

# Improved Visual Story Generation with Adaptive Context Modeling

Zhangyin Feng<sup>1</sup>, Yuchen Ren<sup>2</sup>, Xinmiao Yu<sup>1</sup>,  
Xiaocheng Feng<sup>1,3</sup>, Duyu Tang, Shuming Shi, Bing Qin<sup>1,3</sup>

<sup>1</sup> Harbin Institute of Technology,

<sup>2</sup> Renmin University of China,

<sup>3</sup> Peng Cheng Laboratory

{zyfeng, xmyu, xcfeng, qinb}@ir.hit.edu.cn

siriusren@ruc.edu.cn

## Abstract

Diffusion models developed on top of powerful text-to-image generation models like Stable Diffusion achieve remarkable success in visual story generation. However, the best-performing approach considers historically generated results as flattened memory cells, ignoring the fact that not all preceding images contribute equally to the generation of the characters and scenes at the current stage. To address this, we present a simple method that improves the leading system with adaptive context modeling, which is not only incorporated in the encoder but also adopted as additional guidance in the sampling stage to boost the global consistency of the generated story. We evaluate our model on PororoSV and FlintstonesSV datasets and show that our approach achieves state-of-the-art FID scores on both story visualization and continuation scenarios. We conduct detailed model analysis and show that our model excels at generating semantically consistent images for stories.

## 1 Introduction

Diffusion models trained on broad text-image data (Rombach et al., 2022; Bao et al., 2022; Feng et al., 2022; Ramesh et al., 2022; Saharia et al., 2022; Balaji et al., 2022; Nichol et al., 2021) achieved remarkable success in text-to-image generation and showed strong abilities to synthesize photorealistic images of high resolution and great semantic consistency to text prompts. Such a huge success drives the extension of modern diffusion text-to-image models into more scenarios like visual story generation, which is to generate a series of images for a story of multiple sentences.

A recent work, AR-LDM (Pan et al., 2022), which is built upon open-sourced Stable Diffusion, achieves the state-of-the-art FID on the benchmark datasets for visual story generation. AR-LDM encodes previous text-image context as a sequence of additional conditions, which is then attended by






Id	Text	Image
1	Poby talks and gathers his hands. Poby, Loopy and Pororo are clapping their hands.	
2	There are eight glasses of different colors of fluids. Eddy is standing in front of the glasses.	
3	Eddy is holding two sticks and talking. Poby, Loopy, Pororo and Crong are sitting around the table.	
4	Poby, Loopy, Pororo and Crong are clapping.	
5	Eddy is standing in front of the eight glasses containing different colors of fluids.	

Figure 1: A motivating example of a story with five sentences. Blue and purple lines indicate the dependencies between images.

the UNet decoder for image generation. Despite its remarkable success, one limitation is that previous text-image pairs of the same story are all flattened as conditioning memories. This is different from the fact that not all the scenes/characters of sentences in the same story are closely related. Take Figure 1 as an example. The scene of the fourth sentence is not related to either the second or the third sentence. On the contrary, the generation of the fifth image should depend more on the second and third images than others. From this example, we can see that the dependency between images could be largely measured by the semantic relations between sentences.

In this work, we present a simple approach<sup>1</sup> that selectively adopts historical text-image data from the same story in the generation of an image. Specifically, we freeze the text and image representations produced by off-the-shell encoders, and adaptively compute conditioning vectors of context

<sup>1</sup>We name our model as **ACM-VSG** (Adaptive **C**ontext **M**odeling for Visual Story **G**eneration).

by considering the semantic relation between the current sentence and all the history. Such resulting conditioning vectors will be queried by UNet in a traditional way. Furthermore, based on the consideration that images should have similar scenes and characters if their corresponding sentences are similar, we further add context-aware guidance like the use of classifier guidance or CLIP guidance (Nichol et al., 2021) in standard text-to-image generation.

To validate the effectiveness of our approach, we evaluate our model on story visualization and continuation tasks. Experimental results on PororoSV and FlintstonesSV datasets show that both adaptive encoder and guidance improve the quality of the generated images as well as the global consistency of the visual story. The contributions of this work are as follows:

- We present a diffusion model that adaptively uses context information in the encoder and sampling guidance for visual story generation.
- Our approach achieves state-of-the-art results on benchmark datasets for both story visualization and continuation tasks.
- We show that our model excels at generating semantically consistent images for stories.

## 2 Related work

### 2.1 Text-to-Image Generation

We group modern text-to-image generation approaches into three categories. The **first** category is generative adversarial network (Goodfellow et al., 2014; Reed et al., 2016; Zhang et al., 2017). They jointly learn a generator and a discriminator, where the generator is trained to generate images to fool the discriminator and the discriminator is trained to distinguish between real and (generated) fake images. The **second** category is encoder-decoder plus discrete variational autoencoder (dVAE). Methods are developed based on a well-trained discrete variational autoencoder (Van Den Oord et al., 2017), which is capable of mapping an image to discrete tokens and reconstructing an image from discrete tokens. Thus, the task of text-to-image generation could be viewed as a special translation task that converts natural language tokens to image tokens. Autoregressive models (Ramesh et al., 2021a; Ding et al., 2021; Gafni et al., 2022; Yu et al., 2022) typically use Transformer (Vaswani et al., 2017) to

generate a visual token conditioned on the previously generated tokens, resulting in high latency in the inference stage. Muse (Chang et al., 2023) is a non-autoregressive model that tremendously speeds up the inference stage by generating image tokens in parallel. The **third** category is diffusion models — image generation is considered as an iterative refinement process, where two ends of the spectrum are the Gaussian noise and the real image, respectively. Some studies adopt a variational autoencoder to compress an image to the latent space and learn the diffusion process in the latent space of images (Rombach et al., 2022; Bao et al., 2022; Feng et al., 2022). Some works (Ramesh et al., 2022; Saharia et al., 2022; Balaji et al., 2022) directly learn the diffusion model over pixels and typically include cascaded up-sampling models (e.g., from  $64\times 64$  to  $256\times 256$  and from  $256\times 256$  to  $1024\times 1024$ ) to produce high-resolution images.

