

TARN-VIST: Topic Aware Reinforcement Network for Visual Storytelling

Weiran Chen, Xin Li, Jiaqi Su, Guiqian Zhu, Ying Li, Yi Ji*, Chunping Liu*

School of Computer Science and Technology, Soochow University, Suzhou, China

{wrchen2023, 20224227052}@stu.suda.edu.cn, czu_lixin@163.com

{gqzhu, ingli, jiyi, cpliu}@suda.edu.cn

Abstract

As a cross-modal task, visual storytelling aims to generate a story for an ordered image sequence automatically. Different from the image captioning task, visual storytelling requires not only modeling the relationships between objects in the image but also mining the connections between adjacent images. Recent approaches primarily utilize either end-to-end frameworks or multi-stage frameworks to generate relevant stories, but they usually overlook latent topic information. In this paper, in order to generate a more coherent and relevant story, we propose a novel method, **Topic Aware Reinforcement Network for Visual StoryTelling (TARN-VIST)**. In particular, we pre-extracted the topic information of stories from both visual and linguistic perspectives. Then we apply two topic-consistent reinforcement learning rewards to identify the discrepancy between the generated story and the human-labeled story so as to refine the whole generation process. Extensive experimental results on the VIST dataset and human evaluation demonstrate that our proposed model outperforms most of the competitive models across multiple evaluation metrics.

Keywords: Visual Storytelling, Topic Information, Reinforcement Learning

1. Introduction

Nowadays, visual storytelling has garnered increasing attention from the fields of both computer vision (CV) and natural language processing (NLP) due to its significance and practicality in some applications such as image retrieval, image subtitling, and blind navigation (Fan et al., 2021). As opposed to visual captioning, visual storytelling also involves exploring the corresponding relationships between object pairs in adjacent images. Additionally, when humans tell stories, they usually revolve around a specific central topic. Therefore, to generate high-quality stories, visual storytelling models also require taking the topic information into account.

In the field of visual storytelling, existing methods can be divided into two main categories: end-to-end-based methods (Huang et al., 2016; Yu et al., 2017; Kim et al., 2018; Wang et al., 2018, 2019; Hu et al., 2020; Xu et al., 2021; Braude et al., 2022; Chen et al., 2022; Yang and Jin, 2023; Li et al., 2023) and multi-stage-based methods (Hsu et al., 2020; Chen et al., 2021; Chu et al., 2021). Both them, end-to-end-based methods typically map directly from the input image sequence to the output story, while multi-stage-based methods employ different modules trained independently in distinct stages, and the results of the previous stage are always used as the input for the subsequent one. The common idea behind the end-to-end-based methods is to use a convolutional neural network (CNN) as an encoder to extract high-dimensional image features and overall image-stream features. These

representational feature vectors are then fed into a long short term memory (LSTM) to construct the story. These approaches can always yield outstanding stories with high score in automatic metrics. On the other hand, multi-stage-based methods, known as planning and writing strategies, advocate for separating the generation process into several steps. Generally speaking, the first step is always to employ an object detection module to detect salient concepts in the given images and the next step is to generate a related story through a transformer-based architecture. This kind of methods can generate stories that reflect human preferences.

Despite their remarkable progress, there are still several technical limitations in the visual storytelling task. One of the drawbacks is that few models consider the latent topic information of the generated story. The topic of the story serves as its central idea, which is the core of the story. It is well known that maintaining a coherent and engaging storyline is crucial, and centering the narrative around a specific topic aids in achieving this coherence. Besides, telling a story around a concrete topic can enhance readers' comprehension as well as facilitate better content retention. Nevertheless, once the generated story goes off-topic, it will become incoherent and lack narrative variety. To be noted, reinforcement learning provides a unique technique to guide the model towards generating stories that not only capture the correct relationships between images but also ensure a more consistent and accurate topic information. Hence, we incorporate reinforcement learning into our methodology to enhance the overall quality of visual storytelling.

*Corresponding author

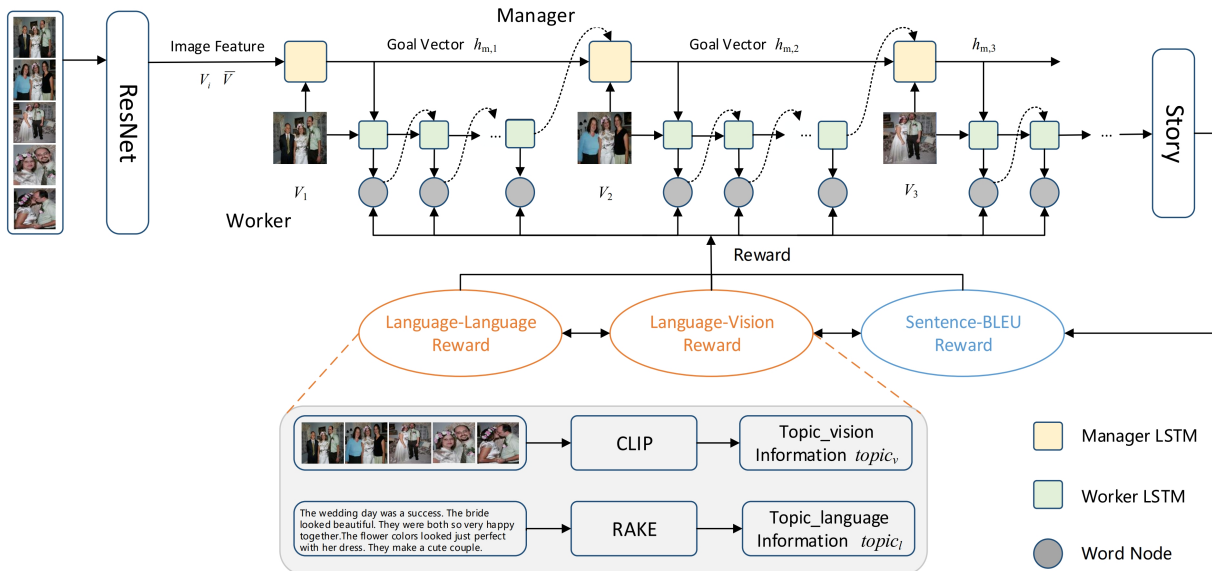


Figure 1: Overview of TARN-VIST. In our model, image features are obtained by the pre-trained ResNet and then fed into the hierarchical decoder which consists of a manager LSTM and a worker LSTM to generate a sample story. Once the candidate story is generated, the two topic consistency rewards are combined to refine the generation process. Furthermore, we also set a classical sentence-level BLEU reward to control the fluency of the generated story.

