

StoryImager: A Unified and Efficient Framework for Coherent Story Visualization and Completion

Ming Tao^{1,2}, Bing-Kun Bao^{1,2}[†], Hao Tang³, Yaowei Wang², and Changsheng Xu^{2,4,5}

¹ Nanjing University of Posts and Telecommunications

² Peng Cheng Laboratory

³ Peking University & Carnegie Mellon University

⁴ MAIS, Institute of Automation, Chinese Academy of Sciences

⁵ School of Artificial Intelligence, University of the Chinese Academy of Sciences
bingkunbao@njupt.edu.cn

Abstract. Story visualization aims to generate a series of realistic and coherent images based on a storyline. Current models adopt a frame-by-frame architecture by transforming the pre-trained text-to-image model into an auto-regressive manner. Although these models have shown notable progress, there are still three flaws. 1) The unidirectional generation of auto-regressive manner restricts the usability in many scenarios. 2) The additional introduced story history encoders bring an extremely high computational cost. 3) The story visualization and continuation models are trained and inferred independently, which is not user-friendly. To these ends, we propose a bidirectional, unified, and efficient framework, namely StoryImager. The StoryImager enhances the storyboard generative ability inherited from the pre-trained text-to-image model for a bidirectional generation. Specifically, we introduce a Target Frame Masking Strategy to extend and unify different story image generation tasks. Furthermore, we propose a Frame-Story Cross Attention Module that decomposes the cross attention for local fidelity and global coherence. Moreover, we design a Contextual Feature Extractor to extract contextual information from the whole storyline. The extensive experimental results demonstrate the excellent performance of our StoryImager. Code is available at <https://github.com/tobran/StoryImager>.

Keywords: Generative Model · Story Visualization · Story Completion

1 Introduction

The last few years have witnessed the great success of large pre-trained generative models for a variety of applications. Among them, text-to-image synthesis is one of the important tasks of generative models. Due to its practical value, text-to-image synthesis has become an active research area, leading to the development of large pre-trained text-to-image autoregressive and diffusion models,

[†] Corresponding Author.



Fig. 1: (a) Existing models adopt the auto-regressive generative approach, which restricts usability in many scenarios. And the users need to switch between models to meet their current requirements (b) Our proposed StoryImager unifies different tasks into one model, which is more comprehensive to tackle various generative requirements.

e.g., DALL-E [31] and LDM [32]. Based on the powerful image-generative ability, some recent work [22, 26, 29] tried to extend pre-trained text-to-image models beyond generating a single image to generating a sequence of story images, a task referred to as story visualization.

Unlike single text-to-image generation, story visualization aims to generate coherent and visually appealing story images based on a sequence of story descriptions. The task requires the model to capture the essence of the story and the relationships between different story descriptions and then translate them into realistic images that tell a story in a coherent manner. Some work expanded this task by introducing story continuation to generate subsequent story images based on both given story descriptions and historical frames. The importance of story visualization and continuation lies in its potential applications in entertainment, education, and multimedia storytelling.

Although impressive results have been presented in previous work [22, 26, 29], they still suffer from three flaws. First, the capabilities of current models are limited. The autoregressive generative architecture they adopted [22, 26, 29] only supports frame-by-frame generation due to unidirectional attention. It restricts the applicability of models in scenarios that require referencing future or bidirectional frames, such as story backtracking, insertion, and replacement. This limitation prevents users from making arbitrary modifications and extensions based on existing story frames. Second, existing models are trained and inferred independently for different tasks. The current story image generation task includes Story Visualization and Story Continuation. However, existing methods tackle them as two independent tasks. They train two independent models and infer separately. Figure 1(a) shows that independent models require users to store, load, and specify different large models for each requirement. Thus, it increases the complexity of usage and increases hardware requirements. Third, current models introduce large multimodal models such as BLIP [16] to encode history story captions and frames, then stack extracted hidden states directly, as shown in Figure 2(a). The additional large models and stacked long hidden states enlarge the story model and bring an extremely high computing budget.

To address the above issues, we propose a novel unified story visualization and completion framework named StoryImager. For the first issue, we further extend the scope of the tasks and categorize them into story visualization and comple-




Fig. 2: (a) Existing models introduce large models to encode the history information for auto-regressive generation. (b) The storyboard generative ability of Stable Diffusion [32] learned from pretraining process. (c) Our proposed StoryImager inherits the storyboard generative ability and unifies different tasks through a masking strategy.

tion based on the different given conditions. Story visualization is only conditioned on story captions, whereas story completion is conditioned on both story captions or images provided at any arbitrary frame. The new scope fulfills the synthesis requirements and makes the task more comprehensive. Under the new scope, we propose a Storyboard-based Generation (Storyboard-Gen). As shown in Figure 2(b), the large pretraining enables the stable diffusion to generate plausible storyboards. The Storyboard-Gen inherits the pre-trained ability and enhances it for story visualization and completion through Parameter-Efficient Fine-Tuning (PEFT) [23]. The Storyboard-Gen enables the bidirectional synthesis of story images in the storyboard. As shown in Figure 1(b), it expands our model in the story image generation task and allows it to be applied in a broader range of scenarios.