### 2.2 Visual Story Generation

Visual story generation includes two settings: story visualization and story continuation. Story visualization was firstly introduced by Li et al. (2019), who proposes the StoryGAN model for sequential text-to-image generation. Based on the GAN network, they proposed to combine image and story discriminators for adversarial learning. To improve the global consistency across dynamic scenes and characters in the story, Zeng et al. (2019) jointly considers story-to-image-sequence, sentence-to-image, and word-to-image-patch alignment by proposing an aligned sentence encoder and attentional word encoder. Li et al. (2020) includes dilated convolution in the discriminators to expand the receptive field of the convolution kernel in the feature maps and weighted activation degree to provide a robust evaluation between images and stories. To improve the visual quality, coherence and relevance of generated images, Maharana et al. (2021a) extends the GAN structure by including a dual learning framework that utilizes video captioning to reinforce the semantic alignment between the story and generated images, and a copy-transform mechanism for sequentially consistent story visualization. Maharana and Bansal (2021a) improves the generation quality by incorporating constituency parse trees, commonsense knowledge, and visual structure via bounding boxes and dense captioning. Unlike the story visualization task, whose input only contains the text story, the story continuation

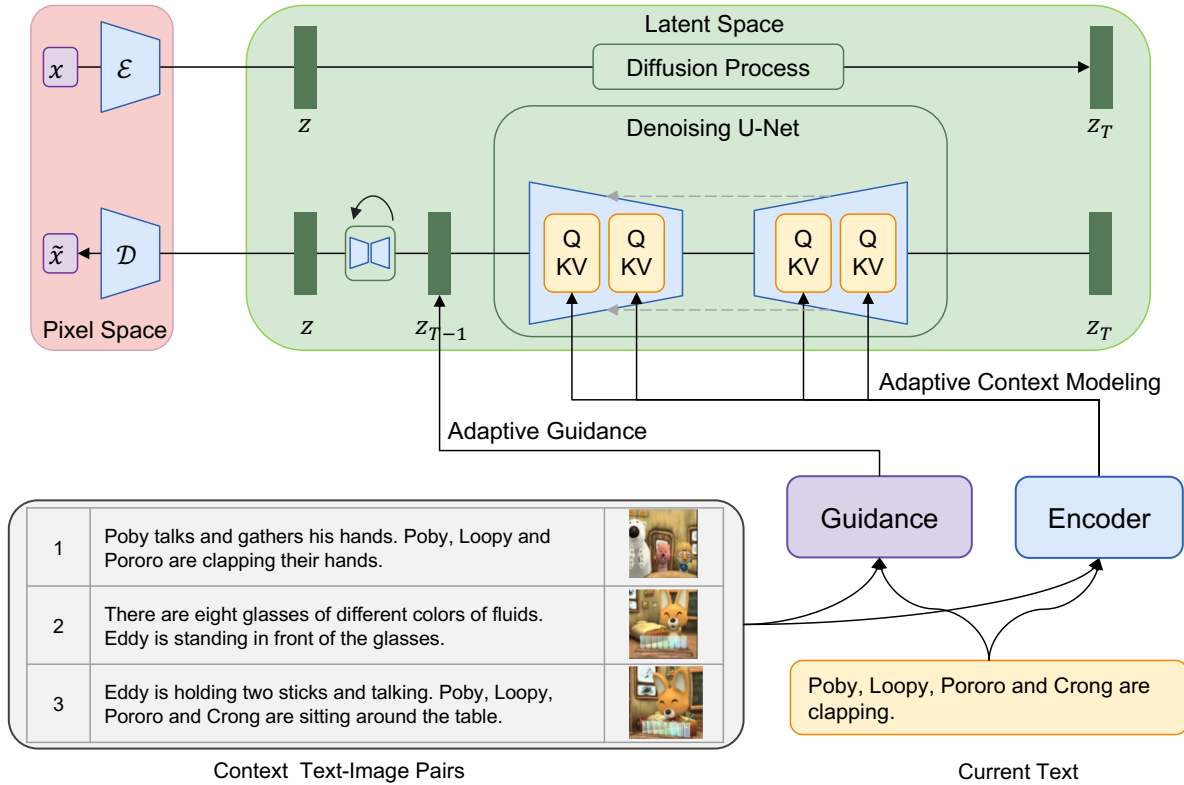


Figure 2: An overview of our model architecture. Based on the latent diffusion model (Rombach et al., 2022), we propose an adaptive encoder and an adaptive guidance. Adaptive encoder is used to get the adaptive context vectors. Conditional diffusion module transforms context vectors to image. Adaptive guidance aims to guide diffusion sampling process with adaptive context information.

task also includes the first image as input. Maharana et al. (2022) introduces story continuation and modifies the pre-trained text-to-image model DALL-E (Ramesh et al., 2021b) by adding a cross attention module for story continuation. Pan et al. (2022) employs a history-aware encoding to incorporate previously generated text-image history to diffusion model for visual story generation.

### 3 Model

We introduce our approach in this section. We first present the model architecture of our approach (§3.1), and then describe three important components: adaptive encoder (§3.2), conditional diffusion model (§3.3) and adaptive guidance (§3.4).

#### 3.1 Model Architecture

An overview of the approach is depicted in Figure 2. It includes an adaptive encoder, a conditional diffusion model, and an adaptive guidance. Based on current text prompt and historical text-image context, the adaptive encoder represents them as conditional vectors. Then the conditional diffusion model transforms these vectors into the correspond-

ing image. During the diffusion sampling process, the adaptive guidance component further guides each diffusion step by comparing it to similar preceding images in the current story to enhance the global consistency of the generated images.

