

Every picture tells a story: Image-grounded controllable stylistic story generation

Holy Lovenia*, Bryan Wilie*, Romain Barraud*,
Samuel Cahyawijaya, Willy Chung, Pascale Fung
Center for Artificial Intelligence Research (CAiRE)
The Hong Kong University of Science and Technology
(hlovenia, bwilie, rmbarraud)@connect.ust.hk

Abstract

Generating a short story out of an image is arduous. Unlike image captioning, story generation from an image poses multiple challenges: preserving the story coherence, appropriately assessing the quality of the story, steering the generated story into a certain style, and addressing the scarcity of image-story pair reference datasets limiting supervision during training. In this work, we introduce Plug-and-Play Story Teller (PPST) and improve image-to-story generation by: 1) alleviating the data scarcity problem by incorporating large pre-trained models, namely CLIP and GPT-2, to facilitate a fluent image-to-text generation with minimal supervision, and 2) enabling a more style-relevant generation by incorporating stylistic adapters to control the story generation. We conduct image-to-story generation experiments with non-styled, romance-styled, and action-styled PPST approaches and compare our generated stories with those of previous work over three aspects, i.e., story coherence, image-story relevance, and style fitness, using both automatic and human evaluation. The results show that PPST improves story coherence and has better image-story relevance, but has yet to be adequately stylistic.

1 Introduction

Enabling machine-generated stories based on visual cues opens up promising directions, and leads language models (LMs) to be viewed as an interface, allowing its involvement in artistic tasks such as advertisement creation and AI-generated movie scripting (McIntyre and Lapata, 2009; Ji et al., 2022; Xu et al., 2019; Hao et al., 2022).

In that direction, vision-language understanding and generation works succeed in leveraging image as well as text as cross-modal knowledge to solve various tasks (Kafle et al., 2019; Zhou et al., 2020; Yu et al., 2021). One fundamental task, image captioning, which involves the model to generate an

*The authors contributed equally to this work.

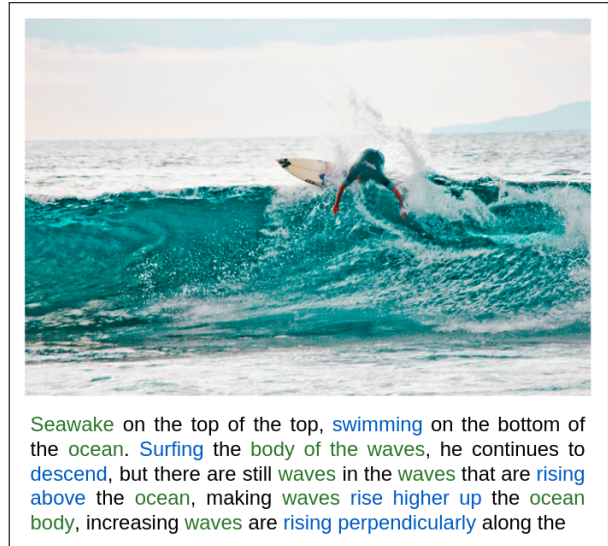


Figure 1: A story generated by action-styled PPST. **Green** denotes the words associated with the image. **Blue** denotes the words associated with the style.

informative textual caption according to a given image, opens up a venue for creativity to be explored. Humans can compose concise descriptions of pictures by focusing on what they find important. Mao et al. (2015); Xia et al. (2021); Mokady et al. (2021); Radford et al. (2021) lay a solid foundation on the current capability of machine learning models to relay cross-modal knowledge for the language models to do generation out of images. Beyond generating captions, creating stories—which utilize linguistics to compose and narrate an inter-related series of events (Li et al., 2018; Peng et al., 2018; Chandu et al., 2019)—according to a single input image offers even possibilities for creativity-based tasks (Wang et al., 2020b; Yang et al., 2019; Hsu et al., 2018).

From the recent advancement on the image-to-story task, it is evident that multiple challenges still remain to be properly solved. One of the main challenges is that model-generated stories tend to lose their coherence as their length increases. Further-

more, the generated text needs to go beyond the pure description of an image as captioning does. Data scarcity, in this context the lack of ready-to-use datasets of image associated with a short story, is also a challenge. Lastly, to the best of our knowledge, there is still limited control over generated stories aside from their relevance to the corresponding image, especially with regards to style (Alabdulkarim et al., 2021). Style has a role to convey a message or story through certain variations of diction and ways of delivery appropriate for a specific context (Ficler and Goldberg, 2017; Shen et al., 2017; Rishes et al., 2013).

In this work, we introduce Plug-and-Play Story Teller (PPST). We take a step towards generating a stylistic story from an image while alleviating the data scarcity issue by leveraging large pre-trained models such as CLIP (Radford et al., 2021) and GPT-2 (Radford et al., 2019), and to add the possibility to control the rendered style through plug-and-play adapters, explored in (Madotto et al., 2020) and (Radford et al., 2021). PPST yields improved natural and on-topic stories, and the resulting stylistic stories also have a strong image-story relevance. Our results highlight the performance of PPST, especially in story coherence and image-story relevance, improving the previous state-of-the-art performance. Lastly, we provide an analysis on the generated stories, including the occurring issues such as repetition and lack of common sense. We present an example of our generated stories using PPST in Figure 1.

2 Related work

2.1 Vision-language generation

In vision-language generation, we exploit both image and text as cross-modal knowledge to address various tasks. Taking on the fact that humans can prepare concise descriptions of pictures by focusing on what they find important, Mao et al. (2015) explore this direction by developing a multimodal recurrent neural network model (RNN) to generate novel image captions. Xia et al. (2021) build a method of cross-modal generative pre-training for text-to-image caption generators through multiple generation tasks. Huang et al. (2019) build an Attention on Attention (AoA) module, which extends conventional attention mechanisms to determine the relevance between attention results and queries. In the encoder, AoA helps to rectify model relationships among different objects in the image; in the

decoder, AoA filters out irrelevant attention results and keeps only the useful ones.

