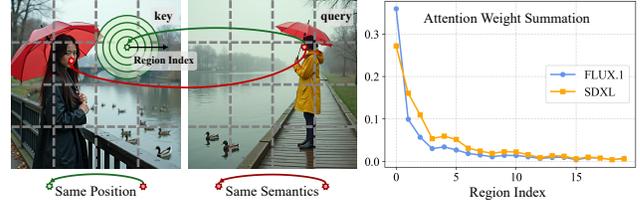


Figure 3. **The locality bias in inter-image attention.** Both SDXL [18] and FLUX.1 [12] tend to allocate more attention weights to spatially nearby regions instead of more semantically related ones.

exhibits a locality bias, tending to allocate more attention weights to spatially nearby regions rather than to more semantically related ones. As a result, the model struggles to maintain consistency across varying layouts, failing to capture distant but contextually important features.

The inconsistency observed in previous methods, as shown in Fig. 2, could largely stem from this locality bias. When characters in different images undergo large motion variations or occupy different positions within the image, their relative distance could increase significantly. As a result, the inter-image attention yields lower response intensity for these distant regions, making it challenging to gather sufficient information to ensure consistency. Building on this observation, our method begins with point tracking and incorporates several techniques to mitigate the effects of the positional shift, ensuring more robust and consistent representations across varying image layouts.

FLUX.1 FLUX.1 [12] is a latent-space DiT-based T2I model with a multimodal transformer as its core architecture. It includes two types of transformer blocks: the double block and the single block, with both blocks taking text and image tokens as input.

Although the two blocks differ in structure and parameter allocation, their core mechanism remains the same: the concatenated text and image tokens undergo global attention. As a training-free method, our primary focus is on the attention process within the model. Therefore, we omit the structural differences between the two types of blocks and formulate the attention process in both as a unified form:

$$X_{mm} = \text{cat}[X_t, X_i] \in \mathbb{R}^{(l+hw) \times d}, X \in \{Q, K, V\}, \quad (1)$$

$$W_{mm} = \frac{Q_{mm} \cdot K_{mm}^T}{\sqrt{d}} \in \mathbb{R}^{(l+hw) \times (l+hw)}, \quad (2)$$

$$H_{mm} = \text{Softmax}(W_{mm}) \cdot V_{mm} = \text{cat}[H_t, H_i], \quad (3)$$

where $\text{cat}[\cdot]$ refers to concatenation. The subscripts mm , t and i stand for “multimodal”, “text”, and “image”, respectively. l is the length of text tokens. d is the dimension of each attention head. (h, w) represents the resolution of the image tokens, which is $1/16$ of the generated images. The number of attention head is omitted for simplicity.