#### 3.2 Adaptive Encoder

Given a story  $\mathbf{S}$  which consists of a sequence of text prompts:  $\mathbf{S} = \{s_1, s_2, \dots, s_L\}$ . Story visualization aims to generate a sequence of images  $\mathbf{X} = \{x_1, x_2, \dots, x_L\}$ . Each image corresponds to a text prompt. Different from text-to-image generation, which only generates one isolated image for the text prompt, story visualization requires global consistency between the generated images. A natural idea is to combine historical text-image context when generating the current image.

$$\begin{aligned}
 P(\mathbf{X}|\mathbf{S}) &= \prod_{i=1}^L P(x_i|\hat{x}_{<i}, \mathbf{S}) \\
 &= \prod_{i=1}^L P(x_i|\tau_\theta(\hat{x}_{<i}, s_{\leq i}))
 \end{aligned}$$

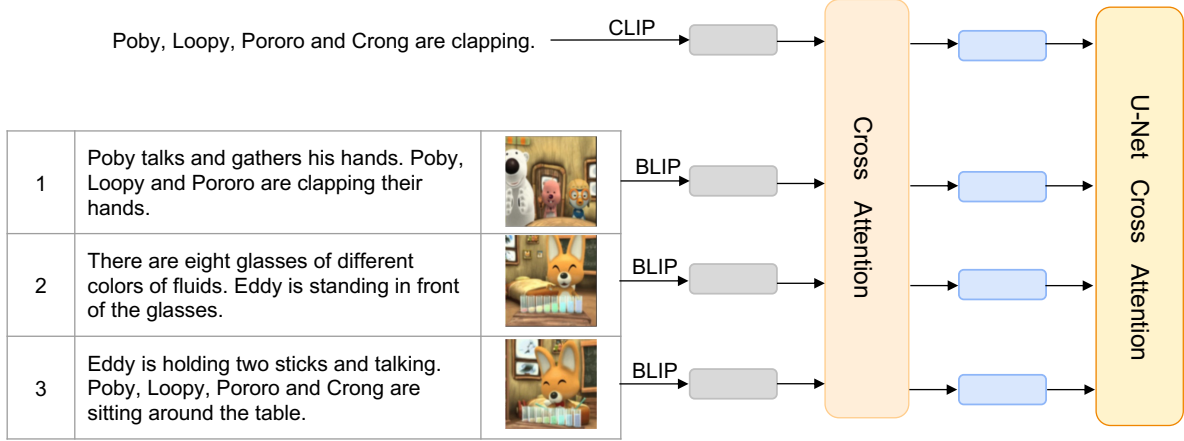


Figure 3: Adaptive encoder consists of three modules: (1) CLIP text encoder is used to encode text prompt. (2) BLIP encoder is used to encode historical text-image pair. (3) Cross attention module is used to filter useful information adaptively.

where  $\tau_\theta$  denotes the history-aware conditioning encoder.

As shown in Figure 1, we find that some images in the history of the same story are similar to the current image, and some images are even completely irrelevant. The purpose of the adaptive encoder is to automatically find the relevant historical text-image pairs, and then encode them into the condition vectors. As shown in Figure 3, adaptive encoder consists of a CLIP text encoder, a BLIP text-image encoder and a cross attention module. Both CLIP (Radford et al., 2021) and BLIP (Li et al., 2022a) are multimodal pre-trained models. The difference is that CLIP encodes text and image respectively, and BLIP can jointly represent text-image pair. We use CLIP to get the current text prompt vector  $v_i$ , and BLIP to get the historical vectors  $\{h_0, \dots, h_{i-1}\}$ . Then a cross attention is equipped to filter history information and we can obtain the updated vectors  $\{\hat{h}_0, \dots, \hat{h}_{i-1}\}$ . In the cross attention module, the text vector  $v_i$  is the query, and each historical vector  $h_{<i}$  is the key and value. Finally, we concatenate the current text vector and history vectors to get the final condition vector  $c = [v_i; \hat{h}_0; \dots; \hat{h}_{i-1}]$ .

### 3.3 Conditional Diffusion Model

Denosing diffusion probabilistic models are a class of score-based generative models, which have recently gained traction in the field of text-to-image generation (Ho et al., 2020). A diffusion model typically contains forward and reverse processes. Given a data  $x_0$  sampled from a real-world data distribution  $q(x)$ , the forward process is imple-

mented as a predefined Markov chain that gradually corrupts  $x_0$  into an isotropic Gaussian distribution  $x_T \sim \mathcal{N}(0, I)$  in  $T$  steps:

$$x_t = \sqrt{\alpha_t}x_{t-1} + \sqrt{1 - \alpha_t}\epsilon_t, \quad t \in \{1, \dots, T\}$$

where  $\epsilon_t \sim \mathcal{N}(0, I)$ , and  $\{\alpha_t \in (0, 1)\}_{t=1}^T$  is a predefined noise variance schedule. The reverse process aims to learn a denoising network  $\epsilon_\theta(\cdot)$  to reconstruct the data distribution  $x_0$  from the Gaussian noise  $x_T \sim \mathcal{N}(0, I)$ . We can express an arbitrary sample  $x_t$  from the initial data  $x_0$ :

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon,$$

where  $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$  and  $\epsilon \sim \mathcal{N}(0, I)$ . The denoising network  $\epsilon_\theta(\cdot)$  is trained to recover  $x_0$  by predicting the noise  $\epsilon_\theta(x_t, t)$ . The corresponding learning objective can be formalized as a simple mean-squared error loss between the true noise and the predicted noise:

$$\mathcal{L} = \mathbb{E}_{x_0, \epsilon, t, c} [\|\epsilon - \epsilon_\theta(x_t, t, c)\|_2^2],$$

where  $t$  is uniformly sampled from  $\{1, \dots, T\}$ ,  $c$  is condition and  $\epsilon \sim \mathcal{N}(0, I)$ .

The denoising network  $\epsilon_\theta(\cdot)$  is typically implemented by U-Net (Ho et al., 2020). To make the diffusion process conditional on the input, condition  $c$  is fed into  $\epsilon_\theta(\cdot)$  via a cross-attention layer implementing  $Attention(Q, K, V) = softmax(\frac{QK^T}{\sqrt{d}}) \cdot V$ , where the intermediate representations of the U-Net acting as the query  $Q$ , and the condition embeddings  $c$  acting as the key  $K$  and value  $V$ .