Further, Pan et al. (2020) introduce a unified X-Linear attention block, that fully employs bilinear pooling to selectively capitalize on visual information or perform multimodal reasoning to leverage high order intra- and inter-modal interactions. Cornia et al. (2020) build a meshed transformer with memory architecture that improves both the image encoding and the language generation steps. It explores a multi-level representation of the relationships between image regions integrating learned a priori knowledge, and uses a mesh-like connectivity at decoding stage to exploit low- and high-level features. Mokady et al. (2021) show the effectiveness of the encoding from a recent advancement on vision-language pre-training approach, CLIP (Radford et al., 2021) encoding as a prefix to the caption for image captioning.

2.2 Modeling on low-resource data

Modeling on low-resource data tends to lead to overfitting, which results in non-robust and overly-specific models. This problem is often solved by using augmentation methods. Different augmentation methods and toolkits for various data formats have been developed to better regularize models and increase robustness (Perez and Wang, 2017; Park et al., 2019; Dhole et al., 2021; Lovenia et al., 2022).

With the rise of large pre-trained models, astonishing progress has been made for handling low-resource data. Large pre-trained models, such as BERT (Devlin et al., 2019), GPT2 (Radford et al., 2019), and CLIP (Radford et al., 2021) have shown to be effective for handling multiple low-resource tasks (Wilie et al., 2020; Cahyawijaya et al., 2021; Winata et al., 2022, 2021). The labelled data for image-to-story task is also scarce, hence we extend these large pre-trained models to allow a more robust image-to-story generation.

2.3 Image-grounded story generation

In the image-grounded story generation task (Rameshkumar and Bailey, 2020; Wang et al., 2020a, 2018; Concepción et al., 2016; Ferraro et al., 2019; Mitchell et al., 2018; Min et al., 2021), the widely adopted pipeline includes: 1) extracting captions from an image, 2) encoding the caption, 3) altering the caption with pre-trained encoded stories, and 4) decoding the resulting story. Skip-thought vectors (Kiros et al., 2015)




Figure 2: The inference pipeline of Plug-and-Play Story Teller (PPST) for controllable story generation based on a single image. Visual input encoder encodes the image into a rich semantic embedding, which is then projected into a fixed length visual prefix to be fed to the stylistic language model in order to generate a styled story.

and a sentence encoder-decoder have been used to build an image-to-story generator or to align books and movies (Zhu et al., 2015). Ba et al. (2016) design alternative pipelines by chaining a convolutional neural network (CNN) to extract feature and a recurrent neural network (RNN) with attention for story generation.

Other previous works have explored the use of a graph-based architecture (Wang et al., 2020b) for visual storytelling by modeling the two-level relationships on scene graphs. Yang et al. (2019) present a commonsense-driven generative model, which aims to introduce commonsense from an external knowledge base for visual storytelling. Hsu et al. (2018) propose an inter-sentence diverse beam search to produce expressive stories. One of the latest works in the field is Image2Story (Min et al., 2021), which will be further explained in §4.3.

2.4 Controllable text generation

One important aspect required in natural language generation is the control over the produced result. Recent approaches on style generation control have shown promising results. Dathathri et al. (2020) develop plug-and-play language models (PPLM), which combine a pre-trained LM with one or more simple attribute classifiers that guide text generation without any further training of the LM. Smith et al. (2020) adapt (Weston et al., 2018; Roller et al., 2021), and compare it with some of the previously mentioned approach on controlling the styles of generative models to match one among about 200 possible styles.

While Smith et al. (2020) mention that PPLM-style approach is cheaper at train time, Madotto et al. (2020) highlight its considerable computa-

tional overhead. Madotto et al. (2020) tackle this issue by developing a plug-and-play conversational model (PPCM) that uses residual adapters (Houlsby et al., 2019) and discards the need of further computation at decoding time and any fine-tuning of a large LM. At the same time, the generation result using PPCM is also more fluent and style-consistent. For this reason, we adapt PPCM to introduce style controllability into our method.

3 Plug-and-Play Story Teller (PPST)

We present the overview of our approach: Plug-and-Play Story Teller (PPST) during inference in Figure 2. To generate stories out of an image, PPST involves two main components: visual input encoder (*Enc*) and plug-and-play stylistic story decoder (*Dec*). We use two datasets: an image captioning dataset $\mathcal{D} = \{(v_i^{\mathcal{D}}, c_i^{\mathcal{D}})\}_{i=1}^n$, where $v^{\mathcal{D}}$ denotes image as the visual content and $c^{\mathcal{D}}$ denotes the caption with the textual description of the respective image, and a book passage collection $\mathcal{B} = \{(p_i^{\mathcal{B}}, g_i^{\mathcal{B}})\}_{i=1}^n$, where $p^{\mathcal{B}}$ denotes the passage chunk and $g^{\mathcal{B}}$ denotes its style (genre).

3.1 Visual input encoder

Initially, PPST needs to be able to grasp what the image depicts on a factual basis (e.g., objects, performed actions, and the implied associations) so it should have prior knowledge to develop the story on. For this purpose, we use CLIP (Radford et al., 2021), which learns and accumulates knowledge of visual concepts through a wide variety of image-sentence pairs. CLIP builds its comprehension of text-image alignment by pre-training an image encoder and a text decoder together, and employs a contrastive learning objective to maximize the cosine similarity for the correct image-sentence

pairings. Leveraging the text-image alignment capability provided by CLIP, we utilize its image encoder as the visual input encoder (Enc) to produce rich semantic embeddings $\mathcal{R}^\mathcal{E} = \{r_i^\mathcal{E}\}_{i=1}^n$ from the images $\{v_i^D\}_{i=1}^n$.