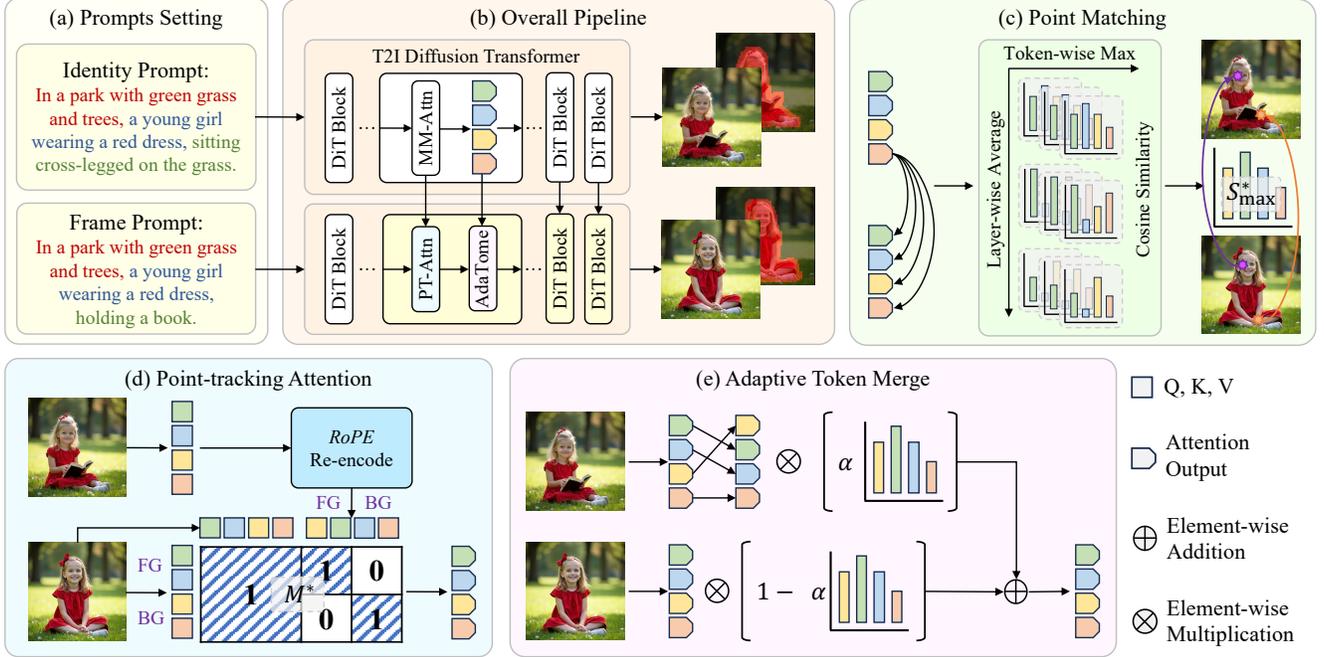


Figure 4. **An overview of our CharaConsist.** (a) shows an example of the textual prompts. (b) is the overall pipeline of our method, in which we replace the original multimodal attention with our proposed point-tracking attention, and further introduce an adaptive token merge module, as detailed in (d) and (e). (c) shows the point matching strategy of our method.

3.2. CharaConsist

Pipeline Definition We divide the consistent generation pipeline into two parts: identity image generation and frame images generation. The identity image generation follows the original generation process without modifications, and stores intermediate variables such as keys, values, and attention outputs from specific layers at certain timesteps. The frame images then achieve consistency with the identity image by accessing these stored variables.

Unlike previous methods that require generating at least 2 to 4 batched images in parallel to obtain identity information, *CharaConsist* needs only a single identity image. This advantage largely arises from the tracking mechanism in our method, which effectively uses the identity information from a single image, eliminating the need for parallel processing for augmentation. As a result, this approach significantly reduces the additional GPU memory overhead.

Point Matching The goal of point matching is to identify, for each point in the frame image, the semantically corresponding point within the identity image, thereby guiding subsequent attention and token merging processes. For example, the facial features of the character in the frame image should align with those in the identity image. Previous work DIFT [29] demonstrates that semantic correspondence can be determined by measuring the similarity of intermediate features from the UNet in diffusion models at a specific timestep. However, we observe that this

approach failed in FLUX.1. Specifically, we test the cosine similarity of outputs from both intermediate transformer blocks and attention layers in FLUX.1. For the transformer block outputs, we found that the matching points remain fixed at the same position, failing to capture semantic correspondence. In contrast, for attention layer outputs, the matching points exhibit significant fluctuations across layers, timesteps, and samples, making stable and accurate matching difficult. Furthermore, we observe that while matching relationships across layers vary significantly, the similarity distributions exhibit consistent characteristics. Based on this observation, we average the similarity scores across layers at the same timestep, formulated as:

$$S^* = \frac{1}{N} \sum_n \frac{H_{i,frm,n} \cdot H_{i,id,n}^T}{\|H_{i,frm,n}\|_2 \cdot \|H_{i,id,n}\|_2}, \quad (4)$$

$$map^*(\cdot), S_{\max}^* = \operatorname{argmax}(S^*), \max(S^*), \quad (5)$$

where N and n are the number and index of layers, respectively. $map^*(\cdot)$ is the matching results and $map^*(j) = k$ means the j -th point in frame image is matched with the k -th point in identity image. S_{\max}^* is the averaged cosine similarity of the matched token pair. This approach provides stable semantic similarity results and accurate point matching relationships. Results and further details can be found in Fig. 6 and the supplementary materials.