Classifier-free guidance (Ho and Salimans, 2022) is a widely used technique to improve sample quality while reducing diversity in conditional diffusion models, which jointly trains a single diffusion model on conditional and unconditional objectives via randomly dropping  $c$  during training (e.g. with 10% probability). During sampling, the output of the model is extrapolated further in the direction of  $\epsilon_\theta(x_t|c)$  and away from  $\epsilon_\theta(x_t|\emptyset)$  as follows:

$$\hat{\epsilon}_\theta(x_t|c) = \epsilon_\theta(x_t|\emptyset) + \gamma \cdot (\epsilon_\theta(x_t|c) - \epsilon_\theta(x_t|\emptyset))$$

where  $\gamma \geq 1$  is the guidance scale.

### 3.4 Adaptive Guidance

Previous work (Dhariwal and Nichol, 2021; Nichol et al., 2021; Li et al., 2022b) in text-to-image generation have explored to utilize a classifier or a CLIP (Radford et al., 2021) model to improve a diffusion generator. A CLIP model consists of two separate pieces: an image encoder and a caption encoder. The model optimizes a contrastive cross-entropy loss that encourages a high dot-product if the image is paired with the given caption, or a low dot-product if the image and caption correspond to different parts in the training data. The denoising diffusion process can be perturbed by the gradient of the dot product of the image and caption.

One of the primary challenges of visual story generation is to maintain consistent background and character appearances throughout the story. In order to achieve this goal, during the diffusion sample stage, we propose an adaptive guidance, which explicitly requires that the image generated later should be consistent with the preceding generated images. Considering that images whose corresponding sentences are similar should have similar scenes/characters. When generating the image  $x_i$  in the story, we first use clip text encoder to calculate the similarity score for each historical text in  $\{s_1, \dots, s_{i-1}\}$  with the current text  $s_i$ . After that, we select the text-image pair with the highest similarity score. When the similarity score exceeds the threshold, we believe that the selected image and the image to be generated currently have high similarity, and we can use this image to guide the sampling process of the diffusion model. When the similarity score is lower than the threshold, we think that images in history are not similar to the image to be generated at present, and do not add sampling guidance.

The previous CLIP guided model (Nichol et al., 2021) needs to train an additional noisy CLIP model. It’s time and computation costly, and difficult to classify noised image. Following UPainting (Li et al., 2022b), we use normal CLIP for guidance, and modify the CLIP inputs as follows:

$$\hat{x}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}}(x_t - \sqrt{1 - \bar{\alpha}_t}\epsilon_t)$$

$$x_{in} = \sqrt{1 - \bar{\alpha}_t}\hat{x}_0 + (1 - \sqrt{1 - \bar{\alpha}_t})x_t$$

The denoising diffusion process can be formulated as follows:

$$\hat{\epsilon}'_\theta(x_t|c) = \epsilon_\theta(x_t|\emptyset) + \gamma \cdot (\epsilon_\theta(x_t|c) - \epsilon_\theta(x_t|\emptyset))$$

$$-g\sqrt{1 - \bar{\alpha}_t}\nabla_{x_t}(f(x_{in}) \cdot f(x_h))$$

where  $\gamma \geq 1$  is the classifier weight,  $g \geq 0$  is adaptive guidance weight,  $f(\cdot)$  is the CLIP image encoder and  $x_h$  is the most similar image to the current image  $x_t$ .

## 4 Experiments

### 4.1 Datasets and Metrics

We carried out experiments on both story visualization and story continuation tasks. Given a sequence of sentences forming a narrative, story visualization is the task of generating a corresponding sequence of images. Story continuation, including an initial ground truth image as input, is a variant of story visualization. We use two popular datasets, PororoSV (Li et al., 2019) and FlintstonesSV (Gupta et al., 2018), to evaluate our model. We give the statistics of two datasets in Table 1 and show main characters in Figure 4 and Figure 5 to help the reader understand our examples.

	Train	Valid	Test
PororoSV	10,191	2,334	2,208
FlintstonesSV	20,132	2,071	2,309

Table 1: Statistics for PororoSV and FlintstonesSV datasets.

We adopt the automatic evaluation metrics following existing works and report results using the evaluation script provided in prior work<sup>2</sup>.

**Frechet Inception Distance (FID)** captures the level of similarity between two groups based on statistical analysis of visual features in their respective

<sup>2</sup><https://github.com/adymaharana/VLCStoryGan>



Figure 4: Main characters in PororoSV dataset.




Figure 5: Main characters in FlintstonesSV dataset.

raw images, using the inception v3 model. Lower FID scores indicate higher resemblance between the predicted images and the ground-truth images.

**Character F1 score** calculates the proportion of characters present in the generated images that exactly match the characters in the story inputs. To achieve this, a pretrained inception v3 model (Szegedy et al., 2015) is fine-tuned on each dataset using a multi-label classification loss, enabling it to make predictions of characters in test images.

**Frame Accuracy** evaluates whether all characters from a story are correctly represented in the corresponding images, utilizing the same model employed in the Character F1 score. While the Character F1 score measures the proportion of characters captured in a story, Frame Accuracy quantifies the percentage of samples where all characters are appropriately included.

## 4.2 Implementation Details

Our model is fine-tuned from the pre-trained Stable Diffusion text-to-image generation model. We use CLIP base model and BLIP base model. We only train the parameters of diffusion model and cross attention module, and freeze the parameters of variational auto-encoder, CLIP and BLIP, which could speed up training and save GPU memory. Follow previous work, we train our model for 50 epochs. We use Adam optimizer and set learning rate to  $1e-4$ . For  $\gamma$ , we used the default value in stable diffusion without adjustment. For the threshold of similarity score, we randomly sampled 50 stories

to calculate the similarity score, then manually observed the relationship between the similarity score and the image similarity, and finally set the threshold to 0.65. For the  $g$ , we chose the best value 0.15 from 0.1, 0.15, 0.2, 0.5.