Inspired by the above idea, we propose a novel method called **Topic Aware Reinforcement Network for Visual StoryTelling (TARN-VIST)**. As depicted in Fig. (1), our model is an encoder-decoder architecture based on reinforcement learning. To be specific, we first use contrastive language-image pre-training (CLIP) and rapid automatic keyword extraction (RAKE) to extract topic information of the stories in the dataset from both visual and linguistic perspectives, respectively. Subsequently, we harness the extracted topic information along with cosine similarity to design the topic-consistent reinforcement learning rewards. The entire framework aggregates the aforementioned rewards and employs the reinforcement learning algorithm to optimize them.

To summarize, our contributions in this paper are as follows:

- We first take advantage of CLIP and RAKE together to extract topic information of stories from both visual and linguistic perspectives.
- To make full use of the topic information, we design reinforcement learning rewards for topic consistency based on the extracted topic information and cosine similarity.
- Experimental results on the VIST dataset and human evaluation demonstrate that our proposed model outperforms most of the leading models on multiple evaluation metrics.

2. Related Work

2.1. Visual Storytelling

Visual storytelling is the task of generating a reasonable paragraph-level story with the image sequence selected from a photo stream as input. It necessitates a deeper understanding of the event progression in the image stream. Based on their technical characteristics, we can classify the current visual storytelling techniques into end-to-end frameworks and multi-stage frameworks. Where end-to-end frameworks mainly make use of CNN structure such as VGG (Simonyan and Zisserman, 2015) or ResNet (He et al., 2016) to extract visual features and then apply LSTM to generate story. For instance, Huang et al. (2016) encode an image sequence by running a recurrent neural network (RNN) over the fc7 vectors of each image and use gated recurrent units (GRUs) for both the image encoder and story decoder. Yu et al. (2017) design a model composed of three hierarchically-attentive RNNs to encode the album photos, select representative photos, and compose the story. Kim et al. (2018) put forward a deep learning network model called GLAC net that combine global-local (glocal) attention and context cascading mechanisms together to generate story. Wang et al. (2019) employ a scene encoder and a photo encoder to detect the scene changes and meanwhile aggregate the scene information. Yang et al. (2019) present a commonsense-driven generative model, which in-

roduces related crucial commonsense from the external knowledge base. Wang et al. (2020) raise a novel graph-based architecture for visual storytelling by modeling the two-level relationships on scene graphs. Jung et al. (2020) propose to explicitly learn to imagine the storyline that bridges the visual gap. Chen et al. (2021) introduce two novel modules that consider both the correlation among candidate concepts and the image-concept correlation. Qi et al. (2021) present a novel Latent Memory-augmented Graph Transformer (LMGT), which directly inherits the merits from the Transformer structure. Xu et al. (2021) put forward a novel imagine-reason-write generation framework (IRW), which employs a relational reasoning module to fully exploit the external knowledge and task-specific knowledge. Braude et al. (2022) develop a novel message-passing-like algorithm for ordered image attention (OIA) that collects interactions across all the images in the sequence. Yang and Jin (2023) propose a multi-tasking memory-augmented framework, which is jointly trained on factual visual storytelling data and unpaired style corpus. Gu et al. (2023) put forward a coherent visual storytelling (CoVS) framework. It introduces an image sequence encoder module and a new parallel top-down attention module. This kind of approaches can handle the mapping relationship between images and text information well. Their model structures are relatively easy and simple to deploy. However, such methods lack the capacity to model the global structure, and the generated stories usually lack diversity.

In comparison to end-to-end frameworks, multi-stage frameworks customarily separate the training process into multiple steps and the outputs of the previous step are often used as the input of the subsequent step. For example, Hsu et al. (2020) put forward a distill-enrich-generate framework. It distills a set of representative words from the input prompts and enriches the word set by using external knowledge graphs. Chu et al. (2021) introduce PR-VIST to represent the input image sequence as a story graph in which it finds the best path to form a storyline. Chen et al. (2022) propose a new sentiment-aware generative model for VIST called SentiStory. It uses a multi-layered sentiment extraction module (MLSEM). Besides, some researchers have attempted to extract the topic information of stories and generate higher-quality content that closely related to the topic information. Li et al. (2020) directly use the query vocabulary of the dataset itself as the topic information of the story, but since the query vocabulary is relatively extensive and a topic vocabulary sometimes corresponds to hundreds of sample stories, it cannot accurately represent the theme of the story vocabulary. In addition, visual storytelling is essentially a multi-modal task where

the input image sequence also contains plentiful information, but previous studies do not consider the topic vocabulary from the visual perspective.

2.2. Reinforcement Learning

Reinforcement learning is an important branch of machine learning where an agent learns how to make optimal decisions and maximize the return of rewards obtained by interactions with a complicated environment (Nie et al., 2023; Plaat et al., 2023). The core concept of reinforcement learning is to learn the most appropriate policy directly through trial and error experiences. The agent performs an action in the environment (Action), and then observes the feedback of the environment (Reward or Punishment), and modifies its behavior in response to the feedback, so as to increase the chances of receiving better rewards in the future (Mnih et al., 2013). In recent years, reinforcement learning has been widely used in numerous fields, such as autonomous driving (Kiran et al., 2022), robot control (Orr and Dutta, 2023), gaming (Perolat et al., 2022) and healthcare (Yu et al., 2023), etc.

