

Keep it Consistent: Topic-Aware Storytelling from an Image Stream via Iterative Multi-agent Communication

Ruize Wang¹, Zhongyu Wei^{2,3*}, Ying Cheng¹, Piji Li⁴,
Haijun Shan⁵, Ji Zhang⁵, Qi Zhang⁶, Xuanjing Huang⁶

¹ Academy for Engineering and Technology, Fudan University, China

² School of Data Science, Fudan University, China

³ Research Institute of Intelligent and Complex Systems, Fudan University, China

⁴ Tencent AI Lab, China, ⁵ Zhejiang Lab, China

⁶ School of Computer Science, Fudan University, China

{[rzwang18](mailto:rzwang18@fudan.edu.cn), [zywei](mailto:zywei@fudan.edu.cn), [chengy18](mailto:chengy18@fudan.edu.cn), [qz](mailto:qz@fudan.edu.cn), [xjhuang](mailto:xjhuang@fudan.edu.cn)}@fudan.edu.cn;

lipiji.pz@gmail.com; shanhaijun@zhejianglab.com; Ji.Zhang@usq.edu.au

Abstract

Visual storytelling aims to generate a narrative paragraph from a sequence of images automatically. Existing approaches construct text description independently for each image and roughly concatenate them as a story, which leads to the problem of generating semantically incoherent content. In this paper, we propose a new way for visual storytelling by introducing a topic description task to detect the global semantic context of an image stream. A story is then constructed with the guidance of the topic description. In order to combine the two generation tasks, we propose a multi-agent communication framework that regards the topic description generator and the story generator as two agents and learn them simultaneously via iterative updating mechanism. We validate our approach on VIST dataset, where quantitative results, ablations, and human evaluation demonstrate our method’s good ability in generating stories with higher quality compared to state-of-the-art methods.

1 Introduction

Image-to-text generation is an important topic in artificial intelligence (AI) which connects computer vision (CV) and natural language processing (NLP). Popular tasks include image captioning (Karpathy and Fei-Fei, 2015; Ren et al., 2017; Vinyals et al., 2017) and question answering (Antol et al., 2015; Yu et al., 2017b; Singh et al., 2019), aiming at generating a short sentence or a phrase conditioned on certain visual information. With the development of deep learning and reinforcement learning models, recent years witness promising improvement of these tasks for single-image-to-single-sentence generation.

Visual storytelling moves one step further, extending the input and output dimension to a sequence of images and a sequence of sentences. It requires the model to understand the main idea of an image stream and generate coherent sentences. Most of existing methods (Huang et al., 2016; Liu et al., 2017; Yu et al., 2017a; Wang et al., 2018a; Wang et al., 2020) for visual storytelling extend approaches of image captioning without considering topic information of the image sequence, which causes the problem of generating semantically incoherent content.

An example of visual storytelling can be seen in Figure 1. An image stream with five images about a car accident is presented accompanied with two stories. One is constructed by a human annotator and the other is produced by an automatic storytelling approach. There are two problems with the machine generated story. First, the sentiment expressed in the text is inappropriate. In face of a terrible car accident, the model uses some words with positive emotion, i.e., “*excited*” and “*fun*”. Second, some sentence is uninformative. The sentence “*we got to see a lot of different things at the event*” provides little information about the car accident. This example shows that topic information about the image sequence is important for the story generator to produce an informative and semantically coherent story.

In this paper, we introduce a novel task of topic description generation to detect the global semantic context of an image sequence and generate a story with the guidance of such topic information. In

* Corresponding author

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.



Topic: car accident

Ground-truth: i came across a terrible car accident. one of the vehicles was completely destroyed. they had to bring in a tow truck to remove the wreck. the other car was badly damaged as well. it took them a while to clear that part of the street again.

Existing methods: the police were in the accident. the cars was damaged. the police were very excited to see the car. we got to see a lot of different things at the event. the car truck was a lot of fun.