## 4.3 Baselines

**StoryGAN** (Li et al., 2019) uses the standard GAN technique, which includes a recurrent text encoder, an image generation module, and two discriminators - image and story discriminator.

**StoryGANc** (Maharana et al., 2022) follows the general framework of the StoryGAN model and adds the source image as input for the story continuation task.

**CP-CSV** (Song et al., 2020) tries to better preserve character information with three modules: story and context encoder, figure-ground segmentation, and figure-ground aware generation.

**DUCO-StoryGAN** (Maharana et al., 2021b) utilizes a video captioning model to generate an additional learning signal forcing the alignment of image and text, and a memory-augmented transformer to model complex interactions between frames.

**VLC-StoryGAN** (Maharana and Bansal, 2021b) incorporates constituency parse trees, common-sense information and visual information, including bounding boxes and dense captioning, to enhance the visual quality and image consistency.

**Word-Level** (Li and Lukasiewicz, 2022) incorporates word information and extends word-level spatial attention to focus on all words and visual spatial locations in the entire story.

**StoryDALL-E** (Maharana et al., 2022) modifies the pre-trained text-to-image model DALL-E by adding a cross attention module for story continuation.

**AR-LDM** (Pan et al., 2022) employs a history-aware encoding module to incorporate the current text prompt and previously generated text-image history to diffusion model for visual story generation.

## 5 Results

### 5.1 Story Visualization

We evaluate our model on PororoSV dataset for story visualization task. Results are shown in Table 3. We can observe that diffusion-based model

Model	PororoSV			FlintstonesSV		
	FID ↓	Char-F1↑	F-Acc↑	FID ↓	Char-F1↑	F-Acc↑
StoryGANc(BERT)	72.98	43.22	17.09	91.37	70.45	55.78
StoryGANc (CLIP)	74.63	39.68	16.57	90.29	72.80	58.39
StoryDALL-E(prompt)	61.23	29.68	11.65	53.71	42.48	32.54
StoryDALL-E (finetuning)	25.90	36.97	17.26	26.49	73.43	55.19
MEGA-StoryDALL-E	23.48	39.91	18.01	23.58	74.26	54.68
AR-LDM	17.40	-	-	19.28	-	-
<b>ACM-VSG (Ours)</b>	<b>15.36</b>	<b>45.71</b>	<b>22.62</b>	<b>18.41</b>	<b>94.95</b>	<b>88.89</b>

Table 2: Results on the test sets of PororoSV and FlintstonesSV datasets from various models. Scores are based on FID , character classification F1, and frame accuracy evaluations.

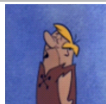

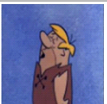



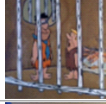
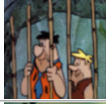

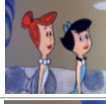
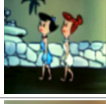
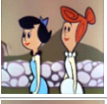



Id	Text	Gold	Ours	AR-LDM
1	Barney is standing in a room. He speaks and looks tired.			
2	Fred and Barney are in jail. Fred is explaining something to Barney while the two of them are standing in a cell behind bars.			
3	Fred and Barney are in jail. Fred opens his arms and speaks. Then Barney responds.			
4	Wilma and Betty are walking through a yard together.			
5	Wilma and Betty are happily walking next to each other outside. Betty is talking while Wilma is listening.			

Figure 6: Example of generated images from previous model AR-LDM and our proposed model.

Model	FID ↓
StoryGAN	158.06
CP-CSV	149.29
DUCO-StoryGAN	96.51
VLC-StoryGAN	84.96
VP-CSV	65.51
Word-Level SV	56.08
AR-LDM	16.59
<b>ACM-VSG (Ours)</b>	<b>15.48</b>

Table 3: Story visualization FID score results on PororoSV dataset.

outperforms the prior methods by a large margin, and our proposed ACM-VSG achieves the best FID score 15.48, indicating our model is able to generate high-quality images.

## 5.2 Story Continuation

Table 2 shows the results for story continuation task. As we can see, our model can achieve the best results on both datasets, 15.36 and 18.41 FID for PororoSV and FlintstonesSV, respectively. And our model can greatly preserve characters to improve the consistency of the story. In addition, we show an example on FlintstonesSV and pororoSV dataset in Figure 6 and Figure 7. We can observe that our model is able to maintain the text-image alignment and consistency across images.

## 5.3 Ablation Study

Table 4 shows ablation studies to ensure that each component in the our proposed method benefits visual story generation. -Guidance means removing the adaptive guidance. -Attention means removing the cross attention module in the adaptive encoder.

Id	Text	Gold	Ours	AR-LDM
1	Eddy keep holding his picture and explains his idea to Poby. Because of the snowy weather we have no choice but to rescue Pororo and Crong by airship. Therefore Eddy resorts to Poby that they need Poby's help.			
2	Poby seems surprised because Poby doesn't expect that Poby will be needed in this situation. After hearing from Eddy Poby turns his head to the left side.			
3	Pororo and Crong are in the middle of the mountain. They seem tired and exhausted. Pororo close his eyes with long hard thinking.			
4	Pororo closes his eyes with long hard thinking. The weather is snowy and it becomes worse. Pororo can't find any other solutions except rope to get out of this mountain. Pororo and Crong are stuck in this mountain so Pororo tries to use rope.			
5	Pororo and Crong try to pull the rope to overcome this situation. However it is hard to fully apply their force.			

Figure 7: Example of generated images from previous model AR-LDM and our proposed model.

Model	FID ↓	Char-F1↑	F-ACC↑
ACM-VSG	15.36	45.71	22.62
- Guidance	15.96	44.56	22.13
- Attention	16.88	44.27	20.25

Table 4: Ablation study results for story continuation task on PororoSV.

Id	Text	Gold	Ours
1	The green lizard is in the kitchen. He is being held by someone using him as a knife. He speaks to the camera.		
2	Betty is talking on the phone in the bedroom, lying in bed.		
3	Betty is laying in bed in the bedroom.		
4	Wilma is in the room speaking on the phone.		
5	Wilma is in the living room. She is standing by the phone while speaking into the phone receiver.		