For the visual storytelling task, some researchers have already tried to introduce reinforcement learning into the field of visual stories and achieved promising generation results. Chen et al. (2018) propose an adversarial reward learning (AREL) framework to learn an implicit reward function from human demonstrations, and then optimize policy search with the learned reward function. Wang et al. (2018) design the rewards with two critic networks, including a multi-modal and a language-style discriminator to generate relevant and story-style paragraphs. Furthermore, through rethinking about principles that make up high-quality story, Hu et al. (2020) present three assessment criteria which are relevance, coherence and expressiveness, and then employed a reinforcement learning framework called ReCo-RL to capture the essence of these quality criteria. However, to the best of our knowledge, attempts on formulating reward functions for reinforcement learning based on topic information are still in the blank stage. Hence, we first extract the topic information of the dataset and design topic-consistent reinforcement learning reward functions to improve the overall generation process.

3. Approach

In this paper, we define the visual storytelling task as follows: given an image sequence $I = \{i_1, i_2, \dots, i_m\}$, the task aims to produce a human-like story $S = \{s_1, s_2, \dots, s_m\}$ where s_i is a sequence of words describing the i -th image.

3.1. Topic Information Extraction

In this section, we first describe the topic information extraction process. As the visual storytelling task is a multi-modal task, we utilize CLIP (Radford et al., 2021) and RAKE (Rose et al., 2010) to extract the topic information of the story from the visual and linguistic perspectives separately.

We use CLIP to extract the topic information of the story from the visual perspective as shown in algorithm 1. The algorithm mainly includes four steps: candidate-concept extraction, image encoding, text encoding and similarity calculation.

In the candidate concept extraction step, the *Clarifai's* image detection API¹ is applied to retrieve the concepts that appear in the input image sequence. Different from other common object detection algorithms, this API can not only accurately detect the objects, but also effectively identify the scene and text information in the image stream. At this point, it should be noticed that for each picture, only the top three concepts are taken. So we need to remove some noise concepts (such as "people", "person", "men", etc.). Next, we use the sentence pattern "The topic of this photo is {concept}" to assemble the extracted {concept} into the sentence pattern. For example, for the concept "graduation", the sentence formed is "The topic of this photo is graduation". The reason for assembling concepts into sentences is to use CLIP for zero-shot prediction, and the pre-training of CLIP is performed on a large number of "sentence-picture" collections. Apart from this, after assembling into sentences, the text encoder can acquire richer text feature information, which is helpful and beneficial for subsequent model processing.

For the image encoding and text encoding processes, we primarily leverage CLIP's text encoder, image encoder, and zero-sample prediction function. The zero-shot prediction function is realized through the semantic knowledge learned in the pre-training phase, and it can perform image and text classification without the need for additional sample data. Besides, CLIP can comprehend image content through text annotation and image classification. At this stage, the text features of the above text information and the visual features of the input image are extracted by using the text encoder and the image encoder of CLIP, respectively. The overall image features are then calculated by average weighting and adding the individual image features. Finally, we take advantage of cosine similarity to calculate the information with the highest similarity between text features and image features, and then extract the top 1 element from the tensor Similarity[0]. It is worth noting that the topk function is commonly used to retrieve the top k elements

¹<https://clarifai.com/clarifai/main/models>

Algorithm 1 Visual Perspective Topic Information Extraction Process

Input: Image Sequence I with 5 images and $Candidate - Concept$

Output: Topic Information

```
1: Initialise  $Candidate - Concept \leftarrow []$ 
2: /* Candidate-Concept Extraction */
3: for  $i = 1$  to 5 do
4:   Extract top-3 concept  $c_i^j$  from Image  $i$  with
    $Clarifai's$  Image Detection API
5:   Filter out some useless concepts and assemble
   concepts into sentences  $s_i^j$ 
6:    $Candidate - Concept.append(s_i^j)$ 
7: end for
8: /* Image Encoding */
9: for  $i = 1$  to 5 do
10:  Image-feature  $i = CLIP-Image-Encoder(Image\ i)$ 
11: end for
12: Image-mean-feature = mean (Image-feature  $i$ )
13: /* Text Encoding */
14: Text-features=CLIP-Text-Encoder( $Candidate-Concept$ )
15: /* Similarity Calculation */
16: Similarity = Image-mean-feature @ Text-features
17: Topic Information = Similarity[0]. topk(1)
18: return Topic Information
```

from a tensor in practice and the eventual extracted element is a word or phrase, which is the desired image topic information.

Furthermore, we also apply RAKE to extract the topic vocabulary of the stories as the topic information from the linguistic perspective in the dataset. Compared to latent dirichlet allocation (LDA), RAKE can not only rapidly extract keywords from extensive textual data but also excels in identify and extract multi-word phrases, rather than limited to some single keywords. It is versatile, effective, fast and applicable for various text types, including both long and short texts. The core idea behind RAKE is to identify words or phrases with high frequency and importance in the text as the topic information.

3.2. Topic Aware Storytelling Model

The Encoder We first feed the image sequences into ResNet (He et al., 2016) and extract their high-level image features. For the follow-up requirements, we compute the average value of all the image features v_i as follows:

$$\bar{V} = \frac{1}{n} \sum_i v_i \quad (1)$$

Then, the image features v_i and overall image sequences features \bar{V} are combined by direct con-

catenation along the channel dimension and sent into decoder together.

The Decoder The decoder is hierarchical which involves a manager LSTM and a worker LSTM (Huang et al., 2019). The manager LSTM serves as a supervisor to control the overall flow of the story, which is denoted as follows:

$$h_{m,i} = \text{LSTM}_M([\bar{V}; v_i; h_{w,i-1}^T], h_{m,i-1}) \quad (2)$$

When describing i -th image, the manager LSTM will take three kinds of information into consideration, which are: 1) overall information of the image sequence \bar{V} ; 2) image information in the i -th image v_i ; 3) sentences generated for previous image $h_{w,i-1}^T$. Then, the manager LSTM will predict a hidden state $h_{m,i}$ as the goal vector and passes the vector to the worker LSTM.