Figure 1: An example of visual storytelling from VIST dataset consisting of five images with the topic of car accident. Two stories presented are from an automatic approach and a human annotator respectively. Wavy lines highlight the inappropriate sentiment of the machine-generated sentence. Underlines indicate that the sentence provides little on-topic information about the image sequence.

practice, we propose a framework named **Topic-Aware Visual Story Telling (TAVST)** to tackle the two generation tasks: (1) topic description generation for the given image stream; (2) story generation with the guidance of the topic information. To effectively combine these two tasks, we propose a multi-agent communication framework that regards the topic description generator and the story generator as two agents. In order to enable the interaction of these two agents for information sharing, an iterative updating (IU) module is incorporated into the framework. Extensive experiments on the VIST dataset (Huang et al., 2016) show that our framework achieves better performance compared to state-of-the-art methods. Human evaluation also demonstrates that the stories generated by our model are better in terms of relevance, expressiveness and topic consistency.

2 Approach

Given an image stream $x = (x_1, \dots, x_N)$, where N is the number of images, we aim to output a topic description y_{topic} and N sub-stories to form a complete story $y = (y_1, \dots, y_N)$. The proposed framework mainly includes three stages, namely visual encoding, initial stage of generation and iterative updating (IU). The visual encoder (§2.1) is employed to extract image features as visual context vectors. In the initial stage, we have the initial version of the two generation agents. The initial topic description generator (§2.2) takes visual context vectors as input and generates a topic vector. The initial story generator (§2.3) combines the topic vector and visual context vectors via co-attention mechanism and construct the initial version of story. Considering that the generated story can also benefit the topic description generator, the two agents communicate with each other in the IU module (§2.4) via message passing mechanism as fine tuning. The overall architecture of our proposed model is shown in Figure 2. Each of these modules will be described in details in the following sections.

2.1 Visual Encoder

Given an image stream x with N images, we first extract the high-level visual features f_i of each image x_i ($i \in 1, \dots, N$) through a CNN model. Then for the whole image stream, following the previous work (Wang et al., 2018b), a bidirectional gated recurrent unit (biGRU) is employed as the visual encoder:

$$\begin{aligned} f_i &= \text{ResNet}(x_i) \\ \overrightarrow{h}_{i,t} &= \overrightarrow{\text{GRU}}(f_i, \overrightarrow{h}_{i,t-1}) \\ \overleftarrow{h}_{i,t} &= \overleftarrow{\text{GRU}}(f_i, \overleftarrow{h}_{i,t+1}) \end{aligned}$$

where $\overrightarrow{h}_{i,t}$ is the forward hidden state at time step t of i -th visual feature f_i , while the $\overleftarrow{h}_{i,t}$ is the backward one. At each time step, the visual features are sequentially fed into the visual encoder to obtain visual context vectors, which have integrated the visual information from the observed images. At last, fused

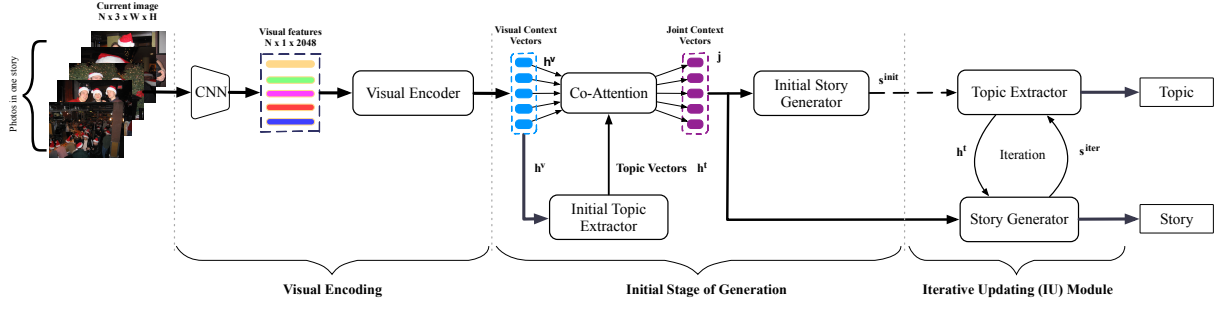


Figure 2: Overall architecture of our TAVST model.

with the original visual representation, the final visual context vector h_i^v can be calculated as:

$$h_i^v = \text{ReLU}([\overleftarrow{h}_i; \overrightarrow{h}_i] + W_f f_i) \quad (1)$$

where W_f is a projection matrix.

2.2 Initial Topic Description Generator

Given the visual context vector extracted from the image sequence, we first learn to generate the topic description. In practice, all visual context vectors h_i^v are concatenated and then fed into the initial topic description generator that employs a gated recurrent unit (GRU) decoder.