Figure 8: The generated image story has global inconsistency.

## 5.4 Error Analysis

Our model significantly improves the performance of visual story generation, but there are still some limits. In order to solve these limitations in fu-

	Text	Gold	Ours
1	Loopy pushes her drink. loopy says Loopy won't drink juice.		
2	Fred and Wilma are standing in a room. Fred speaks while Wilma holds onto his shoulder.		

Figure 9: The bad cases for repetitive character and character action error.

ture work, we analyze the generated images. We randomly sample hundreds of stories from PororoSV and FlintstonesSV datasets and summarize the errors.

**Story Inconsistency.** As shown in Figure 8, there are inconsistencies across generated images. The style of the bed and the color of the quilt between the second and third generated images are inconsistent. The environment around *Wilma* is inconsistent between the fourth and fifth images.

**Repetitive Character.** As shown in Figure 9 case 1, the model may generate the same character repeatedly if it appears multiple times in the text.

**Character Action Error.** As shown in Figure 9 case 2, the subtle actions of characters in the image are misaligned with the text. In case 2, Wilma holds



onto Fred’s shoulder in text, while Fred holds onto Wilma’s shoulder in the generated image.

## 6 Conclusion

In this paper, we explore an effective adaptive context modeling method to improve visual story generation. First, we use an adaptive encoder to select the context closely related to the current image from the historical text-image pairs using the current text. Then we fed the context vectors and the text vector to the diffusion model, and use an adaptive guidance to guide the generation of the current image. Experimental results verify that adaptive context modeling could help generate higher quality images and more consistent stories. In addition, we analyze the generated images and find potential research directions in the future: (1) focus on the global consistency of the story, (2) pay attention to the action and expression of the character, (3) obtain the exact semantics of long stories.

## Limitations

A limitation of this work is that it is only evaluated on synthesized datasets of cartoons with limited characters and scenes. In the real world application, there might be many different scenes/characters, posing new challenges to the proposed approach. Another limitation is the requirement of supervised training data and resources. Despite the number of trainable parameters of our approach (850M) is less than AR-LDM ( $\sim 1.5B$ ), the model still needs many story-level training data and computing resources.

## Acknowledgements

Xiaocheng Feng is the corresponding author of this work. We thank the anonymous reviewers for their insightful comments. Zhangyin Feng, Xinmiao Yu, Xiaocheng Feng and Bing Qin are supported by the National Key R&D Program of China via grant 2020AAA0106502, National Natural Science Foundation of China (NSFC) via grant 62276078, the Key R&D Program of Heilongjiang via grant 2022ZX01A32 and the International Cooperation Project of PCL, PCL2022D01.

## References

Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. 2022. ediffi: Text-to-image diffusion models with

an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*.

Fan Bao, Chongxuan Li, Yue Cao, and Jun Zhu. 2022. All are worth words: a vit backbone for score-based diffusion models. *arXiv preprint arXiv:2209.12152*.

Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. 2023. Muse: Text-to-image generation via masked generative transformers. *arXiv preprint arXiv:2301.00704*.

Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794.

Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. 2021. Cogview: Mastering text-to-image generation via transformers. *Advances in Neural Information Processing Systems*, 34:19822–19835.

Zhida Feng, Zhenyu Zhang, Xintong Yu, Yewei Fang, Lanxin Li, Xuyi Chen, Yuxiang Lu, Jiayang Liu, Weichong Yin, Shikun Feng, et al. 2022. Ernie-vilg 2.0: Improving text-to-image diffusion model with knowledge-enhanced mixture-of-denoising-experts. *arXiv preprint arXiv:2210.15257*.

Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. 2022. Make-a-scene: Scene-based text-to-image generation with human priors. *arXiv preprint arXiv:2203.13131*.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.

Tanmay Gupta, Dustin Schwenk, Ali Farhadi, Derek Hoiem, and Aniruddha Kembhavi. 2018. Imagine this! scripts to compositions to videos. *arXiv: Computer Vision and Pattern Recognition*.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851.

Jonathan Ho and Tim Salimans. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.

Bowen Li and Thomas Lukasiewicz. 2022. Word-level fine-grained story visualization.

Chunye Li, Liya Kong, and Zhiping Zhou. 2020. Improved-storygan for sequential images visualization. *Journal of Visual Communication and Image Representation*, 73:102956.

- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022a. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *arXiv preprint arXiv:2201.12086*.
- Wei Li, Xue Xu, Xinyan Xiao, Jiachen Liu, Hu Yang, Guohao Li, Zhanpeng Wang, Zhifan Feng, Qiaoqiao She, Yajuan Lyu, et al. 2022b. Upainting: Unified text-to-image diffusion generation with cross-modal guidance. *arXiv preprint arXiv:2210.16031*.
- Yitong Li, Zhe Gan, Yelong Shen, Jingjing Liu, Yu Cheng, Yuexin Wu, Lawrence Carin, David Carlson, and Jianfeng Gao. 2019. Storygan: A sequential conditional gan for story visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6329–6338.
- Adyasha Maharana and Mohit Bansal. 2021a. Integrating visuospatial, linguistic, and commonsense structure into story visualization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6772–6786, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Adyasha Maharana and Mohit Bansal. 2021b. Integrating visuospatial, linguistic and commonsense structure into story visualization. *arXiv: Computation and Language*.
- Adyasha Maharana, Darryl Hannan, and Mohit Bansal. 2021a. Improving generation and evaluation of visual stories via semantic consistency. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2427–2442, Online. Association for Computational Linguistics.
- Adyasha Maharana, Darryl Hannan, and Mohit Bansal. 2021b. Improving generation and evaluation of visual stories via semantic consistency. *arXiv: Computation and Language*.
- Adyasha Maharana, Darryl Hannan, and Mohit Bansal. 2022. Storydall-e: Adapting pretrained text-to-image transformers for story continuation. In *European Conference on Computer Vision*, pages 70–87. Springer.
- Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*.
- Xichen Pan, Pengda Qin, Yuhong Li, Hui Xue, and Wenhu Chen. 2022. Synthesizing coherent story with auto-regressive latent diffusion models. *arXiv preprint arXiv:2211.10950*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021a. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021b. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR.
- Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. 2016. Generative adversarial text to image synthesis. In *International conference on machine learning*, pages 1060–1069. PMLR.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*.
- Yun Zhu Song, Zhi Rui Tam, Hung Jen Chen, Huihao Han Lu, and Hong-Han Shuai. 2020. Character-preserving coherent story visualization. *Springer International Publishing eBooks*.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2015. Rethinking the inception architecture for computer vision. *arXiv: Computer Vision and Pattern Recognition*.
- Aaron Van Den Oord, Oriol Vinyals, et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. 2022. Scaling autoregressive models for

content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*.