The worker LSTM attempts to complete the generation of word description based on the goal vector. When generating t -th word for the i -th image, the worker LSTM decoder will take the current image information v_i , word embedding of the previously generated word e_i^{t-1} and goal vector $h_{m,i}$ as input to predict its hidden state $h_{w,i}^t$. Then, the worker LSTM enforces a linear layer $f(\cdot)$ to approximate the probability of choosing the next word:

$$h_{w,i}^t = \text{LSTM}([v_i; h_{m,i}; e_i^{t-1}], h_{w,i}^{t-1}) \quad (3)$$

$$p_\theta(y_i^t | y_i^{1:t-1}, v_i, \bar{V}) = \text{softmax}(f(h_{w,i}^t)) \quad (4)$$

Topic Consistency Rewards Design For a given image sequence, we define that the extracted topic information from the perspective of vision and language are $topic_v$ and $topic_l$. For the story generated by the model, we use RAKE to extract its topic information $topic_c$, and then propose three reward functions r_{bleu} , $r_{topic-cv}$ and $r_{topic-cl}$ based on the text cosine similarity and topic consistency:

$$r_{bleu} = \text{sentence-bleu}(story_c, story_g) \quad (5)$$

$$r_{topic-cv} = \text{cosine-similarity}(topic_c, topic_v) \quad (6)$$

$$r_{topic-cl} = \text{cosine-similarity}(topic_c, topic_l) \quad (7)$$

where $story_c$ means the story generated by our model and $story_g$ represents the corresponding ground-truth story. In summary, the total reward function is:

$$r = \lambda \cdot r_{bleu} + \gamma \cdot r_{topic-cv} + \eta \cdot r_{topic-cl} \quad (8)$$

where λ , γ and η are the hyper-parameters controlling the proportion of each part.

3.3. Training

Reinforcement learning can learn a policy that encourages the model to focus more on those key aspects and then generate stories with more emotion and fluency through maximizing the given reward:

$$J_{RL}(\beta) = \sum_{Y, V \in D'} E_{y_i \sim \pi_i} [(b - r(y_i)) \log \pi_i] \quad (9)$$

$$r(y_i) = \lambda r_{bleu}(y_i) + \gamma r_{topic-cv}(y_i) + \eta r_{topic-cl}(y_i) \quad (10)$$

where $\pi \equiv p_\theta(y_i | v_i, \bar{V})$ is the policy and b is the baseline that reduces the given reward variance.

During training, we find that optimizing only with the reinforcement learning loss has the potential to increase expected rewards while sacrificing the quality of our generative model. To address this concern, we first train our model with maximum likelihood estimation (MLE) and then continue to train the model jointly with reinforcement loss $Loss_{RL}$ and MLE loss $Loss_{MLE}$. This training strategy can strike a balance between reward optimization and model fidelity, resulting in improved storytelling ability. MLE loss and the mixed training loss $Loss_{mixed}$ are defined as follows:

$$Loss_{MLE} = \sum_{Y, V \in D'} \sum_{i=1}^N \sum_{t=1}^T -\log P(W_{i,t} = Y_{i,t}) \quad (11)$$

$$Loss_{mixed} = \omega Loss_{RL} + (1 - \omega) Loss_{MLE} \quad (12)$$

Following Hu et al. (2020), we set the $\omega=0.5$ in eq. (12) as it is a good trade-off between reinforcement learning loss and MLE loss.

4. Experiments

4.1. Dataset and Evaluation Metrics

We train and evaluate our model on the VIST dataset (Huang et al., 2016), which is the most widely used dataset for the visual storytelling task. The VIST dataset¹ includes 10,117 Flickr albums with 21,0819 unique images and it is split into training/validation/testing sets with 400,98/4,988/5,050 samples, respectively.

During the whole experiment process, we utilize two groups of automatic metrics for the quantitative evaluation. One group consists of several surface-level-based similarity measures, including BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), ROUGE (Lin, 2004), CIDEr

¹<https://visionandlanguage.net/VIST/>

Method	BLEU-1	BLEU-2	BLEU-4	METEOR	ROUGE	CIDEr	SPICE
Seq2seq	-	-	3.5	31.4	-	6.8	-
H-Attn-Rank	-	-	-	34.1	29.5	7.5	-
BARNN	-	-	-	33.3	-	-	-
SRT	43.4	21.4	5.2	12.3	-	11.4	-
XE-ss	62.3	38.2	13.7	34.8	29.7	8.7	-
AREL	63.7	39.0	14.0	35.0	29.6	9.5	8.9
HPSR	61.9	37.8	12.2	34.4	31.2	8.0	-
HSRL	-	-	12.3	35.2	30.8	10.7	7.5
SGVST	65.1	40.1	14.7	35.8	29.9	9.8	-
ReCo-RL	-	-	12.4	33.9	29.9	8.6	8.3
INet	64.4	40.1	14.7	35.6	29.6	11.0	-
IRW	66.7	41.6	15.4	35.6	29.6	11.0	-
CKAKS	-	-	12.0	35.4	30.0	10.5	-
LGMT	67.5	41.6	15.1	35.6	29.7	10.0	-
Sentistory	64.8	39.8	14.2	35.3	29.8	9.7	-
TARN-VIST	69.0	43.5	13.4	35.8	29.5	12.1	11.3

Table 1: Quantitative results on the VIST dataset for surface-level-based automatic metrics. For all these metrics, higher score means better performance.

Method	BERTScore	BARTScore	BLEURT
KE-VIST(No KG)	28.25	17.21	43.63
KE-VIST (With OpenIE)	29.12	17.93	46.85
KE-VIST (with VG)	29.16	18.03	47.54
PR-VIST	27.64	18.09	48.92
TARN-VIST	30.47	18.51	49.43

Table 2: Quantitative results on the VIST dataset for semantic understanding evaluation metric. For all these metrics, higher score means better performance.