Mapping network Although Enc and Dec have been pre-trained using natural language supervision, both of them undergo the learning process separately, which leads to develop latent spaces that provide crucial knowledge but are independent from each other. Furthermore, Dec has yet to be familiar with the visual content offered by the representations generated by Enc ($\mathcal{R}^\mathcal{E}$). To align Dec with the latent space where $\mathcal{R}^\mathcal{E}$ is in, the straightforward way is to simply fine-tune Dec on $\mathcal{R}^\mathcal{E}$.

However, this method expands the number of parameters that Dec has and adds a notable amount of computation cost to the training process. Due to this reason, following (Mokady et al., 2021; Li and Liang, 2021), we introduce a mapping network Map to act as a bridge between the latent spaces of Enc and Dec . Using $\mathcal{R}^\mathcal{E}$ as its input, we train Map to produce a fixed length visual prefix $\mathcal{P}^\mathcal{E}$ adjusted to the latent space of Dec , so Dec can receive and understand visual information from the prefix $\mathcal{P}^\mathcal{E}$, making fine-tuning on $\mathcal{R}^\mathcal{E}$ more of an option rather than a necessity. The usage of Map in our pipeline is further explained in §3.2.

3.2 Plug-and-play stylistic story decoder

Borrowing the natural language ability that large pre-trained models possess, we utilize a pre-trained language model LM as a foundation for generating text in our story decoder Dec . Utilizing a pre-trained language model lets the generation leverage a large amount of unlabelled texts with a causal language modeling objective.

To equip our story decoder Dec with stylistic capabilities, we follow PPCM (Madotto et al., 2020) approach, by inserting residual adapters (Houlsby et al., 2019; Bapna and Firat, 2019) on top of each transformer layer of LM . The adapters act as style adapters $StyAdp = \{S_j\}_{j=1}^m$ which are responsible for guiding LM 's text generation according to the style in use. Each adapter block S_j consists of a layer normalization (Ba et al., 2016) for efficient adaptation, followed by an auto-encoder (Hinton and Zemel, 1993) with a residual connection.

For each style from $j = 1$ to m , we first select a subset of \mathcal{B} where g_i^B equals the j -th style, then




Figure 3: Architecture of our plug-and-play stylistic language model SLM using GPT-2 language model LM and style adapters $StyAdp$.

train S_j using frozen LM parameters and trainable S_j parameters on the passages in the subset p^{B_j} . After training, the $StyAdp$ are then utilized to steer the output of the LM distribution at inference time without modifying the original weights. We refer to the LM with the trained $StyAdp$ as plug-and-play stylistic language model (SLM). The architecture of SLM is shown in Figure 3.

Without any modification, an LM is conditioned on a textual input to prompt text generation. To enable Dec to produce texts based on visual representations, we employ Map to translate $\mathcal{R}^\mathcal{E}$ to the input embedding space of SLM . During a forward pass, Map projects $\mathcal{R}^\mathcal{E}$ into fixed length visual prefixes $\mathcal{P}^\mathcal{E} = \{p_i^\mathcal{E}\}_{i=1}^n$ which is then fed to the Dec to perform text generation based on $\mathcal{P}^\mathcal{E}$. By using this pipeline, we train Map using D to allow Map to project meaningful semantic from $\mathcal{R}^\mathcal{E}$ into the input embedding space of LM . By combining Map and $StyAdp$, we enable SLM to ground its text generation based on a visual content under a weak supervision introduced by \mathcal{D} .

4 Experiment

4.1 Dataset

As described in §3, we utilize two types of datasets. The first dataset is related to images and captions. We use MS-COCO (Lin et al., 2014) as our image captioning dataset \mathcal{D} . MS-COCO is a large-scale 328K-image dataset commonly used for object de-

tection, segmentation, and captioning. We use the image-caption pairs to obtain prefixes and text embeddings to train the mapping network (see §3). Due to our computing resource limitation, we utilize only 10% of MS-COCO total data.

The second dataset is related to books and genres. For the passage collection \mathcal{B} , we use BookCorpus (Zhu et al., 2015) to enable the adaptation of generated stories to a prompted genre. BookCorpus is a large dataset composed of 11,038 books adding up to nearly 985 millions words (1.3 millions unique words) used to train large models such as BERT (Devlin et al., 2019). We obtain the styles of the books by matching the book titles in BookCorpus with the genres in 2021 Smashwords (Bandy and Vincent, 2021) dataset. Smashwords is a dataset listing the e-books available on the Smashwords platform and recording their title, language, price, publication date, URL, and genre.

As a result, we classify the books in 16 genres: romance, fantasy, science fiction, new adult, young adult, thriller, mystery, vampires, horror, teen, adventure, literature, humor, historical, themes, and other. Finally, we split the book texts based on paragraphs, select the text chunks that consist of 30-60 words as passages, and discard the rest. The total number of passages in our dataset nears 7.7M.

4.2 Experiment setup

We use a pre-trained CLIP with Vision Transformer encoder to obtain text-image alignment representation as *Enc*. We note that different from the settings used in (Madotto et al., 2020), where they use open-domain generic dialogues to serve as a prefix to trigger the responses, here we use a visual prefix to trigger the generation in our experiments. Due to the difference in use case, and to enable tendency towards longer generation responses, we use a GPT-2 model instead of the proposed utilisation of DialoGPT (Zhang et al., 2020b) in (Madotto et al., 2020). In detail, as for the *LM* in §3.2, we utilize a pre-trained GPT-2 with 124M parameters, and employ the same model architecture and size as well for the adaptation of (Madotto et al., 2020).