Mask Extraction Our mask extraction strategy is built on comparing the attention weight of image tokens to foreground text tokens against the weight to background text

tokens. Specifically, we format the textual prompt as the background description first, followed by the foreground description, as the example shown in Fig. 4. We then obtain their respective token lengths, denoted as l_{bg} for background and l_{fg} for foreground. Next, we extract the segment of attention weights where image tokens serve as queries and text tokens as keys, represented as:

$$A_{i2t} = \text{Softmax}(W_{mm}[l_{:}, : l_{bg} + l_{fg}]). \quad (6)$$

We apply the Softmax operation only within the actual length of the text sequence to prevent the impact of the image keys and the special text key “<s>”. Here, we observe a similar issue to the point matching process: the attention weights corresponding to each word exhibit significant fluctuations across layers. To address this, we average the attention weights within the foreground and background text sequence separately, and also across different layers. By comparing these averaged attention weights, we derive the robust foreground mask, represented as:

$$A_{i2bg}^* = \frac{1}{N \times l_{bg}} \sum_n \sum_{j=0}^{l_{bg}-1} A_{i2t,n}[:, j], \quad (7)$$

$$A_{i2fg}^* = \frac{1}{N \times l_{fg}} \sum_n \sum_{j=l_{bg}}^{l_{bg}+l_{fg}-1} A_{i2t,n}[:, j], \quad (8)$$

$$M = A_{i2fg}^* > A_{i2bg}^* \in \{0, 1\}^{hw}. \quad (9)$$

Point-tracking Attention After obtaining the point-matching relationships and the foreground masks, we introduce point-tracking attention. Specifically, we first extract the foreground token indices in the frame image, along with the matched token indices in the identity image:

$$C_{fg,frm} = \text{nonzero}(M_{frm}), \quad (10)$$

$$C_{fg,id} = \text{map}^*(C_{frm}), \quad (11)$$

where $\text{nonzero}(\cdot)$ refers to get the indices of non-zero elements. Based on the indices, we extract the keys and values from the identity image and re-encode the positional embedding of the keys as follows:

$$X_{fg,id} = X_{i,id}[C_{fg,id}], X \in \{K^0, V\}, \quad (12)$$

$$K'_{fg,id} = \text{RoPE}(K_{fg,id}^0, C_{fg,frm}), \quad (13)$$

where K^0 represents the keys without positional embedding, RoPE indicates the Rotary Position Embedding [28] used in FLUX.1. We further concatenate the original keys of frame image, the re-encoded foreground keys of identity image, and the background keys of identity image to get the keys of our point-tracking attention, formulated as:

$$K_{frm}^* = \text{cat}[K_{mm,frm}, K'_{fg,id}, K_{bg,id}], \quad (14)$$

$$V_{frm}^* = \text{cat}[V_{mm,frm}, V_{fg,id}, V_{bg,id}]. \quad (15)$$

The operations related to background are optional, depending on whether users want to maintain the background

unchanged. Then the attention process is represented as:

$$W_{frm}^* = \frac{Q_{mm,frm} \cdot K_{frm}^{*T}}{\sqrt{d}} + \log(M^*), \quad (16)$$

$$H_{mm,frm}^* = \text{Softmax}(W_{frm}^*) \cdot V_{frm}^*, \quad (17)$$

where M^* is the attention mask, making the foreground tokens in frame image only attend to the foreground tokens in identity image, and vice versa. This approach effectively leverages the model’s characteristics, as RoPE is applied directly at each attention layer, allowing to store the keys without positional embedding and encode it later as wished.

Adaptive Token Merge Although the re-encoding of the positional embedding enhances consistency through attention, it may also cause disruptions in the locally geometric relationships of tokens due to positional rearrangement. Additionally, a few tokens may not be matched perfectly, leading to information loss. Considering that attention output is a weighted sum of all values and carries more global information, it helps to compensate for such information loss. Therefore, we introduce adaptive token merge, where the attention outputs of the frame image are interpolated with those from the identity image.