For the second issue, we introduce a Target Frame Masking Strategy (Target-Mask) for training and inference that addresses both story visualization and completion tasks. As shown in Figure 2(c), during training, the frames are completely masked for story visualization and partially masked randomly for story completion. During inference, users simply need to load a single model and mask the frames they want to synthesize, enabling them to handle various requirements seamlessly without reloading other models. The Target-Mask provides a more flexible generation approach to meet the specific generation goals of users.

For the third issue, we design a decomposed Frame-Story Cross Attention Module (Frame-Story CAM). It leverages the inherent cross-attention module of the pre-trained diffusion model and decomposes it into local frame-level and global story-level cross-attention. This decomposed design ensures both the image quality of each frame and the overall consistency of the story images. Furthermore, the self-attention in the storyboard enables visual feature interactions between frames. Thus, our model can effectively fuse story caption and visual features without additional large multimodal encoders.

Overall, our contributions can be summarized as follows:

- We extend the task of story image synthesis and propose a unified story visualization and completion framework to synthesize high-fidelity and coherent story images.
- We propose a Storyboard-based Generation, which enables bidirectional synthesis of story images in the storyboard.

- We introduce a Mask-based training and inference strategy that consolidates various synthesis requirements into masked story image prediction.
- We design a Frame-Story Cross Attention Module, which ensures the image fidelity of individual frames and the overall coherence of the entire story.
- Extensive qualitative and quantitative experiments on two challenging datasets demonstrate that the proposed StoryImager outperforms existing models.

2 Related Work

2.1 Text-to-Image Synthesis

Text-to-image generative models greatly impact story synthesis models and are always adopted as the generative backbone. Contemporary advances in text-to-image synthesis have primarily centered around three primary frameworks: Generative Adversarial Networks, Auto-regressive models, and Diffusion models. Text-to-Image GANs [38–40, 42, 43] employ adversarial training strategies between generators and discriminators. Large-scale autoregressive models, such as DALL·E [31], Make-A-Scene [4], and Parti [41], have exhibited commendable scalability and exceptional proficiency in synthesizing images. Diffusion models [3, 10, 11, 24, 34], including VQ-Diffusion [7], GLIDE [25], DALL-E2 [30], Latent Diffusion Models (LDM) [32], and Imagen [33], have garnered significant interest within the research community.

2.2 Story Visualization and Continuation

StoryGAN [17] first proposes the story visualization task and provides a GAN-based sequence generation model consisting of a deep RNN-based context encoder and two discriminators for images and stories. Subsequently, some works follow and refine this network. For example, CP-CSV [37] introduces a foreground segmentation generation module to optimize the consistency of characters and backgrounds in the story. Both DUCO [21] and VLC [20] enhance semantic consistency by dual learning. Still, the former also considers inter-image sequence consistency via copy-transform, while the latter focuses more on textual information via introducing external common-sense information to complement textual details. Word-Level SV [14] and Clustering GAN [15] simplify the two-level GAN network of StoryGAN [17] and further optimize story quality through word-level fine-grained features and clustering learning, respectively. In addition to work based on GAN models, VQ-VAE-based VP-CSV [2], pre-trained model DALL-E-based StoryDALL-E [22], and diffusion model-based AR-LDM [26] can also tackle story visualization. Notably, to improve task generalization, StoryDALL-E [22] defines a new task story continuation, which introduces the source frame to generate unseen plots and characters. AR-LDM [26] and Make-A-Story [29] transform the diffusion model into an autoregressive story image generative manner to synthesize image sequence. The AR-LDM introduces

the large pre-trained BLIP to encode history stories and frames, and the Make-A-Story proposes a visual memory module that captures the context information. There are some works [6, 18, 44] that focus on open-ended story visualization. TaleCrafter [6] proposes an interactive story visualization by incorporating sketch and layout controls. The CogCartoon [44] proposes a character-plugin generation to alleviate dependence on data and storage. The Intelligent Grimm [18] collects a diverse open-ended story dataset from YouTube and E-books.

3 The Proposed Method

In this work, we expand the scope of story image generation and propose a unified and contextually coherent framework. This framework effectively addresses both the story visualization and completion tasks simultaneously, providing comprehensive coverage for the generation process. To synthesize high-fidelity and coherent story images under different tasks, we propose: (i) a Storyboard-Gen approach that enables bidirectional synthesis of story images in a storyboard. (ii) a Target-Mask strategy for training and inference that consolidates diverse synthesis requirements into masked story image prediction. (iii) a Frame-Story CAM with a Context Feature Extractor (Context-FE) ensures both the visual fidelity of individual frames and the overall coherence of the story. In the following of this section, we first present a comprehensive overview of our StoryImager. Following that, we provide detailed explanations of the proposed Storyboard-Gen, Target-Mask, FrameStory-CAM, and Context-FE in detail.