We conduct the experiment using PPST with a non-stylistic setting (without style adapter), referred to as **Non-styled**, and with two stylistic settings, which are **Romance** and **Action**, since they are the styles represented by most amount of samples in the BookCorpus dataset. To filter out the samples that is strongly categorized as **Romance**

and **Action**, we use the first three genres listed by BookCorpus entries to recognize those entries as **Romance** and **Action** entries.

For **Non-styled**, we utilize the same approach described in §3, but instead of using an LM guided by a style adapter, we use a regular pre-trained LM (no style adapter) *LM* directly fine-tuned on the book collection. We employ **Non-styled** as a comparison against the stylistic approaches in terms of a controllable story generation. For **Romance** and **Action**, fine-tuning of the GPT-2 with style adapters on the book collection data is done for a maximum of 10 epochs, with a learning rate of $1e-3$, a batch size of 8, and a maximum sequence length of 512. During the training on image-sentence pairs, we only train the mapping network with a prefix size of 512, a prefix length of 10, and an activation function of \tanh , and freeze the LM.

Our story generation employs beam search with a beam size of 5, a temperature of 0.8, and a top-k of 10. To avoid repetition, we apply a repetition penalty of 0.7 and limit any repetition of 3-gram phrases. To encourage the model to produce a longer story, we apply an exponentially decaying length penalty with a factor of 1.7 after 20 tokens and set a minimum generation length to be 750.

4.3 Baseline

We use **Image2Story** (Min et al., 2021) as our baseline. It combines an RNN and encoder-decoder structure to generate a short story out of an image. The model is built upon skip-thought encoders and structured in a 3-stage pipeline where: 1) a caption based on an input image and a skip-thought vector based on an image-caption dataset are created, 2) a skip-thought vector based on a story dataset is created, and 3) starting from the caption, the vector in 1) is subtracted and the vector in 2) is added so as to obtain a story fitted to the story dataset based on the input image.

4.4 Evaluation setup

PPST relies on visual semantics and information, so we need to ensure that they manage to extract sufficient knowledge from the input image. For this purpose, we use the original captions provided from the MS-COCO dataset as gold references representing the visual content conveyed by the input images for the text-to-text similarity metrics, and the images for the image-to-text similarity metric.

Model	ROUGE-L	ChrF++	MoverScore	BERTScore	BLEURT	BARTScore	CLIPScore
Image2Story	9.06	15.04	50.20	39.65	23.80	-4.00	59.95
PPST Non-styled	9.52	16.81	50.11	39.71	26.51	-4.05	61.99
PPST Romance	10.02	15.30	51.94	46.48	36.70	-3.86	69.02
PPST Action	10.09	15.28	51.94	46.59	36.69	-3.86	69.21

Table 1: Automatic evaluation results on the visual information retention in the generated stories. For image-to-text similarity, i.e., CLIPScore, we compare the generated stories directly with the corresponding images, while for text-to-text similarity metrics we use the original captions provided from the MS-COCO dataset.

Automatic evaluation We compute seven automatic evaluation metrics covering two n-gram-based text-to-text similarity metrics, i.e., ROUGE-L (Lin, 2004) and ChrF++ (Popović, 2017); four model-based text-to-text similarity metrics, i.e., MoverScore (Zhao et al., 2019), BERTScore (Zhang et al., 2020a), BLEURT (Selam et al., 2020), and BARTScore (Yuan et al., 2021); and one image-to-text similarity metric, i.e. CLIPScore (Hessel et al., 2021). For the image-to-text similarity, we compare the text directly with the original image used for generating the story.

Human evaluation To further assess the quality of the generated stories from our system, we conduct a human evaluation in addition to computing the metrics previously mentioned. Each participant is given a questionnaire composed of 10 subsections. Each subsection has 1 image, randomly sampled from our dataset, followed by four stories respectively generated by 1) **Image2Story**, 2) our **Non-Styled** model, 3) our **Romance** model, and 4) our **Action** model. For all models, we ask if "the story makes sense" to assess story coherence, and if "there is a link between the image and the story" to assess image-story relevance. In addition, for our **Romance** and **Action** models, we ask a third question to know if "the story has the given style" to judge style fitness. The participants answer to the questions using a 5-point Likert scale with the choices: "A lot", "A little", "Neutral", "Not really", and "Not at all". The human evaluation is conducted on 13 participants.

5 Result and analysis

5.1 Image-to-story generation quality

As explained in §4.4, we utilize both automatic and human evaluation to measure the quality of the generated story of four models: 1) Min et al. (2021)’s **Image2Story**, 2) our **Non-styled**, 3) our **Romance**, and 4) our **Action**. Table 1 shows all

the automatic evaluation metrics of the generated story. In general, all of our models outperform the baseline **Image2Story** in both n-gram-based text-to-text similarity, model-based text-to-text semantic similarity, and image-to-text semantic similarity metrics. More specifically, the **Romance** and **Action** models perform significantly better on semantic text-to-text and image-to-text similarity metrics by $\sim 7\%$ on the BERTScore, $\sim 10\%$ on the BLEURT, and $\sim 8\%$ on the CLIPScore. The **Non-styled** model performs not as good as the **Romance** and **Action** models but still yields a slightly better score compared to the **Image2Story** model in most metrics. This automatic evaluation result suggests that PPST, with and without the style adapter, can generate a better image-grounded story despite having no direct supervision for the image-to-story generation task itself.

The human evaluation result is shown in Figure 4. In terms of coherence, our evaluation result suggests that stories generated by **Non-styled** surpasses all other models, with an average rating of 3.12, followed by **Romance**, **Image2Story**, and **Action**). This suggests that pre-trained LM is sufficient to generate coherent stories without requiring tuning on the sentence-to-story generation task as incorporated in the prior work (Min et al., 2021), which shows PPST performs well despite the image-story data scarcity issue.