In this process, we first extract the segment attention output corresponding to the foreground image tokens, from both the identity image and frame image. We then rearrange the attention output of identity image based on the point-matching relationships and perform interpolation:

$$H_{fg,id} = H_{i,id}[C_{fg,id}], \quad (18)$$

$$H_{fg,frm}^* = H_{i,frm}^*[C_{fg,frm}], \quad (19)$$

$$H_{fg,frm}^{**} = (1 - \alpha S_{\max}^*) H_{fg,frm}^* + \alpha S_{\max}^* H_{fg,id}', \quad (20)$$

where $\alpha \in [0, 1]$ is a hyper-parameter and decays with the timestep in the generation process. This interpolation process incorporates the averaged similarity S_{\max}^* obtained from Eq. (5) as weights, effectively suppressing tokens with low matching confidence and preventing negative impact.

4. Experiments

4.1. Experimental Setups

Evaluation Datasets This work focuses on consistent generation for characters and environments, while existing benchmarks lack features suited for such tasks. Therefore, we used GPT-4 to generate a series of T2I prompts tailored for such application scenario. Specifically, we instructed GPT-4 to create multiple groups of prompts in the format: “[Environment], [Character], [Action].”, with contents vary across groups. We proposed two evaluation tasks: **background maintaining** and **background switching**. In the former, both “[Environment]” and “[Character]” remain consistent within a group of prompts, while “[Action]” varies. In the latter, only the “[Character]” remains unchanged. Every group consists of 5-8 prompts. For each