(Vedantam et al., 2015) and SPICE (Anderson et al., 2016). In particular, BLEU is a classical metric that computes the geometric average of overlapping n-grams between the generated sentence and the reference sentence. Commonly used ones are BLEU-1, BLEU-2 and BLEU-4. METEOR calculates the sentence-level similarity scores based on the harmonic mean of uni-gram recall and precision. It is sensitive to the text length. ROUGE is a recall-based metric which captures the length of the most common sub-sequence between the generated story and the reference. CIDEr adopts higher order n-grams to account for fluency and uses term frequency-inverse document frequency (TF-IDF) weighting for each n-gram to down-weight commonly occurring ones. SPICE is an automated caption evaluation metric that defined over scene graphs. It can effectively capture and reflect human judgments over model-generated captions.

What’s more, since previous researchers (Chen et al., 2018) have pointed out that current surface-level automatic metrics may correlate poorly with human judgments, we also use BERTScore (Zhang et al., 2020), BARTScore (Yuan et al., 2021) and BLEURT (Sellam et al., 2020) in the experiments. Among them, BERTScore measures the BERT embedding similarity between each token as a

more semantic and robust measure rather than common string-match-based similarity measure. BARTScore is an unsupervised metric that does not require human judgment. It can better evaluate the generated text from various aspects. BLEURT is a transfer learning based metric for natural language generation. It excels in capturing significant semantic similarities between sentences.

4.2. Implementation Details

We implement our framework in PyTorch (Paszke et al., 2019). In the phase of encoding, we leverage the official pre-trained Resnet-152 model¹ to extract the deeper image features. Additionally, in the phase of decoding, the hidden size of manager LSTM and worker LSTM are both set to 512. When calculating the cosine similarity, we employ the Bert-base model from HuggingFace². Furthermore, during the model training with MLE loss, we set learning rate to 0.0001 and the dropout rate to 0.6. Our experiments are conducted on GeForce RTX 3090 GPU with a batch size of 128.

¹<https://huggingface.co/microsoft/resnet-152>

²<https://huggingface.co/google-bert/bert-base-uncased>

Model	BLEU-4	METEOR	CIDEr	SPICE
Baseline	12.40	33.90	8.60	8.30
Baseline+ r_{bleu}	12.82	35.36	11.58	10.84
Baseline+ $r_{bleu}+r_{topic-cv}$	13.10	35.62	11.70	11.37
Baseline+ $r_{bleu}+r_{topic-cl}$	13.44	35.80	10.99	11.36
TARN-VIST	13.46	35.88	12.07	11.25

Table 3: Ablation experiment results on different combinations of the reward functions. Note that here our basic model is ReCo-RL.

γ	η	BLEU-1	BLEU-2	BLEU-3	BLEU-4	ROUGE	METEOR	CIDEr	SPICE
0.3	0.7	68.43	42.48	23.51	12.85	29.14	35.52	11.23	10.76
0.4	0.6	68.38	42.55	23.56	12.88	29.21	35.62	11.69	10.85
0.5	0.5	69.01	43.56	24.27	13.46	29.28	35.88	12.07	11.36
0.6	0.4	68.58	42.79	23.77	13.09	29.24	35.63	11.91	10.94
0.7	0.3	68.24	42.46	23.57	12.92	29.12	35.53	11.44	10.78
0.8	0.2	68.24	42.58	23.66	12.90	29.24	35.48	11.69	10.77

Table 4: Ablation experiment results of TARN-VIST with different γ and θ .

Aspect	KE-VIST	TARN-VIST	Tie	PR-VIST	TARN-VIST	Tie
Relevance	19%	67%	14%	30%	60%	10%
Coherence	22%	50%	28%	18%	62%	20%
Information Richness	28%	64%	8%	33%	50%	17%

Table 5: Human Pairwise Evaluation between TARN-VIST and other models. For each pairwise comparison, each of the three columns stands for the percentage of volunteers that prefer this story to the other one, and consider both stories are of equal quality.

4.3. Quantitative Evaluation

We compare our model with the following baselines: (1) Seq2seq (Huang et al., 2016), a classical end-to-end sequence learning approach; (2) H-Attn-Rank (Yu et al., 2017), a model composed of three hierarchically-attentive recurrent neural nets (RNNs); (3) BARNN (Liu et al., 2017), an attention-based RNN with a skip gated recurrent unit (GRU) which leverages the semantic relation between photo streams and stories; (4) Show, Reward and Tell (SRT) (Wang et al., 2018), a hierarchical generative model with reinforcement learning and adversarial training; (5) XE-ss and AREL (Chen et al., 2018), an adversarial reward learning framework with imitation learning and GAN; (6) HPSR (Wang et al., 2019), a novel model with a hierarchical photo-scene encoder and a reconstructor; (7) HSRL (Huang et al., 2019), a structured reinforcement learning approach with hierarchical decoder; (8) SGVST (Wang et al., 2020), a framework based on scene graphs with GCN and TCN; (9) ReCo-RL (Hu et al., 2020), a reinforcement learning algorithm while the reward is relevance, coherence and expressiveness; (10) INet (Jung et al., 2020), a hide-and-tell training scheme; (11) KE-VIST (Hsu et al., 2020), a three-stage generation framework; (12) IRW (Xu et al., 2021), an imagine-reason-write generation framework; (13) CKAKS (Chen et al., 2021), an extract-enrich-generate framework; (14)

PR-VIST (Chu et al., 2021), a plot and rework framework; (15) LMGIT (Qi et al., 2021), a Transformer based framework for visual story generation; (16) Sentistory (Chen et al., 2022), a sentiment-aware generative model.