The relevance between the image and the story aligns with the automatic evaluation result. Our models, especially the stylistic **Romance** and **Action**, outperform the baseline **Image2Story** by a large margin, achieving a rating score 3.5 compared to only 2.77, which suggests a better text-image alignment compared to the prior work. For style fitness, we find that our **Action** model achieves an adequate style-story score of 2.78, while the **Romance** model, only obtain a romance style-story score of 1.91. We further explicate this phenomenon in §5.2.




Figure 4: Human evaluations of the generated stories from all models in terms of story coherence (**left**), image-story relevance (**middle**), and style-story fitness (**right**).

5.2 Analysis on the generated stories

Aside from the automatic and human evaluations, we manually inspect the stories to gather insights regarding the behaviors of our models. Table 2 provides 2 examples of our image-to-story generation. Similar to the majority of the book passages used in the training step, our generated stories are inclined to lean towards *describing* the visual aspects of the input image and slowly building the occurring events from there, which notably accounts for PPST’s higher image-story relevance scores, rather than recounting a chain of events or actions in a straightforward manner as the baseline. We also find that while the surrounding contexts of our stories are relevant to the respective images, this relevance deteriorates as the stories grow longer.

Furthermore, we observe that our styled generation result can contain repetitions and tends to use a few words more often than the others. This aligns with the drawbacks of PPCM described by Madotto et al. (2020), which are mainly caused by the restricted use of vocabulary for generating attribute consistent responses. It is also mentioned that this abuse of restricted vocabulary harms fluency, because it cannot always fit within a given context. All these limitations negatively impact the coherence and fluency, therefore the overall quality of the generated stories. Finally, in spite of the proven capability of controlling the generation, the style-story score, on the right plot of Figure 4, shows that there is still potential for improvement. We leave this exploration for future work, specifically on realizing more generation control, in this case by improving the generated stories to be more related to the styles being adapted.

6 Discussions

Our works have moved image-grounded story generation forward by improving the generated story coherence and image-story relevance, and

by adding a layer of style control on top of it. However, as explained in §5, the current progress still leaves room for improvement.

Story coherence Taking inspiration from recent works, a few strategies to refine story coherence can be implemented as the next step, for instance an unsupervised hierarchical story infilling (Ippolito et al., 2019), a semantic dependency skeleton generation to extract key information (Xu et al., 2018) or storyline (Yao et al., 2019), a deeper understanding of causal and temporal relations of events through commonsense knowledge (Mostafazadeh et al., 2016), the utilization of both sentence-level and discourse-level prefix information for decoding (Guan et al., 2020), and making use of a story dataset with rich and fine-grained annotations (Akoury et al., 2020).

Image-story relevance For the relevance, rather than simple embedding concatenation, other ways to incorporate visual information to textual (Liu et al., 2019) and deepen visual comprehension (Fang et al., 2015; Huang et al., 2016) can be further investigated.

Style control We also highlight the interesting directions in advancing the realization of control over stylistic story generation. Our exploration underlines the importance of improving generated stories to relate more to the styles being adapted. Improved and new approaches to control the generated stories with more specific, descriptive, even depicted by a short passage, styles will open up interesting venues on controllable text generations to assist artistic and creative tasks, whether these methods include a pre-trained model (Keskar et al., 2019; Gan et al., 2017; Hu et al., 2017) or not (Hu et al., 2022).

Image	Generated stories
	<p>Image2Story: a a cat on a laptop top . Alice lay down , trying to decipher . She was a cat , too , so he did n't even know what to do . She resisted , yet she found herself wanting to keep her head on top of a computer . With every step she laid on top of a cat that had taken on as long as you could , I snapped my laptop shut . I lay on my computer , a reminder of a cat that had a job .</p> <p>PPST Non-styled: A cute cat reading a computer and a pad on a desk in a office. 'This is the best I've seen in a long while.' "I don't know how to say this, but I'm not going to let you do this to me. You're the one that's going to have to deal with me. I don</p> <p>PPST Romance: a cats pawing a tree is lying on a table. "I don't have a computer or computer monitor. I do have some photographs taken during the day. I am trying to find some time to spend together. I have been working on some photographs that show how long the cat has been sitting in the same room as the computer. Next to</p> <p>PPST Action: there is a cat on a table. and next to the table, a moving hand is used a hand. and the monitor is open. The cat is lying on the table with its legs folded up. The monitor is tilted slightly upward toward the table itself. The table itself is folded up slightly slightly from the table's edge to the floor surface.</p>
	<p>Image2Story: a a man is nothing . I was wrestling a picnic , he shouted . I had known he was going to throw his leg around my ankle , and Nate gave it to me . It s a kid , as a friend of mine . In the beginning , I realized I had been carrying her to the field and carrying him on a leash , making it all feel like a baseball game . I would practically twitch before he was responsible .</p> <p>PPST Non-styled: A human being that is having fun! It was the first time I've seen someone that I really like. I hope I'm doing the right thing. It wasn't long ago I was a stranger that someone I lost my virginity to."Well, I'm not sure, but I'm sure it's not the same. I</p> <p>PPST Romance: A person who is in the Frisbee field with two Frisbees in hand. In contrast to the frisbee riguring to catch a Frisbier. Behind them, however, there appears to be nothing unusual happening. Rather, it appears that there is little fuss that happens this year. However, nothing unusual has transpired this year</p> <p>PPST Action: A person who is in the Frisbee field with two frisbees. In the air they are holding them. Behind them stands a young man holding a frisb. They are holding Frisbees in their legs. Hands are placed over their necks to allow them to sit comfortably. They sit comfortably in their chairs to sit upright. They</p>

Table 2: Samples of image-to-story generation result generated from **PPST Non-styled**, **PPST Romance**, and **PPST Action** against the baseline **Image2Story**.