The output p_{topic}^{init} of this decoder is a sequence of probability distribution over the whole topic vocabulary \mathbb{V}_t . The training loss of initial topic description generator is the cross-entropy $L_{topic(mle)}^{init}$ between the generated description p_{topic}^{init} and the ground-truth topic description p_{topic} .

Note that at each time step, the decoder produces the a hidden state h_i^t . Once the last topic hidden state h_M^t is obtained, we concatenate all topic hidden states $h^t = [h_1^t, \dots, h_M^t]$, $M > 1$ as the topic memory, which are fed into the story generation module.

2.3 Initial Story Generator with Co-attention Network

The initial story generator is responsible for generating the story with the guidance of the topic description constructed by the initial topic description generator.

Co-attention Encoding In order to combine both visual information and topic information for story generation, we adopt a co-attention mechanism (Jing et al., 2018) for context information encoding. Specifically, given visual context vectors h^v and topic vectors h^t , the affinity matrix C is calculated by

$$C = \tanh(h^{tT} W_b h^v) \quad (2)$$

where W_b is the weight parameter. After calculating this matrix, we compute attentions weights over the visual context vectors and the topic vectors via the following operations:

$$\begin{aligned} H^v &= \tanh(W_v h^v + (W_t h^t) C) \\ H^t &= \tanh(W_t h^t + (W_v h^v) C^T) \\ a^v &= \text{softmax}(w_{hv}^T H^v) \\ a^t &= \text{softmax}(w_{ht}^T H^t) \end{aligned} \quad (3)$$

where W_v , W_t , w_{hv}^T , w_{ht}^T are the weight parameters. Based on the attention weights, the visual and semantic attentions are calculated as the weighted sum of the visual context vectors and the topic vectors:

$$a_{att}^v = \sum_{n=1}^N a_n^v h_n^v, a_{att}^t = \sum_{m=1}^M a_m^t h_m^t \quad (4)$$

At last, we concatenate the visual and semantic attentions as $[a_{att}^v; a_{att}^t]$, and then use a fully connected layer W_{fc} to obtain the joint context vector:

$$j = W_{fc}[a_{att}^v; a_{att}^t] \quad (5)$$

Story Decoding In the story decoding stage, each joint context vector j_i is fed into a GRU decoder to generate a sub-story sentence y_i for the corresponding image. Formally, the generation process can be written as:

$$s_t^i = \text{GRU}(s_{t-1}^i, [w_{t-1}^i, j_i]) \quad (6)$$

$$p(w_t^i | w_{1:t-1}^i) = \text{softmax}(\text{MLP}(s_t^i)) \quad (7)$$

where s_t^i denotes the t -th hidden state of i -th GRU. We concatenate the previous word token w_{t-1}^i and the context vector j_i as the input at each step. The output p is a probability distribution over the whole story vocabulary \mathbb{V}_s .

Loss Function for Training We define two different loss functions including cross-entropy (MLE) and reinforce (RL). MLE loss is show in Equation 8:

$$L_{story(mle)}^{init}(\theta_1) = - \sum_{t=1}^T \log(p_{\theta_1}(y_t^* | y_1^*, \dots, y_{t-1}^*)) \quad (8)$$

where θ_1 is the parameter of story generator; y^* is the ground-truth story and y_t^* denotes the t -th word in y^* .

Recently, reinforcement learning has shown effectiveness for training text generation model via introducing automatic metrics (e.g., METEOR) to guide the training process (Wang et al., 2018b). We also explore the RL-based approach to train our generator. The reinforcement learning (RL) loss can be written as:

$$L_{story(rl)}^{init}(\theta_1) = -E_{y \sim p_{\theta_1}}(r(y; y^*) - b)^2 \quad (9)$$

where r is a sentence-level metric for the sampled sentence y and the ground-truth y^* ; b is the baseline which can be an arbitrary function but a linear layer in our experiments for simply. To stabilize the RL training process, a simple way is to linearly combine MLE and RL objectives as follows (Wu et al., 2018):

$$L_{story(com)}^{init} = \alpha L_{story(rl)}^{init} + (1 - \alpha) L_{story(mle)}^{init} \quad (10)$$

where hyper-parameter α is employed to control the trade-off between MLE and RL objectives.