In table (1), we observe that our proposed model TARN-VIST achieves state-of-the-art on the BLEU-1, BLEU-2, METEOR, CIDEr and SPICE. Specifically, compared with the previous best model, TARN-VIST makes the improvement by 1.5% on BLEU-1, 1.9% on BLEU-2, 0.2% on METEOR, 1.1% on CIDEr and 2.4% on SPICE. In addition, the performance of BLEU-4 and ROUGE is also very competitive. In table (2), we find that TARN-VIST achieves state-of-the-art on all the metrics. The reason for the excellent performance of TARN-VIST can be attributed to our proposed composite reward function, which can enable the model to effectively achieve topic alignment and generate high-quality stories.

4.4. Ablation Study

To analyze the effectiveness of our proposed reward functions, we carry out a number of ablation experiments on different combinations of the reward functions. The experimental results are shown in Table (3). Note that our basic model is ReCo-RL (Hu et al., 2020), which is an encoder-decoder framework trained through reinforcement

Images	
Topic Information	Topic_vision: Seashore; Topic_language: Good Day;
Images	
Topic Information	Topic_vision: Birthday; Topic_language: Party;
Images	
Topic Information	Topic_vision: Parade; Topic_language: Parade;

Figure 2: Examples of extracted topic information

Method	
Topic Information	Topic_vision: Wedding; Topic_language: Wedding day;
Ground Truth	The wedding day was a success. The bride looked beautiful. They were both so very happy together. The flower colors looked just perfect with her dress. They make a cute couple.
PR-VIST	[female] and [male] took a family photo together . it was [male] 's birthday party . the groom and his wedding were cut off . their mom was so proud of them . her mom got married today , but [female] was very excited to be there . [female] is now getting married today .
KE-VIST	the bride and groom were getting ready for their wedding . they looked so happy to be there . it was a beautiful night . everyone danced . then came out after .
TARN-VIST	the bride and groom are married . The bride is taking a picture with her friends .she was happy and her husband was so happy . after the wedding , we had a great time with her wedding . they were very excited to get a picture of the wedding.

Figure 3: Example story generated from TARN-VIST and several competitive baselines.

learning. From the table, we can see that when BLEU is used as the reward function of reinforcement learning, the performance of the model has been improved to a certain extent compared with the baseline. For instance, the evaluation indicators of CIDEr and SPICE have increased by 2.98% and 2.54%, respectively. Besides, when using $r_{topic-cv}$ or $r_{topic-cl}$ alone, the model also has different degrees of improvement. In addition, when r_{bleu} and $r_{topic-cv}$ are combined, the relative improvements of CIDEr and SPICE are 0.12% and 0.53%. Furthermore, when r_{bleu} , $r_{topic-cl}$ and $r_{topic-cv}$ are integrated together, the model performs best on most evaluation metrics.

In eq. (10), λ , γ and η are three important hyper-parameters. Therefore, we also implement different combinations of the hyper-parameters. To be no-

ticed, λ is set to 0.5 by the control variable method and the sum of γ and η is set to 1. We carry out several experiments on different value pairs of γ and η . The experimental results are shown in Table (4). It can be obviously seen that when γ and η are all equal to 0.5, the model can better balance the effects of the three reward functions, thus making the model reach a local optimal result.

4.5. Human Evaluation

Previous visual storytelling work has highlighted that current string-match-based automatic metrics are not a perfect way to evaluate the model's performance (Chen et al., 2018). Therefore, we further conduct pairwise human evaluation to assess the subjective quality of our model. In particular,

we randomly selected 100 generated sample stories from the test set and invited 20 well-educated volunteers to evaluate results based on the story relevance, the story coherence, and the story information richness. All the evaluators have received college degree or above. Meanwhile, the choices are shuffled before evaluation to avoid any bias. Our model is compared with KE-VIST (Hsu et al., 2020) and PR-VIST (Chu et al., 2021). For each sample pair, given a photo stream and the two stories generated by two models, the volunteers are asked to perform pairwise evaluation based on the relevance, coherence, and information richness. Among them, relevance assesses how the story accurately describes what is happening in the image sequence. Coherence evaluates whether the story is semantically coherent with other sentences. Information richness assesses whether the information is fully expressed in the story. To be more fair, we also provide a neutral option for cases that volunteers consider the two stories to be equally good on one particular criterion. The experimental results are displayed in Table (5). From the table, we can observe that our proposed model significantly outperforms other models in terms of these three metrics.

4.6. Qualitative Evaluation

In order to provide a more intuitive analysis on the rationality of our proposed topic information extraction method, we randomly choose several image sequences, and the topic information $topic_v$ and $topic_l$ extracted from the visual and linguistic perspectives are also displayed. The experimental results are shown in Fig. (2). It can be found that both the topic information can well summarize the content of the image sequence. Take the first set of image sequences as an illustration, the topic information extracted from the visual perspective is "Seashore" and the topic information extracted from the linguistic perspective is "Good Day". We infer that when extracting the topic information from the visual perspective, multiple beach-related vocabulary would be extracted when using the image recognition module to extract the concept of the input image sequence, so the final topic information is "Seashore". Meanwhile, when extracting topic vocabulary from the linguistic perspective, since the story in the dataset describes the picture from the protagonist's point of view, the content is probably "spent a good day at the beach" and other information. Therefore, the extracted topic information is "Good Day".

In addition, to qualitatively evaluate the TARN-VIST, we also compare the sample stories generated by this model with the other two models in a specific image sequence. The results are exhibited in Fig. (3). We can see that the image sequence

topics extracted by our model are closely related to weddings. For the stories generated by the PR-VIST model, "birthday party", "their room was so proud of them" and "her mom got married today" are less pertinent to the topic of the input image sequence. For the stories generated by the KE-VIST model, the two short sentences "everyone danced" and "then came out after" have serious expression problems such as grammatical errors and low coherence. On the contrary, the stories generated by the TARN-VIST are more relevant than those generated by the other two models. For instance, in each sentence, there are words such as "bride", "groom" and "wedding". The presence of a story brings the tone of the story and the topic information closer together so that the stories generated by TARN-VIST are more coherent.