3.1 Model Overview

As illustrated in Figure 3, our proposed StoryImager consists of a pretrained Text Encoder, a pair of Image Encoder \mathcal{E} and Decoder \mathcal{D} of pretrained autoencoder, a Contextual Feature Extractor (Context-FE), a pretrained diffusion Model [32] with Parameter-Efficient Fine-Tuning (PEFT) [23], a Target Frame Masking Strategy, and Frame-Story Cross Attention Modules (Frame-Story-CAM). The story images are first encoded into latent space, and the text encoder encodes the text descriptions into word embeddings. The Context-FE extracts global contextual information and predicts a frame-aware latent prior. The Target Frame Masking Strategy masks target story images for training and inference. Then the U-Net of the pretrained diffusion model takes the latent features, word embeddings, and contextual information as input and then fuses them through the Frame-Story Cross Attention Module at each layer. The whole model is trained by predicting the noise of the masked parts. Finally, after reversing the diffusion process multisteps, the latent features are decoded to the target images.

3.2 Storyboard-based Generation

Existing story image generation models adopt an auto-regressive architecture for synthesizing multiple images from story captions. They transform the pre-trained text-to-image model into a frame-by-frame generative model, enabling

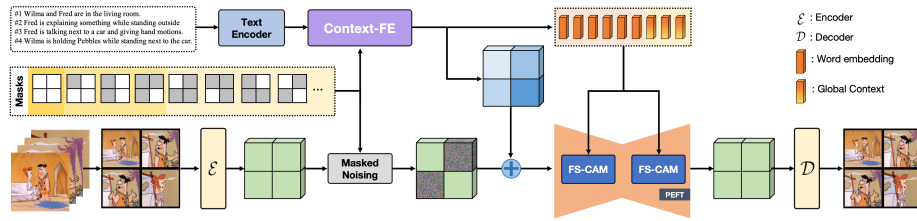


Fig. 3: The architecture of StoryImager for story visualization and completion. It adopts a Storyboard-based Generation approach to enable bidirectional story image generation and unifies different tasks through the Target Frame Masking Strategy.

the generation of story image sequences. However, the limitation of the auto-regressive architecture lies in its unidirectional generation, which hampers its applicability in various scenarios. During the process of creating story images, users may need to modify a specific frame within the story or incorporate new story elements into an existing narrative. These requirements necessitate a model that can refer to the content of existing frames bidirectionally, allowing for effective interaction with the existing story. Nevertheless, unidirectional generation disregards requirements such as story insertion, replacement, and backtracking, significantly constraining the usability of these models in such contexts.

In recent years, large-scale pretrained models have demonstrated remarkable zero-shot capabilities in various natural language processing and computer vision tasks. They keep the structure of pre-trained models while transforming the downstream tasks to closely resemble the approaches used in pretraining. For example, Prompt learning [1, 27, 28] leverages the construction of suitable prompts or contexts to enable GPT [1] to perform text classification tasks and specific text generation [1, 5, 19]. By setting up downstream tasks to closely resemble the pretraining process, we can better utilize the knowledge acquired during pretraining. Inspired by this, we explore the potential of large pretrained diffusion models in generating storyboards. As shown in Figure 2(b), we find that the pretrained diffusion model is able to synthesize plausible storyboards. Thus, a question arises: Can we transform the task of story visualization and completion into the generation of storyboard images? Motivated by it, we introduce Storyboard-based Generation (Storyboard-Gen) as our story image synthesis approach. The Storyboard-Gen arranges the story images in sequential order and places them into panels, forming a storyboard image. In contrast to the auto-regressive generative approach, where diffusion models would need to learn story generation from scratch. Our storyboard-based Generation leverages the inherent storyboard image generative ability of pretrained models and trains them in a manner that aligns closely with their original pretraining objectives. Furthermore, compared with the auto-regressive generative approach, our Storyboard-Gen enables the bidirectional synthesis of story images in a storyboard. It forms the foundation for constructing a unified framework for this task.

3.3 Target Frame Masking Strategy

Previous story image generative models focus on story visualization and continuation. Story continuation is a variant of story visualization that shares the same goal but incorporates a source frame (i.e., the first frame) to guide the generation of subsequent frames. However, the information in a single source frame is limited, and such a setup restricts the application scenarios of the model. Moreover, existing methods tackle story visualization and continuation as two independent tasks. They train two independent models separately. This not only increases the training burden, but also introduces complexity for users, who must constantly switch between models to meet their specific requirements when generating long story images. Therefore, we extend the story image generative task and introduce the Target Frame Masking Strategy (Target-Mask) for training and inference to unify story visualization and completion.