Gangyan Zeng, Zhaohui Li, and Yuan Zhang. 2019. Pororogan: an improved story visualization model on pororo-sv dataset. In *Proceedings of the 2019 3rd International Conference on Computer Science and Artificial Intelligence*, pages 155–159.

Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. 2017. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 5907–5915.

## A PororoSV Cases

### Case 1:

1. Tongtong opens the door. Crong is now on the Pororo's car. They are entering the Tongtong's house. Tongtong tries to find where the magic wand is.
2. The pink magic wand is located behind the chair. Then the magic wand becomes to come out.
3. Tongtong finally finds out the magic wands.
4. As Tongtong finds out the magic wand Tongtong is confident to make Pororo normal. Tongtong says that Pororo will turn back to normal.
5. Pororo jumps high after Tongtong's promise. Tongtong asks Pororo not to mess up.

### Case 2:

1. Pororo and Crong is in Pororo's house. They are standing next to the bed. Pororo is pointing Crong. Crong looks sad.
2. Pororo and Crong is in Pororo's house standing next to the bed. Pororo is pointing drawer and Crong looks at it with a sad face.
3. Poby is in Pororo's house. Poby is approaching a drawer. Above the drawer there is a book which is slightly open.
4. Poby is in Pororo's house searching for something. Poby leans his head down to take a look.
5. Poby is in Pororo's house. Poby is thinking something standing still putting his right hand on his jaw.

### Case 3:

1. Pororo dust the snow off from Poby.
2. Petty is holding the green block.
3. Harry explains situation. Pororo walks toward chair.
4. Pororo sits and joins play.
5. Eddy and Crong are yelling.

### Case 4:

1. Poby looks at Harry and Harry is talking to Poby while sitting on Poby's shoulder.
2. Petty talks and smiles and mops the floor.
3. Petty is smiling and mobbing the floor.
4. Petty smiles and puts the stuff on the table.
5. Poby talks and opens Poby mouth.

### Case 5:

1. Pororo says it was so stinky.
2. feeling embarrassed Poby waved Poby hands.
3. Pororo thinks who farted just before.
4. everyone saw Crong pinching everyone noses.
5. Crong is sitting on the toilet.

### Case 6:

1. Poby is leaving Loopy house.
2. Poby says bye to Loopy.
3. Poby thinks Pororo might have fixed broken chair.
4. Poby smiles. Loopy walks toward the chair.
5. Loopy is satisfied with fixed chair.

### Case 7:

1. Poby is tired so Poby says to Harry that Poby wants to go to bed with sleepy eyes.
2. Harry is surprised. Harry looks at the window
3. there are two cactus on the shelf. outside the window is dark already.
4. light is turned off and Poby and Harry finish ready to sleep. Harry say good night to Poby.
5. light is turned off and Poby and Harry finish ready to sleep. Poby lays down on the bed.

### Case 8:

1. Poby notices that someone is skiing down.
2. Pororo is skiing away. Loopy is chasing.
3. Pororo notices that Poby is waiting for Pororo.
4. Loopy and Poby are lying down. Eddy approaches.
5. Poby and Loopy stands up.

### Case 9:

1. Poby comes out of the house and find out his friends.
2. Poby explains to the friends that Poby must have fallen asleep inside.
3. Eddy is happy to see that Poby is also already at Eddy's house along with other friends.
4. Eddy is inviting his friends to come into his house. everyone follows him.
5. Eddy is leading his friends into his house. everyone is getting inside.

### Case 10:

1. Pororo and friends came running toward Poby. Poby is watching them.
2. Pororo and friends are talking to Poby. Poby has no idea why his friends are acting like this.
3. Harry sat on Poby's head. Harry is saying sorry to Poby. Poby looks surprised. Harry looks sad.
4. Harry is feeling guilty. meanwhile Poby has no idea what Harry is talking about. Harry is sitting on Poby's head.
5. Harry is feeling very sorry to Poby. Harry is talking to Poby on Poby's head.

### Case 11:

1. Poby is brushing pole to Poby's noses for making Poby

itchy. Pororo tries sneezing Poby to take out of the air.

2. Pororo shows that small helicopter is still working properly.

3. Pororo continually suggests Poby trying to sneeze again.

4. Poby tries to sneeze to get out of air from his body. However sneezing with Poby's free will is really difficult.

5. Poby tries to sneeze to get out of air from his body. However sneezing with Poby's free will is really difficult. Poby gives up sneezing and says Poby can't sneeze anymore.

**Case 12:**

1. Poby was keep singing. Suddenly Poby falls down.

2. Poby feels ashamed and wants that nobody saw him falling down.

3. Seeing Poby through the telescope Eddy secretly smiles and talks to himself that Eddy saw Poby falling down.

4. Eddy is interested in seeing things and friends through telescope. Eddy brings telescope and goes to the mountain to observe his friends more.

5. Up on the mountain Eddy chooses a target. It is Pororo. Eddy looks through the telescope.

## B FlintstonesSV Cases

**Case 1:**

1. Fred and Barney stand outside holding blue lunch boxes. Fred talks to Barney.

2. Fred stands in the kitchen having a friendly conversation with someone.

3. Fred is standing in a room. He is speaking while looking over his shoulder and smiling.

4. Fred is trying to kiss Wilma in a room. Wilma is holding a type of plant in her hand.

5. Wilma is in a living room adjusting the leafs on a house plant that is sitting on a table while talking then she stands up straight and turns her head.