5. Conclusion and Future Work

In this paper, we propose TARN-VIST, an innovative topic aware reinforcement network for visual storytelling to generate more cogent story. We first utilize CLIP and RAKE to mine the topic information of the story from both visual and linguistic perspectives. Subsequently, we employ the reinforcement learning and design topic consistency rewards to refine the generation process. Extensive experimental results on the benchmark dataset demonstrate that our model outperforms most competitive baselines across a number of evaluation metrics.

In future work, we will explore grammar and discourse structure in visual storytelling tasks, as they play a key role in the accuracy, coherence, and readability of the generated stories. Besides, it is also interesting and practical to further analyse linguistic style to improve the quality and diversity of generated stories.

6. Acknowledgements

This work was partially supported by Postgraduate Research & Practice Innovation Program of Jiangsu Province, National Natural Science Foundation of China (NSFC Grant No. 61773272, 61272258, 61301299, 615-72085, 61170124, 61272005, 62376041), Provincial Natural Science Foundation of Jiangsu, China (Grant No. BK20151254, BK201512-60), Science and Education Innovation based Cloud Data fusion Foundation of Science and Technology Development Center of Education Ministry, China (2017B03112), Six talent peaks Project in Jiangsu Province, China (DZXX-027), Key Laboratory of Symbolic Computation and Knowledge Engineering of Ministry of Education, Jilin University, China (Grant No.93K172016K08), and Collaborative Innovation Center of Novel Software Technology and Industrialization, China.

7. Bibliographical References

- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. SPICE: semantic propositional image caption evaluation. In *Computer Vision - ECCV 14th European Conference*, pages 382–398, Amsterdam, The Netherlands. Springer.
- Satanjeev Banerjee and Alon Lavie. 2005. ME-TEOR: an automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization@ACL 2005*, Ann Arbor, Michigan, USA, June 29, 2005, pages 65–72, Ann Arbor, Michigan, USA. Association for Computational Linguistics.
- Tom Braude, Idan Schwartz, Alexander G. Schwing, and Ariel Shamir. 2022. Ordered attention for coherent visual storytelling. In *MM '22: The 30th ACM International Conference on Multimedia*, pages 3310–3318, Lisboa, Portugal. ACM.
- Hong Chen, Yifei Huang, Hiroya Takamura, and Hideki Nakayama. 2021. Commonsense knowledge aware concept selection for diverse and informative visual storytelling. In *Thirty-Fifth AAAI Conference on Artificial Intelligence*, pages 999–1008, Virtual Event. AAAI Press.
- Wei Chen, Xuefeng Liu, and Jianwei Niu. 2022. Sentistory: A multi-layered sentiment-aware generative model for visual storytelling. *IEEE Trans. Circuits Syst. Video Technol.*, 32(11):8051–8064.
- Wenhu Chen, Xin Wang, 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, ACL, Volume 1: Long Papers*, pages 899–909, Melbourne, Australia. Association for Computational Linguistics.
- Yun-Wei Chu, Chi-Yang Hsu, Ting-Hao Kenneth Huang, and Lun-Wei Ku. 2021. Plot and rework: Modeling storylines for visual storytelling. *CoRR*, abs/2105.06950.
- Ruichao Fan, Hanli Wang, Jinjing Gu, and Xianhui Liu. 2021. Visual storytelling with hierarchical BERT semantic guidance. In *MMAAsia '21: ACM Multimedia Asia*, pages 24:1–24:7, Gold Coast, Australia. ACM.
- Jinjing Gu, Hanli Wang, and Ruichao Fan. 2023. Coherent visual storytelling via parallel top-down visual and topic attention. *IEEE Trans. Circuits Syst. Video Technol.*, 33(1):257–268.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 770–778, Las Vegas, NV, USA. IEEE Computer Society.
- Chao-Chun Hsu, Zi-Yuan Chen, Chi-Yang Hsu, Chih-Chia Li, Tzu-Yuan Lin, Ting-Hao Kenneth Huang, and Lun-Wei Ku. 2020. Knowledge-enriched visual storytelling. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, pages 7952–7960, Honolulu, Hawaii, USA. AAAI Press.
- Junjie Hu, Yu Cheng, Zhe Gan, Jingjing Liu, Jianfeng Gao, and Graham Neubig. 2020. What makes A good story? designing composite rewards for visual storytelling. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, pages 7969–7976, New York, NY, USA. AAAI Press.
- Qiuyuan Huang, Zhe Gan, Asli Celikyilmaz, Dapeng Oliver Wu, Jianfeng Wang, and Xiaodong He. 2019. Hierarchically structured reinforcement learning for topically coherent visual story generation. In *The Thirty-Third AAAI Conference on Artificial Intelligence*, pages 8465–8472, Honolulu, Hawaii, USA. AAAI Press.
- Ting-Hao (Kenneth) Huang, Francis Ferraro, Nasrin Mostafazadeh, Ishan Misra, Aishwarya Agrawal, Jacob Devlin, Ross B. Girshick, Xiaodong He, Pushmeet Kohli, Dhruv Batra, C. Lawrence Zitnick, Devi Parikh, Lucy Vanderwende, Michel Galley, and Margaret Mitchell. 2016. Visual storytelling. In *NAACL HLT The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1233–1239, San Diego California, USA. The Association for Computational Linguistics.
- Yunjae Jung, Dahun Kim, Sanghyun Woo, Kyungsu Kim, Sungjin Kim, and In So Kweon. 2020. Hide-and-tell: Learning to bridge photo streams for visual storytelling. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI*, pages 11213–11220, New York, NY, USA. AAAI Press.
- Taehyeong Kim, Min-Oh Heo, Seonil Son, Kyoung-Wha Park, and Byoung-Tak Zhang. 2018. GLAC net: Glocal attention cascading networks for multi-image cued story generation. *CoRR*, abs/1805.10973.
- B. Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A. Al Sallab, Senthil Kumar