In the forward process, the Target-Mask randomly samples a binary mask m to indicate the desired frames for generation. In the case of story visualization, the mask consists entirely of True, indicating that all frames are target generation frames. For story completion, only the frames requiring completion are marked True in m , while the rest of the frames remain unmasked. This strategy enables us to generate specific frames based on current requirements selectively. Then the Target-Mask applies a masked noising process on the latent of storyboard x . We define $x_0 = x$ and only add noise on the latent of masked frames instead of the whole latent:

$$\begin{aligned}\tilde{x}_t &= \sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon \\ x_t &= \tilde{x}_t \odot m + x_0 \odot (1 - m),\end{aligned}\tag{1}$$

where $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and t is the timestep in the forward process. The Target-Mask enables the model to utilize visual information in given story frames and learn to recover masked frames $x_0 \odot m$. This ensures that generated frames in the mask m are consistent with the given frames. Following [10] we train a network ϵ_θ to predict noise ϵ from noisy x_t :

$$\mathcal{L}_{\text{DM}} = \mathbb{E}_{\epsilon \sim \mathcal{N}(0, \mathbf{I})} [\|m \odot \epsilon - m \odot \epsilon_\theta(x_t, t, s)\|_2^2].\tag{2}$$

where s is the story caption. The typical training scheme used in stable diffusion is to predict the loss on the whole latent. Since our target is to predict masked story frames according to given frames and story captions, we only calculate the loss on masked frames.

In the inference phase, we apply a masked noising process on the target frames in the storyboard $x_T = \epsilon \odot m + x_0 \odot (1 - m)$, where T is the number of sampling steps. The unmasked frames are kept from the denoising process at each step. Then we reverse the diffusion process and obtain the completed storyboard x_0 .

By incorporating Target-Mask, our approach enhances the flexibility and adaptability of the model for both story visualization and completion tasks. It allows us to generate complete storyboards or fill in missing frames seamlessly,




Fig. 4: (a) The architecture of Frame-Story Cross Attention Module. It decomposes the cross-attention module into story-level and frame-level cross-attention to enable local image fidelity and global story coherence. (b) The proposed Contextual Feature Extractor summarizes the whole text information, extracts global contextual information, and predicts the frame-aware latent prior for the U-Net.

based on the provided context and requirements. Compared with previous methods, our masking strategy unifies different tasks in one framework and extends the ability of the story image generative model. It simplifies the training complexity and user interaction process, ensuring a streamlined and user-friendly experience in creating coherent story images.

3.4 Frame-Story Cross Attention Module

To ensure that the generated story images in the storyboard possess both realistic visual content and overall narrative coherence, we introduce a novel module called Frame-Story Cross Attention Module (Frame-Story CAM). As illustrated in Figure 4(a), the Frame-Story CAM decomposes the pretrained cross-attention modules of the diffusion model in each layer into two components: the cross-attention mechanism between local frames and the corresponding word embeddings, and the cross-attention mechanism between the entire storyboard and the extracted contextual story features.

At the frame level, the cross-attention mechanism divides the encoded latent features into four smaller frame-specific latent features at each layer, aligned with the sequence of frames in the storyboard. Then these frame-specific latent features are flattened to the batch dimension. Simultaneously, the word embeddings extracted from each story caption are also arranged in the order of frames and flattened to the batch dimension. Finally, we fuse the flattened frame-specific latent features and their corresponding word embeddings through a cross-attention mechanism. This process incorporates the textual information from each story caption into the visual features of each local story frame, thereby facilitating seamless alignment between the visual and textual information for each local story frame.

At the storyboard level, the cross-attention mechanism performs feature fusion between the entire encoded latent features and the global contextual features extracted by the Contextual Feature Extractor (Context-FE). This fusion ensures consistency between the overall visual representation of the story images and the global storylines.

In particular, we did not introduce separate cross-attention mechanisms from scratch for the frame and storyboard levels. Instead, we adopted the Low-Rank Adaptation (LoRA) [12], which freezes the pretrained model weights and introduces trainable rank decomposition matrices into each layer. Furthermore, both frame-level and storyboard-level cross-attention at each layer utilizes the pretrained parameters of the current layer’s cross-attention. They learn two sets of trainable low-rank decomposition matrices for local and global cross-modal feature fusion. This design of two low-rank decompositions from one pretrained layer allows our Frame-Story CAM to inherit the excellent cross-modal fusion capabilities learned from the pretraining for both local and global fusion, while significantly reducing the number of trainable parameters.

Armed with the proposed Frame-Story CAM, our model can maintain the visual quality of individual frames and ensure the overall coherence of the story. The specifically designed LoRA applied on the frame level and the storyboard level cross attention strikes a balance between model adaptability and parameter efficiency, resulting in an effective and efficient framework for generating high-fidelity and coherent story images.

3.5 Contextual Feature Extractor

Story captions often comprise multiple sentences, and concatenating all of the captions into a single long text can easily exceed the maximum length of the pretrained text encoder. Moreover, the pretrained text encoder is not specifically trained for the combination of multiple sequential story captions, which may result in the loss of sequential information across frames. To effectively extract the contextual information of the entire story, we propose the Contextual Feature Extractor (Context-FE).