**Case 2:**

1. A Lounging creature is lounging around the room and talking.

2. Wilma is ironing a shirt in the laundry room.

3. Wilma is in the living room. Wilma is ironing. Wilma is bobbing her head.

4. Wilma is in a room. She talks to someone.

5. Wilma is in a room. She is talking.

**Case 3:**

1. Fred sits in the living room and speaks to Barney, who waits to respond.

2. Barney stands and talks with someone in the living room.

3. Fred and Barney are in the quarry. Fred is speaking to Barney.

4. Fred and Barney look worried. Fred and Barney are behind the wall in the yard. Barney and Fred are talking to each other.

5. Barney is talking to Fred outside behind a stone fence. Fred begins to slump down and look sad.

**Case 4:**

1. Fred is riding in the car thinking and talking to himself.

2. Fred touches his chin then crosses his arm while outside.

3. Fred is walking outside while speaking out loud.

4. Fred and Barney are riding in the car with golf clubs strapped to the bumper.

5. Fred is standing in the room, talking to someone off screen left.

**Case 5:**

1. Fred and Barney are standing outside next to the stone wall. Barney is wearing an outfit that makes him look like a boy scout. Fred says something to Barney and then points at him.

2. Fred and Barney are standing on a sidewalk. Barney is speaking to Fred, while Fred listens silently with his hands on his hips.

3. Barney is outside talking.

4. Fred and Barney stand in the yard. They speak to each other.

5. Fred and Barney are standing outside talking. They are in front of the wall and Barney has a hate on.

**Case 6:**

1. Fred sits in the living room and speaks to Barney, who waits to respond.

2. Barney stands and talks with someone in the living room.

3. Fred and Barney are in the quarry. Fred is speaking to Barney.

4. Fred and Barney look worried. Fred and Barney are behind the wall in the yard. Barney and Fred are talking to each other.

5. Barney is talking to Fred outside behind a stone fence. Fred begins to slump down and look sad.

**Case 7:**

1. Wilma is in the room, she is talking.

2. Wilma is in the dining room talking to someone then she starts to laugh.

3. Barney slides towards doorway, and opens door.

4. Wilma is sitting in the dining room at the table while talking on the phone.

5. Wilma is in the dining room. She sits at the table on the phone. Fred is wheeled into the room laying in a bed. As Fred enters, Wilma lowers the phone and looks at him with concern and surprise.

**Case 8:**

1. Fred is in a living room kneeling next to a blue chair.

2. Pebbles and Fred are standing in a room talking. Pebbles turns her head and Fred shrugs his shoulders.

3. Fred and Wilma talk in the bedroom, Fred laughs in response to what Wilma says.

4. Fred makes an angry comment while sitting in the room.

5. Fred stands in the kitchen with an ice block on his head. Then someone reaches up and pats the ice. Fred makes a face and the ice starts melting.

**Case 9:**

1. Barney is talking in the living room.
2. Betty is standing in a room, hanging up balloons. She is talking to someone, as she hangs up a green balloon.
3. Betty is in a room. She stands on a stool and holds a balloon in one hand while talking to someone on the ground. The room is decorated for a party.
4. Wilma is in the living room. Wilma is talking.
5. There is a bird that has its head lowered in the room.

**Case 10:**

1. Wilma and Betty are in the room. They are talking to one another while standing.
2. Betty and Wilma are standing in a room. Wilma has bones in her hair. Betty is talking to Wilma.
3. Wilma is wearing a bone curler in her hair while Betty talks to her in a room.
4. Betty and Wilma are talking in a room.
5. Wilma and Betty are standing in a room by the window. They keep looking out the window while Wilma holds the curtain.

**Case 11:**

1. Mr. Slate is driving his car and laughing.
2. A police officer in a police station sits at a desk and talks into a speaker while looking at a stack of papers. He turns the speaker away from his mouth.
3. A Small Policeman behind wheel and a Policeman with Brown Mustache sit in their police car blinking.
4. The officer that is driving the car is speaking to the officer with mustache.
5. Fred and Barney talk as they sit in the car.

**Case 12:**

1. Barney is sitting outside in a chair reading out loud.
2. Barney is reading the news papers in the backyard.
3. The scene begins with no one in the picture. Barney emerges from hiding behind a stone wall that is in front of him. He is standing outside in the yard. The house is to his right. He says something and then points to himself with his thumb.
4. Barney is outside. He is sitting on a stone wall and is talking.
5. Fred and Barney are in the yard. Fred is yelling at Barney. Fred is holding Barney with a fist raised.




Figure 10: Example of ground truths (left 5 frames) and corresponding generated visual stories (right 5 frames) on PororoSV. These cases are under story continuation setting.



Figure 11: Example of ground truths (left 5 frames) and corresponding generated visual stories (right 5 frames) on FlintstonesSV. These cases are under story continuation setting.

## ACL 2023 Responsible NLP Checklist

---

### A For every submission:

- A1. Did you describe the limitations of your work?  
*limitation*
- A2. Did you discuss any potential risks of your work?  
*4*
- A3. Do the abstract and introduction summarize the paper's main claims?  
*1*
- A4. Have you used AI writing assistants when working on this paper?  
*Left blank.*

### B Did you use or create scientific artifacts?

*Left blank.*

- B1. Did you cite the creators of artifacts you used?  
*No response.*
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?  
*No response.*
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?  
*No response.*
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?  
*No response.*
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?  
*No response.*
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.  
*No response.*

### C Did you run computational experiments?

*4*

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

*4*

*The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.*



- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

4

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

4

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

4

**D  Did you use human annotators (e.g., crowdworkers) or research with human participants?**

*Left blank.*

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

*No response.*

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

*No response.*

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

*No response.*

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

*No response.*

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

*No response.*