- Yogamani, and Patrick Pérez. 2022. Deep reinforcement learning for autonomous driving: A survey. *IEEE Trans. Intell. Transp. Syst.*, 23(6):4909–4926.
- Jiacheng Li, Siliang Tang, Juncheng Li, Jun Xiao, Fei Wu, Shiliang Pu, and Yueting Zhuang. 2020. Topic adaptation and prototype encoding for few-shot visual storytelling. In *MM '20: The 28th ACM International Conference on Multimedia*, pages 4208–4216, Virtual Event / Seattle, WA. ACM.
- Xin Li, Chunping Liu, and Yi Ji. 2023. Associative learning network for coherent visual storytelling. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, page 5, RHODES ISLAND, GREECE. IEEE.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics*, page 74–81, Barcelona, Spain. ACL.
- Yu Liu, Jianlong Fu, Tao Mei, and Changen Chen. 2017. Let your photos talk: Generating narrative paragraph for photo stream via bidirectional attention recurrent neural networks. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, (AAAI-17)*, page 8, San Francisco, California USA. AAAI Press.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin A. Riedmiller. 2013. Playing atari with deep reinforcement learning. *CoRR*, abs/1312.5602.
- Mingshuo Nie, Dongming Chen, and Dongqi Wang. 2023. Reinforcement learning on graphs: A survey. *IEEE Trans. Emerg. Top. Comput. Intell.*, 7(4):1065–1082.
- James Orr and Ayan Dutta. 2023. Multi-agent deep reinforcement learning for multi-robot applications: A survey. *Sensors*, 23(7):3625.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, PA, USA. ACL.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, page 12, Vancouver, BC, Canada. Curran Associates, Inc.
- Julien Perolat, Bart De Vylder, Daniel Hennes, Eugene Tarassov, Florian Strub, Vincent de Boer, Paul Muller, Jerome T. Connor, Neil Burch, Thomas Anthony, Stephen McAleer, Romuald Elie, Sarah H. Cen, Zhe Wang, Audrunas Gruslys, Aleksandra Malysheva, Mina Khan, Sherjil Ozair, Finbarr Timbers, Toby Pohlen, Tom Eccles, Mark Rowland, Marc Lanctot, Jean-Baptiste Lespiau, Bilal Piot, Shayegan Omidshafiei, Edward Lockhart, Laurent Sifre, Nathalie Beauguerlange, Remi Munos, David Silver, Satinder Singh, Demis Hassabi, and Karl Tuyls. 2022. Mastering the game of stratego with model-free multiagent reinforcement learning. *Science*, 378(6623):990–996.
- Aske Plaat, Walter A. Kusters, and Mike Preuss. 2023. High-accuracy model-based reinforcement learning, a survey. *Artif. Intell. Rev.*, 56(9):9541–9573.
- Mengshi Qi, Jie Qin, Di Huang, Zhiqiang Shen, Yi Yang, and Jiebo Luo. 2021. Latent memory-augmented graph transformer for visual storytelling. In *MM '21: ACM Multimedia Conference*, pages 4892–4901, Virtual Event. ACM.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML*, pages 8748–8763, Virtual Event. PMLR.
- Stuart J Rose, David W Engel, Nicholas O Cramer, and Wendy E Cowley. 2010. Automatic keyword extraction from individual documents.
- Thibault Sellam, Dipanjan Das, and Ankur P. Parikh. 2020. BLEURT: learning robust metrics for text generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL*, pages 7881–7892, Online event. Association for Computational Linguistics.
- Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *3rd International Conference on Learning Representations, ICLR*, San Diego, CA, USA.

- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 4566–4575, Boston, MA, USA. IEEE Computer Society.
- Bairui Wang, Lin Ma, Wei Zhang, Wenhao Jiang, and Feng Zhang. 2019. Hierarchical photo-scene encoder for album storytelling. In *The Thirty-Third AAAI Conference on Artificial Intelligence*, pages 8909–8916, Honolulu, Hawaii, USA. AAAI Press.
- Jing Wang, Jianlong Fu, Jinhui Tang, Zechao Li, and Tao Mei. 2018. Show, reward and tell: Automatic generation of narrative paragraph from photo stream by adversarial training. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, page 8, New Orleans, Louisiana, USA. AAAI Press.
- Ruize Wang, Zhongyu Wei, Piji Li, Qi Zhang, and Xuanjing Huang. 2020. Storytelling from an image stream using scene graphs. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, pages 9185–9192, New York, NY, USA. AAAI Press.
- Chunpu Xu, Min Yang, Chengming Li, Ying Shen, Xiang Ao, and Ruifeng Xu. 2021. Imagine, reason and write: Visual storytelling with graph knowledge and relational reasoning. In *Thirty-Fifth AAAI Conference on Artificial Intelligence*, pages 3022–3029, Virtual Event. AAAI Press.
- Dingyi Yang and Qin Jin. 2023. Attractive storyteller: Stylized visual storytelling with unpaired text. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11053–11066, Toronto, Canada. Association for Computational Linguistics.
- 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 *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI*, pages 5356–5362, Macao, China. ijcai.org.
- Chao Yu, Jiming Liu, Shamim Nemati, and Guosheng Yin. 2023. Reinforcement learning in healthcare: A survey. *ACM Comput. Surv.*, 55(2):5:1–5:36.
- Licheng Yu, Mohit Bansal, and Tamara L. Berg. 2017. Hierarchically-attentive RNN for album summarization and storytelling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP*, pages 966–971, Copenhagen, Denmark. The Association for Computational Linguistics.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. Bartscore: Evaluating generated text as text generation. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems*, pages 27263–27277, virtual event. MIT Press.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with BERT. In *8th International Conference on Learning Representations, ICLR*, Addis Ababa, Ethiopia. OpenReview.net.