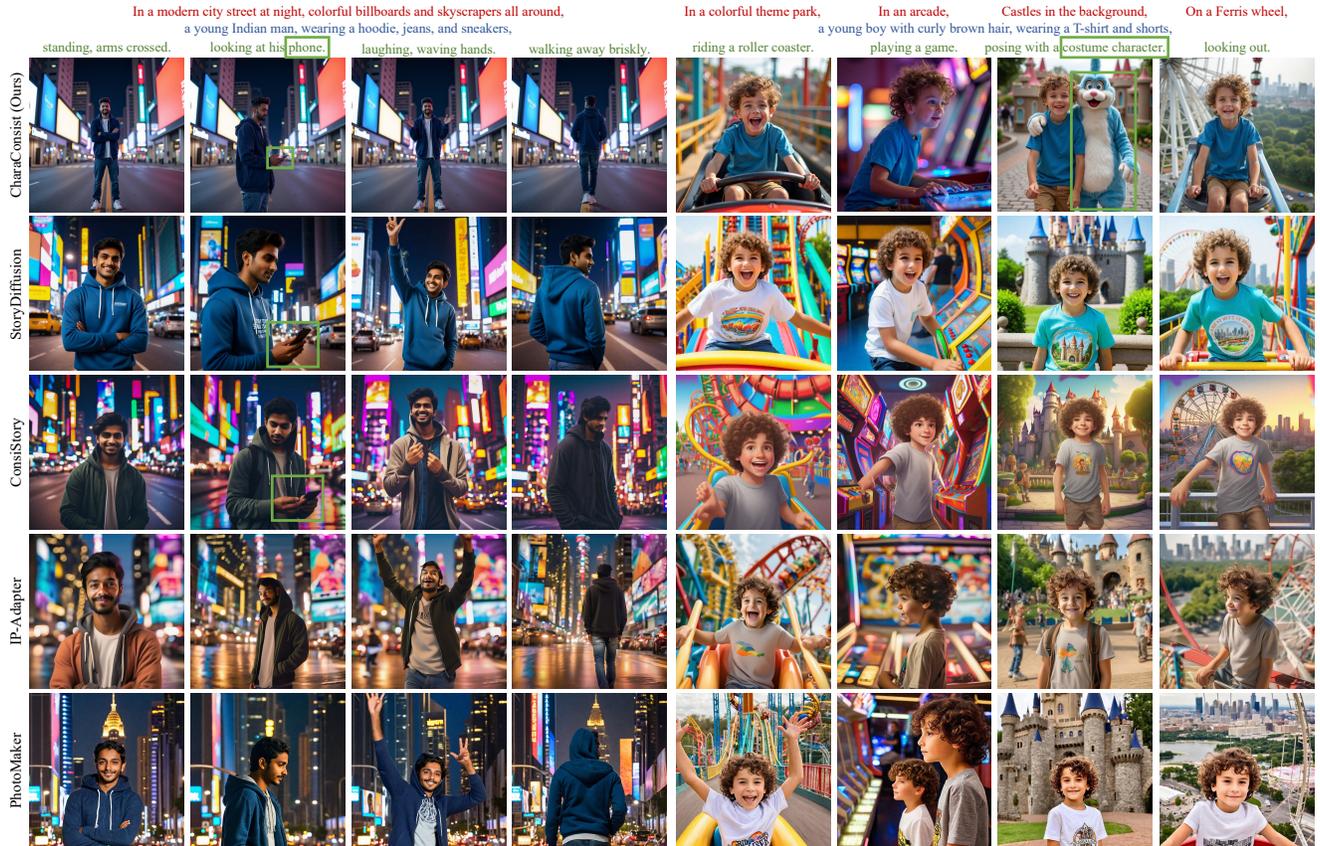


Figure 5. **Qualitative comparisons with previous consistent generation methods and identity-reference methods.** Our method demonstrates superior performance in maintain consistency of multiple aspects such as character identity, clothing and background scenes.

of the two-types tasks, we generated over 200 T2I prompts.

Comparison Methods Comparison methods include both the training-free consistent generation methods StoryDiffusion [35], ConsiStory [30] and the training-dependent identity-reference methods IP-Adapter [34], PhotoMaker [13]. The above generated prompts are shared for all methods. For the identity-reference methods, we first generated the character headshots using the “[Character]” part of prompts, and then cropped the face region using RetinaFace [1] as the input reference face image.

Evaluation Metrics Following previous works [30, 35], we introduce the CLIP text-image similarity (**CLIP-T**) to measure the alignment between images and textual prompts, along with the pairwise CLIP image similarity (**CLIP-I**) to measure the full-image consistency. To better evaluate the characteristics of this work, we introduced three additional metrics. First, we introduce the decoupled CLIP-I, termed as **CLIP-I-fg** and **CLIP-I-bg**, to evaluate the consistency of foreground and background separately. We employ SAM [9] to segment the foreground in each generated image, and compute CLIP similarity separately for the masked foreground and masked background images. We also introduce the identity similarity (**ID Sim**) to evaluate

the identity consistency by compute the facial embedding similarity using RetinaFace [1] and FaceNet [24]. Furthermore, we introduce the image quality score (**IQS**) and image aesthetics score (**IAS**) to evaluate the quality of generated images using Q-Align [32].

4.2. Experimental Results

Qualitative Results In Fig. 5, we present the results of our *CharaConsist* and all comparison methods on both background maintaining and background switching tasks. It can be observed that all comparison methods exhibit varying degrees of clothing inconsistency. In the background maintaining task, these methods can only generate similar scenes but fail to ensure complete consistency. While in the background switching task, these methods focus solely on the foreground subject, leading to deviation from the unique action prompts, e.g., the “*phone*”, “*costume character*” in Fig. 5 and the “*taking photos*” in Fig. 2 are missing. Additionally, we show the points and masks results in Fig. 6, to intuitively demonstrate the effectiveness of our point-matching and mask extraction methods. More qualitative results are shown in the supplementary materials.

Quantitative Results As shown in Tab. 1, compared to

Method	Background Maintaining				ID Sim	Background Switching				Image Quality	
	CLIP-T	CLIP-I	CLIP-I-fg	CLIP-I-bg		CLIP-T	CLIP-I	CLIP-I-fg	ID Sim	IAA	IQA
IP-Adapter [34]	0.358	0.891	0.883	0.881	0.664	0.357	0.833	0.878	0.590	0.811	0.981
PhotoMaker [13]	0.365	0.891	0.881	0.886	0.730	0.371	0.824	0.876	0.674	0.705	0.972
StoryDiffusion [35]	0.370	0.902	0.882	0.895	0.642	0.367	0.853	0.873	0.563	0.777	0.978
ConsiStory [30]	0.368	0.903	0.876	0.895	0.578	0.365	0.846	0.868	0.502	0.815	0.978
<i>CharaConsist</i> (ours)	0.356	0.910	0.883	0.916	<u>0.647</u>	0.360	0.840	0.881	<u>0.587</u>	0.831	0.985

Table 1. **Quantitative comparisons with previous consistent generation methods and identity-reference methods.** The best results within the training-free consistent generation methods are underlined, and the best results across all methods are highlighted in **bold**.