In the initial stage, a combined loss function of $L_{story(com)}^{init}$ and L_{topic}^{init} is computed through:

$$L^{init} = \lambda_1 L_{story(com)}^{init} + (1 - \lambda_1) L_{topic(mle)}^{init} \quad (11)$$

where hyper-parameter λ_1 is employed to balance these losses.

2.4 Iterative Updating Module

Considering that the generated story would also be helpful for the generation of topic description, we design an iterative updating module for the two agents to interact with each other and update iteratively. In IU module, we generate the topic description from the previously generated story, and then use such topic information to further guide story generation. To distinguish the two agents from those of initial version, we call them the IU version.

IU Topic Description Generator We envisage that the generated story is able to provide more accurate information for topic description generation than visual information. Therefore, instead of using visual information as input, the IU version of topic description generator takes the generated story as input. Specifically, the last hidden states s_{iter} of the IU story generator is used as input. Note that the IU topic description generator is initialized as its initial version and keeps training with the same objective.

IU Story Generator The IU story generator shares the structure and parameters with the initial story generator. It takes both topic vector and visual context vector as input. In the decoding process, the story y is the concatenation of the sub-stories y_i generated by IU story generator, and the last hidden states $s_{iter} = [s_{last}^1, \dots, s_{last}^N]$ of the IU story generator will be passed to IU topic description generator for iterative updating.

Loss Function for Training The training losses of the IU story generator are similar to Eq.(8,9,10), including the combination loss $L_{story(com)}^{iter}$ of $L_{story(mle)}^{iter}$ and $L_{story(rl)}^{iter}$.

At each iteration stage n , the IU module loss L_n^{iter} is the weighted sum of IU topic description generation loss $L_{topic(mle)}^{iter}$ and IU story generation loss $L_{story(com)}^{iter}$:

$$L_n^{iter} = \lambda_2 L_{story(com)}^{iter} + (1 - \lambda_2) L_{topic(mle)}^{iter} \quad (12)$$

where hyper-parameter λ_2 is employed to balance these losses.

2.5 Multi-Agent Training

The IU topic description generator and IU story generator communicate with each other iteratively in the IU module until it reaches the given iteration number N_{iter} . The loss for IU module is:

$$L^{iter} = \sum_{n=1}^{N_{iter}} L_n^{iter} \quad (13)$$

Therefore, to train the whole multi-agent learning framework, we introduce a combined loss L which consists of the initial loss L^{init} and IU module loss L^{iter} :

$$L = \beta L^{init} + (1 - \beta) L^{iter} \quad (14)$$

where β is a hyper-parameter to balance these losses. During training, our goal is minimizing L using stochastic gradient descent.

3 Experiments

3.1 Datasets

This paper utilizes VIST dataset (Huang et al., 2016) for experiments. We use the same split settings as previous work (Wang et al., 2018b), including 40,098/ 4,988/ 5,050 samples for training, validation and testing, respectively. Each sample (album) contains five images and a story consisting of five sentences. We use the title of each album as the ground-truth topic description.

3.2 Implementation Details

We use the pre-trained ResNet-152 (He et al., 2016) to extract image features. The vocabulary of story and topic includes words appearing no less than three times in the training set (i.e., story and title). We adopt GRU models for both visual encoder and other decoders, and the hidden size of them is 512. Expect the encoder is bidirectional, the other decoders are unidirectional. The batch size is set as 64 during the training. We use Adam (Kingma and Ba, 2015) with the initial learning rate of 0.0002.

We first pre-train the initial topic description generator using MLE. Then we pre-train both the topic description generator and the story generator jointly using MLE. The number of iteration N_{iter} is set to 2, the weight of RL is $\alpha = 0$, and hyper-parameters in loss optimization are set as $\lambda_1 = 0.7$, $\lambda_2 = 0.7$ and $\beta = 0.3$, which are selected based on validation set. After warm-up pre-training, α and learning-rate are set to 0.8 and 0.00002 to fine-tune using RL. Here we use METEOR scores as the reward. We select the best model which achieves the highest METEOR score on the validation set. The reason is that METEOR is proved to correlate better with human judgment than CIDEr-D in the small references case and superior to BLEU@N all the time (Vedantam et al., 2015; Wang et al., 2018a). During the test stage, we generate the stories by performing a beam-search with a beam size of 3.