As shown in Figure 4(b), Context-FE receives word embeddings from each story caption and incorporates additional learnable query embeddings to summarize the information in the word embeddings of each frame. We employ two types of learnable query embedding: context query embeddings and prior query embeddings. These query embeddings are stacked along the length dimension and fed into a network consisting of two layers of Transformers. Using this network, we obtain context embeddings and prior embeddings of each frame’s story caption. Then, the prior embeddings are fed into an MLP (Multilayer Perceptron) to predict frame-aware latent prior. The frame-aware latent prior is reshaped to match the size of the latent representation, which is $4 \times 64 \times 64$. Afterward, the frame-aware latent prior is added to the input latent representation. This incorporation of the prior information provides the Frame-Story CAM with valuable contextual cues, enabling it to better differentiate between different frames. Simultaneously, the context embeddings of each frame are concatenated and fed into a context summarizer consisting of two layers of Transformers. The global context feature summarized from the context summarizer contains the contextual information of the entire story caption. Incorporating the global context feature allows the model to effectively leverage the comprehensive story context for coherent story image synthesis.

4 Experiments

4.1 Datasets

We evaluate our approach on two challenging datasets: Pororo-SV [13] and Flintstones-SV [8]. Each story in these two datasets contains five consecutive frames. The Pororo-SV dataset consists of 10191/2334/2208 samples in the training, validation, and testing sets, respectively. The Flintstones-SV dataset consists of 20132/2071/2309 samples in the training, validation, and testing sets, respectively. The image resolution of these two datasets is 128×128 . We resize the images to 248 and pad them with 4 pixels, resulting in a final image size of 256×256 . The storyboard consists of four-story images with a resolution of 512×512 . Each image of Pororo-SV and Flintstones-SV corresponds to 1 story description. The partitioning of these two datasets into training, validation, and testing subsets is conducted in line with the established practices in previous studies [22, 26, 29].

4.2 Training and Evaluation Details

Our method is based on the stable diffusion model with publicly available checkpoints v1.5. We freeze the pretrained autoencoder and text encoder and finetune the U-Net through LoRA with $\alpha = 4$, $r = 128$. We use the AdamW optimizer to train our model. We set the learning rate 0.001 for the U-Net and 0.0001 for the Contextual Feature Extractor. All models were trained on $8 \times$ NVIDIA RTX A6000 GPUs. The network is trained 300 epochs and 150 epochs on Pororo-SV and Flintstones-SV for about 20 hours. Following the previous story visualization and continuation works [14, 22, 26, 29], we adopt the Fréchet Inception Distance (FID) [9] to evaluate the image fidelity of synthesized story images. We also adopt the Fréchet Story Distance (FSD) [14, 37] to evaluate the overall quality of the story image sequence. Different from FID, the FSD considers temporal consistency. The FID and FSD provide a robust measure to evaluate the image fidelity and story consistency of the synthesized story image sequences. During the evaluation, we sample images using the DDIM scheduler [35] for 50 inference steps with a guidance scale set to 6.0 and crop the generated images from the storyboard. In addition, we also conduct large-scale human evaluations regarding visual quality, visual consistency, and story relevance.

4.3 Quantitative Evaluation

To evaluate the performance of our proposed StoryImager, we compare it with several state-of-the-art story visualization and continuation methods [14, 20, 22, 26, 37], which have achieved impressive results. The comparison results for Pororo-SV and Flintstones-SV are shown in Table 1. From the table, we can observe that our StoryImager achieves better FID and FSD compared to other models on both story visualization and continuation tasks. Our model also surpasses the recently proposed Causal-Story [36], which even adopts additional

Table 1: The results of FID and FSD compared with the state-of-the-art models on the test set of Pororo-SV and Flintstones-SV. \times indicates that the model does not support this task.

Dataset	Method	Story Visualization		Story Continuation		Story Completion	
		FID ↓	FSD ↓	FID ↓	FSD ↓	FID ↓	FSD ↓
Pororo-SV [13]	StoryGAN [17]	78.64	94.53	\times	\times	\times	\times
	CP-CSV [37]	67.76	71.51	\times	\times	\times	\times
	DUCO [21]	95.17	171.70	\times	\times	\times	\times
	VLC [20]	94.30	122.07	\times	\times	\times	\times
	WL-SV [14]	56.08	52.50	\times	\times	\times	\times
	StoryDALL-E [22]	\times	\times	25.90	45.70	\times	\times
	MEGA-StoryDALL-E [22]	\times	\times	23.48	-	\times	\times
	Make-A-Story [29]	27.33	51.20	22.66	44.22	\times	\times
	AR-LDM [26]	16.59	35.33	17.40	37.52	\times	\times
	Causal-Story [36]	16.28	-	16.98	-	\times	\times
StoryImager (Ours)	15.63	28.13	15.45	27.10	14.72	26.77	
Flintstones-SV [8]	StoryGAN [17]	90.55	122.71	\times	\times	\times	\times
	StoryGANc [22]	-	-	90.29	-	\times	\times
	WL-SV [14]	72.37	91.30	\times	\times	\times	\times
	StoryDALL-E [22]	\times	\times	26.49	54.30	\times	\times
	MEGA-StoryDALL-E [22]	\times	\times	23.58	-	\times	\times
	Make-A-Story [29]	36.55	53.10	23.74	52.08	\times	\times
	AR-LDM [26]	23.59	39.70	19.28	43.32	\times	\times
	Causal-Story [36]	-	-	19.03	-	\times	\times
StoryImager (Ours)	22.27	36.51	18.32	35.33	18.11	34.20	

Table 2: The computational cost and time requirements of AR-LDM and our StoryImager on Pororo-SV dataset.