Figure 6. **Point matching and mask extraction results.** The points in the first image are manually selected, while those in the subsequent images are determined by the matching relationship.

the consistent generation methods, across two evaluation tasks, our *CharaConsist* achieves superior performance on most consistency metrics, especially in CLIP-I-bg and ID Sim. These results strongly validate the superiority of our method in maintaining character and scene consistency. As the only exception, with the varying environments in the background switching task, the lower full-image CLIP-I score of our method is reasonable. Compared to the identity-reference methods, *CharaConsist* offers significant advantages in CLIP-I-fg and CLIP-I-bg scores, indicating our method’s ability in preserving the details in character clothing and background environments. Although our method does not surpass these methods on identity similarity, which is the training target for these methods, this is not contradictory. As a training-free method, *CharaConsist* can complement the identity-reference generation by addressing their shortcomings in clothing and environment consistency. We will further discuss this in Sec. 5.

Additionally, one concern is our method has lower CLIP-T score compared to other methods. This could be largely due to the domain differences stemming from the base model, as FLUX.1’s original results also show lower CLIP-T, which is further detailed in the ablation study and Tab. 2.

Ablation Study Since this work is built on a different base model compared to the baseline methods, we conducted ablation study to analyze the impact of the base model itself and our proposed method. For our *CharaConsist*

Metrics	Base Model			Increment		
	CsSt	StDiff	Ours	CsSt	StDiff	Ours
CLIP-T	.368	.377	.358	.0	-.007	-.002
CLIP-I	.867	.882	.874	.036	.020	.036
CLIP-I-fg	.837	.855	.842	.039	.027	.041
CLIP-I-bg	.868	.877	.883	.027	.018	.033
ID Sim	.448	.546	.476	.130	.096	.171

Table 2. **Consistency increments comparison.** The “CsSt” and “StDiff” refer to ConsiStory and StoryDiffusion respectively

and the comparison methods, we evaluate both the original results of the corresponding base models and the increments on consistency. The base models are SDXL, RealVisXL4.0, and FLUX.1 for ConsiStory, StoryDiffusion, and our *CharaConsist*, respectively. The results in Tab. 2 demonstrate that, **first**, compared to the original results of FLUX.1, our method significantly enhances consistency, indicating its effectiveness. **Second**, our method achieves the largest increment in consistency. This indicates that the superior consistency results shown in Tab. 1 are not due to the new base model. Instead, although FLUX.1 provides a notable improvement in image quality, its original results exhibit lower consistency compared to RealVisXL4.0 in almost all metrics. **Furthermore**, The FLUX.1 exhibits a considerably lower CLIP-T score than other two base models, which may due to domain gaps in visual style.

5. Conclusion and Discussion

We propose *CharaConsist* to enhance consistent text-to-image generation. By addressing the locality bias with point-tracking attention and adaptive token merge, it ensures fine-grained consistency under character action variations. It also enables controllable background preserving or switching, broadening real-world applicability.

Limitation As a method focused on training-free consistent generation, our *CharaConsist* can produce consistent contents aligned with textual prompts, but cannot take input identity as reference. An ideal solution would be combining the trained identity-reference model with our method, in which the former could provide the ability to accept input identities and maintain higher facial similarity, while our method could compensate for their limitations in consistent background environments and character clothing. This is an important direction for our future exploration.

Acknowledgement. This project was supported by the National Natural Science Foundation of China (No.92470203), National Natural Science Foundation of China (No.62472104), Beijing Natural Science Foundation (No.L242022), and Fundamental Research Funds for the Central Universities (2024XKRC082).