3.3 Models for Comparison

We compare our proposed methods with several baselines for visual storytelling as follows:

seq2seq (Huang et al., 2016): It generates caption for each single model via classic sequence-to-sequence model and concatenate all captions to form the final story.

h-attn-rank (Yu et al., 2017a): On top of the classic sequence-to-sequence model, it adds an additional RNN to select photos for story generation.

HPSR (Wang et al., 2019a): It introduces an additional RNN stacked on the RNN-based photo encoder to detect the scene change. Information from both RNNs are fed into an RNN for story generation.

Methods	B-1	B-2	B-3	B-4	R-L	C	M
MLE							
seq2seq (Huang et al., 2016)	–	–	–	3.5	–	6.8	31.4
h-attn-rank (Yu et al., 2017a)	–	–	21.0	–	29.5	7.5	34.1
HPSR (Wang et al., 2019a)	61.9	37.8	21.5	12.2	31.2	8.0	34.4
VST (MLE)	62.3	38.0	21.8	12.7	29.7	7.8	34.3
TAVST w/o IU (MLE)	63.1	38.6	22.9	14.0	29.7	8.5	35.1
TAVST (MLE)	63.6	39.3	23.4	14.2	30.3	8.7	35.4
RL							
AREL (Wang et al., 2018b)	63.7	39.0	23.1	14.0	29.6	9.5	35.0
HSRL (Huang et al., 2019)	-	-	-	12.3	30.8	10.7	35.2
TAVST w/o IU (RL)	63.5	39.2	23.2	14.3	30.0	8.7	35.3
TAVST (RL)	64.2	39.6	23.7	14.6	31.0	9.2	35.7

Table 1: Overall performance of story generation on VIST dataset for different models in terms of BLEU (B), METEOR (M), ROUGE-L (R-L), and CIDEr-D (C). **Bolded** numbers are the best performance in each column.

AREL (Wang et al., 2018b): It is based on the framework of reinforcement learning and the generation of a single word is treated as the policy. The reward model learns the reward function from human demonstrations.

HSRL (Huang et al., 2019): It is based on the framework of hierarchical reinforcement learning. The higher level agent is responsible for generating a local concept for each image as the guidance to the lower level agent for sentence generation .

VST: This is the baseline version of our model without using topic information as guidance.

TAVST w/o IU: This is our proposed *TAVST* method without IU module, which only equipped with initial topic description generator.

TAVST: This is our full model. *TAVST (MLE)* is trained using MLE loss, while *TAVST (RL)* is trained via RL loss.

3.4 Automatic Evaluation Results

We evaluate our model on two generation tasks i.e., story generation and topic description generation, in terms of four automatic metrics: BLEU (Papineni et al., 2002), ROUGE-L (Lin and Och, 2004), METEOR (Banerjee and Lavie, 2005), and CIDEr (Vedantam et al., 2015).

Story Generation The overall experimental results are shown in Table 1. *TAVST (MLE)* outperforms all of the baseline models trained with MLE. This confirms the effectiveness of topic information for generating better stories. Noticeably, compared with the RL-based models, our *TAVST (MLE)* has already achieved a competitive performance and outperforms other RL models (i.e., *AREL* and *HSRL*) in terms of METEOR and BLEU@[2-4] metrics. After equipped with RL, our *TAVST (RL)* model is able to further improve the performance, outperforming the two RL models in terms of all metrics except CIDEr-D. Our full model *TAVST* (both MLE and RL versions) outperforms *TAVST w/o IU*, which directly demonstrates the effectiveness of the IU module. *TAVST w/o IU* achieves better performance than *VST*, which proves that topic description generator can provide guidance for story generation.

Topic Description Generation Table 2 shows the results of topic description generation. *TAVST* achieves higher performance compared to *TAVST w/o IU*, indicating that the generated story is able to provide assistance

Methods	B-1	B-2	R	C	M
TAVST w/o IU	5.1	2.2	12.6	10.9	4.8
TAVST	6.0	2.7	13.4	12.1	5.6

Table 2: Performance of topic description generation.