Method	Training				Inference			
	GPU-Memory	Time	Tasks	Models	GPU-Memory	Speed	Tasks	Models
AR-LDM [26]	38G+38G	52h+52h	2	2	16G+16G	14.5s	2	2
StoryImager	12G	30h	>2	1	8G	8.0s	>2	1

causal reasoning to the AR-LDM framework. The improvements are more obvious in FSD, which evaluates the overall coherence of the synthesized story images. The results demonstrate that our model excels at synthesizing high-fidelity and contextually coherent story images. Furthermore, our model has the ability to perform story completion, a capability that was absent in previous models. In the story completion testing, we randomly mask parts of the story image and then predict the missing portions. This testing process does not include story visualization. We list the performance of StoryImager on the story completion task, which is a unique feature of our model.

To demonstrate the efficiency of our framework, we conducted a comparison with the state-of-the-art AR-LDM framework [26]. AR-LDM utilizes an autoregressive approach for story image generation and has been adopted as a foundational framework [36]. Table 2 presents the cost that is evaluated on $8\times$ and $1\times$ A6000 with batch size 1 for training and inference, respectively. We assess

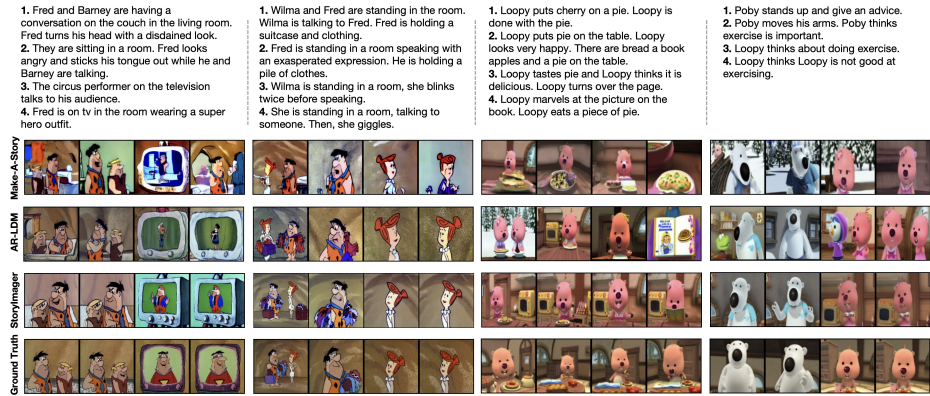


Fig. 5: Comparison of story visualization results between Make-A-Story, AR-LDM, and our proposed StoryImager on Flintstones-SV and Pororo-SV datasets.

resource and time requirements of inference on a single GPU, as it better reflects the usage of most users. As shown in Table 2, our model achieves improved performance and extended tasks while reducing hardware and time requirements. The results indicate that our model is a more effective and efficient framework for story visualization and completion.

4.4 Qualitative Evaluation

Figures 5 and 6 show examples of visual comparisons between our StoryImager, AR-LDM [26], and Make-A-Story [29] in story visualization and continuation tasks, respectively. For story visualization, the characters and backgrounds synthesized by AR-LDM and Make-A-Story are not consistent with other frames. As the synthesized story images are shown in Figure 5 of Pororo-SV, the appearance of “Poby” and the background of “Loopy” are inconsistent with other frames. The same issue also exists in the Flintstones-SV dataset, such as the disappearance of the “couch” in the second frames and the transformation of the “television”. For story continuation, we also observed the inconsistent issue of AR-LDM and Make-A-Story. Additionally, there is a potential for the visual features from the first frame to be disregarded during the generation of subsequent frames. Compared with AR-LDM and Make-A-Story, our StoryImager can synthesize contextual coherent images during story visualization and incorporate the visual features from the first frame during the story continuation.

Moreover, we show the story completion capability of StoryImager, as illustrated in Figure 7. Our StoryImager is the first story completion model. It can generate the first frame based on given frames and the textual description of the target frame, enabling story backtracking functionality. Additionally, our model supports adding new frames to the current story image, achieving story infilling.



Fig. 6: Comparison of story continuation results between Make-A-Story, AR-LDM, and our proposed StoryImager on Flintstones-SV and Pororo-SV datasets.

Table 3: Human evaluation results of story visualization and continuation tasks on Pororo-SV and Flintstones-SV datasets.