References

- [1] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *CVPR*, pages 5203–5212, 2020. 7
- [2] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, pages 12873–12883. Computer Vision Foundation / IEEE, 2021. 3
- [3] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis. In *ICML*. OpenReview.net, 2024. 1, 2, 3
- [4] 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 *ICLR*. OpenReview.net, 2023. 4
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 3
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778. IEEE Computer Society, 2016. 3
- [7] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 3
- [8] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014. 3
- [9] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *ICCV*, pages 4015–4026, 2023. 7
- [10] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*, 2024. 3
- [11] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *CVPR*, pages 1931–1941, 2023. 4
- [12] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 1, 2, 3, 4
- [13] Zhen Li, Mingdeng Cao, Xintao Wang, Zhongang Qi, Ming-Ming Cheng, and Ying Shan. Photomaker: Customizing realistic human photos via stacked id embedding. In *CVPR*, pages 8640–8650, 2024. 3, 4, 7, 8
- [14] Zhimin Li, Jianwei Zhang, Qin Lin, Jiangfeng Xiong, Yanxin Long, Xincheng Deng, Yingfang Zhang, Xingchao Liu, Minbin Huang, Zedong Xiao, et al. Hunyuan-dit: A powerful multi-resolution diffusion transformer with fine-grained chinese understanding. *arXiv preprint arXiv:2405.08748*, 2024. 3
- [15] Ziyi Li, Hao Luo, Xincheng Shuai, and Henghui Ding. AnyI2V: Animating any conditional image with motion control generation. In *ICCV*, 2025. 3
- [16] Tao Liu, Kai Wang, Senmao Li, Joost van de Weijer, Fhad Khan, Shiqi Yang, Yaxing Wang, Jian Yang, and Mingming Cheng. One-prompt-one-story: Free-lunch consistent text-to-image generation using a single prompt. In *ICLR*, 2025. 2, 3
- [17] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *ICCV*, pages 4172–4182. IEEE, 2023. 1, 2, 3
- [18] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: improving latent diffusion models for high-resolution image synthesis. In *ICLR*. OpenReview.net, 2024. 1, 2, 3, 4
- [19] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 3
- [20] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020. 3
- [21] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. 1, 3
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015. 3
- [23] 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 *CVPR*, pages 22500–22510, 2023. 4
- [24] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, pages 815–823, 2015. 7
- [25] Xincheng Shuai, Henghui Ding, Xingjun Ma, Rongcheng Tu, Yu-Gang Jiang, and Dacheng Tao. A survey of multimodal-guided image editing with text-to-image diffusion models. *arXiv preprint arXiv:2406.14555*, 2024. 3
- [26] Xincheng Shuai, Henghui Ding, Zhenyuan Qin, Hao Luo, Xingjun Ma, and Dacheng Tao. Free-form motion control: Controlling the 6d poses of camera and objects in video generation. In *ICCV*, 2025. 3

- [27] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, pages 2256–2265. PMLR, 2015. [3](#)
- [28] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. [6](#)
- [29] Luming Tang, Menglin Jia, Qianqian Wang, Cheng Perng Phoo, and Bharath Hariharan. Emergent correspondence from image diffusion. *NeurIPS*, 36:1363–1389, 2023. [3](#), [5](#)
- [30] Yoad Tewel, Omri Kaduri, Rinon Gal, Yoni Kasten, Lior Wolf, Gal Chechik, and Yuval Atzmon. Training-free consistent text-to-image generation. *ACM TOG*, 43(4):1–18, 2024. [2](#), [3](#), [7](#), [8](#)
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [3](#)
- [32] Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang, Erli Zhang, Wenxiu Sun, et al. Q-align: teaching llms for visual scoring via discrete text-defined levels. In *Proceedings of the 41st International Conference on Machine Learning*, pages 54015–54029, 2024. [7](#)
- [33] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. [3](#)
- [34] 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. [3](#), [4](#), [7](#), [8](#)
- [35] Yupeng Zhou, Daquan Zhou, Ming-Ming Cheng, Jiashi Feng, and Qibin Hou. Storydiffusion: Consistent self-attention for long-range image and video generation. *NeurIPS*, 2024. [2](#), [3](#), [7](#), [8](#)