7 Limitations

We discuss here about the limitations of our work, specifically concerning the chosen heuristic to align style and passages from BookCorpus, the limited amount of data in choice of style for the adapters, and possible biases.

As explained in section 4.1, we choose to split book texts based on paragraphs in order to retain a certain degree of logical fluency throughout each story samples, as a paragraph usually deals with a single theme or idea. While this helps keeping the passages relatively short, one limitation of this

approach is that some passages might not fully reflect the style of narrative that it is classified as. For example, not every paragraph taken out of context from a romance book will exhibit its genre. There could be a sizable amount of passages that focus on world-building and laying the groundwork for the book's main plot to progress.

We decide to focus on romance and action because these styles are the most represented in the dataset used, as well as being more straightforward to capture in terms of style compared to other genres that rely on an underlying plot throughout the book such as historical or adventure. Generaliz-

ing PPST to these styles with a lower amount of resources might require further experiments.

Lastly, previous works have shown that captioning models can exhibit harmful biases, such as gender bias (Hendricks et al., 2018) and racial bias (Zhao et al., 2021). Since we pair those image captions data with written stories from a wide variety of books, those biases can be further amplified. Thus, such generative processes must be used with caution. While tackling unwanted biases in images or captions is a must, the bias exhibited in stories is sometimes justified by the context and the surrounding narrative. Not all stories should be completely neutral, and this balance should be considered carefully in future directions.

8 Conclusion

By leveraging text-image alignment representations to describe the visual content of a given image in words, we can use the resulting semantic embeddings as prior knowledge to generate a short story out of a given picture through a plug-and-play controllable language model approach. It also allows us to tackle the data scarcity issue in this task.

The results show that our Plug-and-Play Story Teller (PPST) generates more consistent and on-topic stories according to the visual information, as well as performing better in relevance and image-story relationship than the previous state-of-the-art. We also found that PPST without style adapters (**Non-styled**) generates more coherent stories, and PPST utilizing style adapters (**Romance** and **Action**) have a similar, if not a slightly better, image-story relationship than the other approaches.

References

- Nader Akoury, Shufan Wang, Josh Whiting, Stephen Hood, Nanyun Peng, and Mohit Iyyer. 2020. Storyum: A dataset and evaluation platform for machine-in-the-loop story generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6470–6484.
- Amal Alabdulkarim, Siyan Li, and Xiangyu Peng. 2021. Automatic story generation: Challenges and attempts. *NAACL HLT 2021*, page 72.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Jack Bandy and Nicholas Vincent. 2021. Addressing "documentation debt" in machine learning research: A retrospective datasheet for bookcorpus. *arXiv preprint arXiv:2105.05241*.
- Ankur Bapna and Orhan Firat. 2019. Simple, scalable adaptation for neural machine translation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1538–1548.
- Samuel Cahyawijaya, Genta Indra Winata, Bryan Wilie, Karissa Vincentio, Xiaohong Li, Adhiguna Kuncoro, Sebastian Ruder, Zhi Yuan Lim, Syafri Bahar, Masayu Khodra, Ayu Purwarianti, and Pascale Fung. 2021. IndoNLG: Benchmark and resources for evaluating Indonesian natural language generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8875–8898, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Khyathi Chandu, Shrimai Prabhumoye, Ruslan Salakhutdinov, and Alan W Black. 2019. "my way of telling a story": Persona based grounded story generation. In *Proceedings of the Second Workshop on Storytelling*, pages 11–21.
- Eugenio Concepción, Pablo Gervás, and Gonzalo Méndez. 2016. Mining knowledge in storytelling systems for narrative generation. In *Proceedings of the INLG 2016 Workshop on Computational Creativity in Natural Language Generation*, pages 41–50.
- Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. 2020. Meshed-memory transformer for image captioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10578–10587.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. Plug and play language models: A simple approach to controlled text generation. In *International Conference on Learning Representations*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kaustubh D. Dhole, Varun Gangal, Sebastian Gehrmann, Aadesh Gupta, Zhenhao Li, Saad Mahamood, Abinaya Mahendiran, Simon Mille, Ashish Srivastava, Samson Tan, Tongshuang Wu, Jascha Sohl-Dickstein, Jinho D. Choi, Eduard H. Hovy, Ondrej Dusek, Sebastian Ruder, Sajant Anand, Naganender Aneja, Rabin Banjade, Lisa Barthe, Hanna Behnke, Ian Berlot-Attwell, Connor Boyle, Caroline Brun, Marco Antonio Sobrevilla Cabezudo, Samuel Cahyawijaya, Emile Chapuis, Wanxiang Che, Mukund Choudhary, Christian Clauss, Pierre Colombo, Filip Cornell, Gautier Dagan, Mayukh