Dataset	Criterion(Avg)	Visualization			Continuation		
		Make-A-Story [29]	AR-LDM [26]	Ours	Make-A-Story [29]	AR-LDM [26]	Ours
Pororo-SV	VQ ↓	2.40	1.89	1.71	2.35	1.87	1.78
	VC ↓	2.63	1.92	1.45	2.60	1.85	1.55
	SR ↓	2.52	1.82	1.66	2.48	1.80	1.72
Flintstones-SV	VQ ↓	2.45	1.85	1.70	2.39	1.86	1.75
	VC ↓	2.74	1.89	1.37	2.70	1.82	1.48
	SR ↓	2.66	1.75	1.59	2.61	1.74	1.65

4.5 Human Evaluation

We conducted human evaluation to assess the Visual Quality (VQ), Visual Consistency (VC), and Story Relevance (SR) of the generated results. For each story, we had 12 evaluators rank the synthesized image sequences from our model and the two other top-performing models [26,29] on a scale of 1 to 3. We provide the ground truth as a reference. Each model generated 300 story image sequences for each dataset and task. As shown in Table 3, our model achieves the best performance in terms of all these three criteria. Owing to the powerful storyboard-Gen and decomposed Frame-Story CAM, StoryImager significantly outperforms other models in Visual Consistency (VC).

4.6 Ablation Study

To verify the effectiveness of different components in the proposed StoryImager, we conduct ablation studies on the story visualization and continuation tasks, respectively. The results of Pororo-SV and Flintstones-SV are shown in




Fig. 7: Examples of story backtracking and infilling synthesized by our StoryImager.

Table 4: The performance of different components of our model on the test set of Pororo-SV and Flintstones-SV.

Task	Method	Pororo-SV		Flintstones-SV	
		FID (\downarrow)	FSD (\downarrow)	FID (\downarrow)	FSD (\downarrow)
Story Visualization	Baseline	21.14	41.18	28.19	42.78
	+ Contextual-FE	18.02	36.34	25.75	39.56
	+ Frame-Story CAM	15.63	28.13	22.27	36.51
Story Completion	Baseline	20.65	40.10	23.13	41.15
	+ Contextual-FE	18.06	34.19	21.82	37.88
	+ Frame-Story CAM	14.72	26.77	18.11	34.20

Table 4. The components being evaluated include Contextual Feature Extractor (Contextual-FE), and Frame-Story Cross Attention Module (Frame-Story CAM). Our baseline is a modified Stable diffusion [32] adapted for story visualization and continuation tasks with a target masking strategy. The baseline concatenates all the story captions into a single long text.

From Table 4, we observe that the proposed Contextual-FE can significantly reduce FID and FSD on these two tasks. If we further introduce the Frame-Story CAM to decompose the cross attention, we can observe a further improvement of FID and FSD. The ablation studies demonstrate the effectiveness of our proposed modules in both story visualization and completion tasks.

5 Conclusion

In this paper, we propose a novel unified and contextual coherent framework for the story visualization and completion model, namely StoryImager. The StoryImager inherits the storyboard generative ability of a large pre-trained text-to-image model and extends the story image generation task. Specifically, we introduce a Target Frame Masking Strategy to unify different story image generation tasks. Furthermore, we propose a Frame-Story Cross Attention Module that decomposes the cross attention for local frame fidelity and global storyboard coherence. Moreover, we design a Contextual Feature Extractor, which extracts global context information and synthesizes context visual features. Our StoryImager achieves significant improvements on two challenging datasets.

Acknowledgements

This work was supported by the National Natural Science Foundation of China under Grants (No.62325206, 61936005), the Key Research and Development Program of Jiangsu Province under Grant BE2023016-4, the Postgraduate Research & Practice Innovation Program of Jiangsu Province under Grant KYCX22_0947, and the Fundamental Research Funds for the Central Universities, Peking University.

References

1. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. In: NeurIPS (2020)
2. Chen, H., Han, R., Wu, T.L., Nakayama, H., Peng, N.: Character-centric story visualization via visual planning and token alignment. arXiv preprint arXiv:2210.08465 (2022)
3. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. NeurIPS (2021)
4. Gafni, O., Polyak, A., Ashual, O., Sheynin, S., Parikh, D., Taigman, Y.: Make-a-scene: Scene-based text-to-image generation with human priors. In: ECCV (2022)
5. Gao, T., Fisch, A., Chen, D.: Making pre-trained language models better few-shot learners. arXiv preprint arXiv:2012.15723 (2020)
6. Gong, Y., Pang, Y., Cun, X., Xia, M., Chen, H., Wang, L., Zhang, Y., Wang, X., Shan, Y., Yang, Y.: Talecrafter: Interactive story visualization with multiple characters. arXiv preprint arXiv:2305.18247 (2023)
7. Gu, S., Chen, D., Bao, J., Wen, F., Zhang, B., Chen, D., Yuan, L., Guo, B.: Vector quantized diffusion model for text-to-image synthesis. In: CVPR (2022)
8. Gupta, T., Schwenk, D., Farhadi, A., Hoiem, D., Kembhavi, A.: Imagine this! scripts to compositions to videos. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 598–613 (2018)
9. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: NeurIPS (2017)
10. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. NeurIPS (2020)
11. Ho, J., Saharia, C., Chan, W., Fleet, D.J., Norouzi, M., Salimans, T.: Cascaded diffusion models for high fidelity image generation. JMLR **23**, 47–1 (2022)
12. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. arXiv preprint arXiv:2106.09685 (2021)
13. Kim, K.M., Heo, M.O., Choi, S.H., Zhang, B.T.: Deepstory: Video story qa by deep embedded memory networks. arXiv preprint arXiv:1707.00836 (2017)
14. Li, B.: Word-level fine-grained story visualization. In: European Conference on Computer Vision. pp. 347–362. Springer (2022)
15. Li, B., Torr, P.H., Lukasiewicz, T.: Clustering generative adversarial networks for story visualization. In: Proceedings of the 30th ACM International Conference on Multimedia. pp. 769–778 (2022)

16. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: International Conference on Machine Learning. pp. 12888–12900. PMLR (2022)
17. Li, Y., Gan, Z., Shen, Y., Liu, J., Cheng, Y., Wu, Y., Carin, L., Carlson, D., Gao, J.: Storygan: A sequential conditional gan for story visualization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6329–6338 (2019)
18. Liu, C., Wu, H., Zhong, Y., Zhang, X., Xie, W.: Intelligent grimm–open-ended visual storytelling via latent diffusion models. arXiv preprint arXiv:2306.00973 (2023)
19. Liu, X., Ji, K., Fu, Y., Tam, W., Du, Z., Yang, Z., Tang, J.: P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers). pp. 61–68 (2022)
20. Maharana, A., Bansal, M.: Integrating visuospatial, linguistic and commonsense structure into story visualization. arXiv preprint arXiv:2110.10834 (2021)
21. Maharana, A., Hannan, D., Bansal, M.: Improving generation and evaluation of visual stories via semantic consistency. arXiv preprint arXiv:2105.10026 (2021)
22. Maharana, A., Hannan, D., Bansal, M.: Storydall-e: Adapting pretrained text-to-image transformers for story continuation. arXiv preprint arXiv:2209.06192 (2022)
23. Mangrulkar, S., Gugger, S., Debut, L., Belkada, Y., Paul, S., Bossan, B.: Peft: State-of-the-art parameter-efficient fine-tuning methods. <https://github.com/huggingface/peft> (2022)
24. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: ICML (2021)
25. Nichol, A.Q., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., Mcgrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In: ICML (2022)
26. Pan, X., Qin, P., Li, Y., Xue, H., Chen, W.: Synthesizing coherent story with auto-regressive latent diffusion models. arXiv preprint arXiv:2211.10950 (2022)
27. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. OpenAI blog **1**(8), 9 (2019)
28. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J., et al.: Exploring the limits of transfer learning with a unified text-to-text transformer. JMLR **21**(140), 1–67 (2020)
29. Rahman, T., Lee, H.Y., Ren, J., Tulyakov, S., Mahajan, S., Sigal, L.: Make-a-story: Visual memory conditioned consistent story generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2493–2502 (2023)
30. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 (2022)
31. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: ICML (2021)
32. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)
33. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S.K.S., Ayan, B.K., Mahdavi, S.S., Lopes, R.G., et al.: Photorealistic text-to-image diffusion models with deep language understanding. arXiv preprint arXiv:2205.11487 (2022)

34. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: ICML (2015)
35. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
36. Song, T., Cao, J., Wang, K., Liu, B., Zhang, X.: Causal-story: Local causal attention utilizing parameter-efficient tuning for visual story synthesis. arXiv preprint arXiv:2309.09553 (2023)
37. Song, Y.Z., Rui Tam, Z., Chen, H.J., Lu, H.H., Shuai, H.H.: Character-preserving coherent story visualization. In: Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16. pp. 18–33. Springer (2020)
38. Tao, M., Bao, B.K., Tang, H., Xu, C.: Galip: Generative adversarial clips for text-to-image synthesis. In: CVPR (2023)
39. Tao, M., Tang, H., Wu, F., Jing, X.Y., Bao, B.K., Changsheng, X.: Df-gan: A simple and effective baseline for text-to-image synthesis. In: CVPR (2022)
40. Xu, T., Zhang, P., Huang, Q., Zhang, H., Gan, Z., Huang, X., He, X.: Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In: CVPR (2018)
41. Yu, J., Xu, Y., Koh, J.Y., Luong, T., Baid, G., Wang, Z., Vasudevan, V., Ku, A., Yang, Y., Ayan, B.K., et al.: Scaling autoregressive models for content-rich text-to-image generation. arXiv preprint arXiv:2206.10789 (2022)
42. Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.N.: Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: ICCV (2017)
43. Zhu, M., Pan, P., Chen, W., Yang, Y.: Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis. In: CVPR (2019)
44. Zhu, Z., Tang, J.: Cogcartoon: Towards practical story visualization. arXiv preprint arXiv:2312.10718 (2023)