- Das, Tanay Dixit, Thomas Dopierre, Paul-Alexis Dray, Suchitra Dubey, Tatiana Ekeinhor, Marco Di Giovanni, Rishabh Gupta, Rishabh Gupta, Louanes Hamla, Sang Han, Fabrice Harel-Canada, Antoine Honore, Ishan Jindal, Przemyslaw K. Joniak, Denis Kleyko, Venelin Kovatchev, and et al. 2021. [NL-augmenter: A framework for task-sensitive natural language augmentation](#). *CoRR*, abs/2112.02721.
- Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, et al. 2015. From captions to visual concepts and back. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1473–1482.
- Francis Ferraro, Ting-Hao Huang, Stephanie Lukin, and Margaret Mitchell. 2019. Proceedings of the second workshop on storytelling. In *Proceedings of the Second Workshop on Storytelling*.
- Jessica Fidler and Yoav Goldberg. 2017. [Controlling linguistic style aspects in neural language generation](#). In *Proceedings of the Workshop on Stylistic Variation*, pages 94–104, Copenhagen, Denmark. Association for Computational Linguistics.
- Chuang Gan, Zhe Gan, Xiaodong He, Jianfeng Gao, and Li Deng. 2017. Stylenet: Generating attractive visual captions with styles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3137–3146.
- Jian Guan, Fei Huang, Zhihao Zhao, Xiaoyan Zhu, and Minlie Huang. 2020. A knowledge-enhanced pre-training model for commonsense story generation. *Transactions of the Association for Computational Linguistics*, 8:93–108.
- Yaru Hao, Haoyu Song, Li Dong, Shaohan Huang, Zewen Chi, Wenhui Wang, Shuming Ma, and Furu Wei. 2022. Language models are general-purpose interfaces. *arXiv preprint arXiv:2206.06336*.
- Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. 2018. Women also snowboard: Overcoming bias in captioning models. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 771–787.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. CLIPScore: a reference-free evaluation metric for image captioning. In *EMNLP*.
- Geoffrey E Hinton and Richard Zemel. 1993. Autoencoders, minimum description length and helmholtz free energy. *Advances in neural information processing systems*, 6.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR.
- Chao-Chun Hsu, Szu-Min Chen, Ming-Hsun Hsieh, and Lun-Wei Ku. 2018. Using inter-sentence diverse beam search to reduce redundancy in visual storytelling. In *First Workshop on Storytelling, NAACL 2018*.
- Zhiqiang Hu, Roy Ka-Wei Lee, Charu C Aggarwal, and Aston Zhang. 2022. Text style transfer: A review and experimental evaluation. *ACM SIGKDD Explorations Newsletter*, 24(1):14–45.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. 2017. Toward controlled generation of text. In *International conference on machine learning*, pages 1587–1596. PMLR.
- Lun Huang, Wenmin Wang, Jie Chen, and Xiao-Yong Wei. 2019. Attention on attention for image captioning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4634–4643.
- Ting-Hao Huang, Francis Ferraro, Nasrin Mostafazadeh, Ishan Misra, Aishwarya Agrawal, Jacob Devlin, Ross Girshick, Xiaodong He, Pushmeet Kohli, Dhruv Batra, et al. 2016. Visual storytelling. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1233–1239.
- Daphne Ippolito, David Grangier, Chris Callison-Burch, and Douglas Eck. 2019. [Unsupervised hierarchical story infilling](#). In *Proceedings of the First Workshop on Narrative Understanding*, pages 37–43, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ziwei Ji, Yan Xu, I Cheng, Samuel Cahyawijaya, Rita Frieske, Etsuko Ishii, Min Zeng, Andrea Madotto, Pascale Fung, et al. 2022. Vscript: Controllable script generation with audio-visual presentation. *arXiv preprint arXiv:2203.00314*.
- Kushal Kafle, Robik Shrestha, and Christopher Kanan. 2019. Challenges and prospects in vision and language research. *Frontiers in Artificial Intelligence*, 2:28.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Ryan Kiros, Yukun Zhu, Russ R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. *Advances in neural information processing systems*, 28.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597.

- Zhongyang Li, Xiao Ding, and Ting Liu. 2018. [Generating reasonable and diversified story ending using sequence to sequence model with adversarial training](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1033–1043, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- Fenglin Liu, Yuanxin Liu, Xuancheng Ren, Xiaodong He, and Xu Sun. 2019. Aligning visual regions and textual concepts for semantic-grounded image representations. *Advances in Neural Information Processing Systems*, 32.
- Holy Lovenia, Bryan Wilie, Willy Chung, Zeng Min, Samuel Cahyawijaya, Dan Su, and Pascale Fung. 2022. [Clozer: Adaptable data augmentation for cloze-style reading comprehension](#). In *Proceedings of the 7th Workshop on Representation Learning for NLP*, pages 60–66, Dublin, Ireland. Association for Computational Linguistics.
- Andrea Madotto, Etsuko Ishii, Zhaojiang Lin, Sumanth Dathathri, and Pascale Fung. 2020. [Plug-and-play conversational models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2422–2433, Online. Association for Computational Linguistics.
- Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, Zhiheng Huang, and Alan Yuille. 2015. Deep Captioning with Multimodal Recurrent Neural Networks (m-RNN). *ICLR*.
- Neil McIntyre and Mirella Lapata. 2009. Learning to tell tales: A data-driven approach to story generation. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 217–225.
- Kyunbok Min, Minh Dang, and Hyeonjoon Moon. 2021. Deep learning-based short story generation for an image using the encoder-decoder structure. *Digital Object Identifier 10.1109/ACCESS.2021.3104276*.
- Margaret Mitchell, Ting-Hao Huang, Francis Ferraro, and Ishan Misra. 2018. Proceedings of the first workshop on storytelling. In *Proceedings of the First Workshop on Storytelling*.
- Ron Mokady, Amir Hertz, and Amit H Bermano. 2021. Clipcap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*.
- Nasrin Mostafazadeh, Alyson Grealish, Nathanael Chambers, James Allen, and Lucy Vanderwende. 2016. Caters: Causal and temporal relation scheme for semantic annotation of event structures. In *Proceedings of the Fourth Workshop on Events*, pages 51–61.
- Yingwei Pan, Ting Yao, Yehao Li, and Tao Mei. 2020. X-linear attention networks for image captioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10971–10980.
- Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin Dogus Cubuk, and Quoc V. Le. 2019. Specaugment: A simple augmentation method for automatic speech recognition. In *INTER-SPEECH*.
- Nanyun Peng, Marjan Ghazvininejad, Jonathan May, and Kevin Knight. 2018. Towards controllable story generation. In *Proceedings of the First Workshop on Storytelling*, pages 43–49.
- Luis Perez and Jason Wang. 2017. The effectiveness of data augmentation in image classification using deep learning. *arXiv preprint arXiv:1712.04621*.
- Maja Popović. 2017. [chrF++: words helping character n-grams](#). In *Proceedings of the Second Conference on Machine Translation*, pages 612–618, Copenhagen, Denmark. Association for Computational Linguistics.
- 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.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Revanth Rameshkumar and Peter Bailey. 2020. Storytelling with dialogue: A critical role dungeons and dragons dataset. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5121–5134.
- Elena Rishes, Stephanie M Lukin, David K Elson, and Marilyn A Walker. 2013. Generating different story tellings from semantic representations of narrative. In *International Conference on Interactive Digital Storytelling*, pages 192–204. Springer.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. [Recipes for building an open-domain chatbot](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325, Online. Association for Computational Linguistics.

- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. Bleurt: Learning robust metrics for text generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. *Advances in neural information processing systems*, 30.
- Eric Michael Smith, Diana Gonzalez-Rico, Emily Dinan, and Y-Lan Boureau. 2020. Controlling style in generated dialogue. *arXiv preprint arXiv:2009.10855*.
- Ruize Wang, Zhongyu Wei, Ying Cheng, Piji Li, Haijun Shan, Ji Zhang, Qi Zhang, and Xuan-Jing Huang. 2020a. Keep it consistent: Topic-aware storytelling from an image stream via iterative multi-agent communication. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2250–2260.
- Ruize Wang, Zhongyu Wei, Piji Li, Qi Zhang, and Xuanjing Huang. 2020b. Storytelling from an image stream using scene graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9185–9192.
- Xin Wang, Wenhui Chen, Yuan-Fang Wang, and William Yang Wang. 2018. No metrics are perfect: Adversarial reward learning for visual storytelling. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 899–909.
- Jason Weston, Emily Dinan, and Alexander Miller. 2018. Retrieve and refine: Improved sequence generation models for dialogue. In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*, pages 87–92.
- Bryan Wilie, Karissa Vincentio, Genta Indra Winata, Samuel Cahyawijaya, Xiaohong Li, Zhi Yuan Lim, Sidik Soleman, Rahmad Mahendra, Pascale Fung, Syafri Bahar, and Ayu Purwarianti. 2020. [IndoNLU: Benchmark and resources for evaluating Indonesian natural language understanding](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 843–857, Suzhou, China. Association for Computational Linguistics.
- Genta Indra Winata, Alham Fikri Aji, Samuel Cahyawijaya, Rahmad Mahendra, Fajri Koto, Ade Romadhony, Kemal Kurniawan, David Moeljadi, Radityo Eko Prasajo, Pascale Fung, Timothy Baldwin, Jey Han Lau, Rico Sennrich, and Sebastian Ruder. 2022. [Nusax: Multilingual parallel sentiment dataset for 10 Indonesian local languages](#).
- Genta Indra Winata, Samuel Cahyawijaya, Zihan Liu, Zhaojiang Lin, Andrea Madotto, and Pascale Fung. 2021. [Are multilingual models effective in code-switching?](#) In *Proceedings of the Fifth Workshop on Computational Approaches to Linguistic Code-Switching*, pages 142–153, Online. Association for Computational Linguistics.
- Qiaolin Xia, Haoyang Huang, Nan Duan, Dongdong Zhang, Lei Ji, Zhifang Sui, Edward Cui, Taroon Bharti, and Ming Zhou. 2021. Xgpt: Cross-modal generative pre-training for image captioning. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 786–797. Springer.
- Jingjing Xu, Xuancheng Ren, Yi Zhang, Qi Zeng, Xiaoyan Cai, and Xu Sun. 2018. [A skeleton-based model for promoting coherence among sentences in narrative story generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4306–4315, Brussels, Belgium. Association for Computational Linguistics.
- Peng Xu, Chien-Sheng Wu, Andrea Madotto, and Pascale Fung. 2019. Clickbait? sensational headline generation with auto-tuned reinforcement learning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3065–3075.
- Pengcheng Yang, Fuli Luo, Peng Chen, Lei Li, Zhiyi Yin, Xiaodong He, and Xu Sun. 2019. Knowledgeable storyteller: A commonsense-driven generative model for visual storytelling. In *IJCAI*, volume 3, page 7.
- Lili Yao, Nanyun Peng, Ralph Weischedel, Kevin Knight, Dongyan Zhao, and Rui Yan. 2019. Plan-and-write: Towards better automatic storytelling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7378–7385.
- Tiezheng Yu, Wenliang Dai, Zihan Liu, and Pascale Fung. 2021. Vision guided generative pre-trained language models for multimodal abstractive summarization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3995–4007.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [Bartscore: Evaluating generated text as text generation](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 27263–27277. Curran Associates, Inc.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020a. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and William B Dolan. 2020b. DIALOGPT:

- Large-Scale Generative Pre-training for Conversational Response Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278.
- Dora Zhao, Angelina Wang, and Olga Russakovsky. 2021. Understanding and evaluating racial biases in image captioning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14830–14840.
- Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M. Meyer, and Steffen Eger. 2019. Moverscore: Text generation evaluating with contextualized embeddings and earth mover distance. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, Hong Kong, China. Association for Computational Linguistics.
- Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. 2020. Unified vision-language pre-training for image captioning and VQA. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13041–13049.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *The IEEE International Conference on Computer Vision (ICCV)*.