In general, the description generator obtains low scores in terms of automatic metrics. observations on the dataset reveal that the length of titles for each album is rela-

tively short, ranging from 2 to 6 words mostly. Given such a short reference, it is difficult for models to obtain high scores in terms of automatic metrics. We further look into the generated descriptions and some of them are actually semantically correct. For example, the reference is “*happy birthday party at my home*” and the generated topic description is “*the birthday gathering*”. Another example is that the reference is “*family feast*” and the generated topic is “*dinner party*”. We believe such kind of topic description with similar meaning can still provide positive guidance for the story generator.

3.5 Human Evaluation

We perform two kinds of human evaluation through Amazon Mechanical Turk (AMT), namely the Turing test and the pairwise comparison. Since we only find one previous work (Wang et al., 2018b) which published the sampled results of their model, we chose it for comparison. In specific, we re-collect human labels for their sample results and stories generated by our models on the same sub-set of albums. A total of 150 stories (750 images) are used, and each of them is evaluated by 3 human evaluators.

Turing Test For the Turing test, we design a survey (as shown in Appendices) that contains an image stream, a generated story by our *TAVST* model and a story written by a human. Evaluators are required to choose the story that is more likely written by a human. The experimental result (Table 3.5) shows that 47.8% of evaluators think the stories generated by our model are written by a human (v.s 38.4% win rate from *AREL*).

Method	Win	Lose	Unsure
VST	22.4%	71.7%	5.9%
AREL	38.4%	54.2%	7.4%
TAVST	42.7%	52.2%	5.1%

Table 3: Human evaluation results on Turing test.

Choice (%)	TAVST vs VST			TAVST vs AREL			TAVST vs GT		
	TAVST	VST	Tie	TAVST	AREL	Tie	TAVST	GT	Tie
Relevance	56.00	42.00	2.00	51.56	47.56	0.89	41.44	48.20	10.36
Expression	60.00	38.44	1.56	46.22	53.33	0.44	42.12	56.31	1.58
Topic Consistency	66.44	31.11	2.44	58.67	38.67	2.67	43.24	46.85	9.91

Table 4: Human evaluation results on pairwise comparison.

Pairwise Comparison A good story for an image stream should have three significant factors: (1) Relevance: the story should be relevant to the image stream. (2) Expressiveness: the story should be concrete and coherent, and have a human-like language style. (3) Topic Consistency: the story should be consistent to the topic. We compare our method with three other methods (*VST*, *AREL*, and *ground-truth (GT)*) in terms of these three metrics. In this annotation task, AMT evaluators need to compare two given stories according to these three factors and choose which story is better in terms of a certain factor. Results are shown in Table 4. Our model performs better than the other two models in terms of relevance and topic consistency. The advantage of topic consistency is more promising. This proves that the topic description generator can help the story generation agent construct a more consistent story.

3.6 Further Analysis on Topic Consistency

We further evaluate the quality of the generated story in terms of topic consistency from the perspective of sentiment. Specifically, we employ a lexicon-based approach using a subjectivity lexicon (Wilson et al., 2005). We count the number of sentiment words in each sentence for the polarity evaluation. The score will be 1,0,-1 if a sentence is positive, neutral and negative, respectively. Based on the score for each sentence, two qualitative experiments are designed to measure the in-story sentiment consistency and topic-story sentiment consistency.

In-story Sentiment Consistency We argue that the sentiment of sentences in a story should be consistent given the album is related to a certain topic. For each story, we obtain a vector with 5 sentiment scores in correspondence to 5 sentences. We then calculate the standard deviation for the vector to represent the divergence score of a story. For each model, we average the divergence scores of all stories

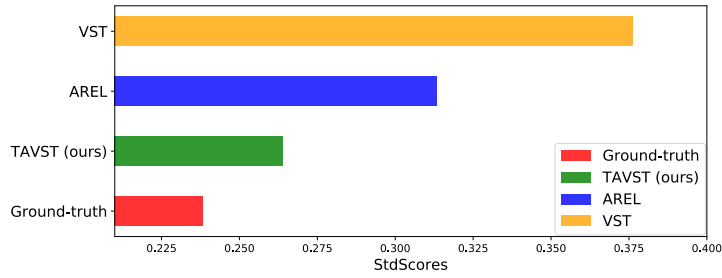


Figure 3: The comparison of in-story sentiment standard deviations among different methods.






						Topic
Ground-truth	i went to the award ceremony yesterday . there were a lot of people there . everyone received an award for their effort . they had a great time . i really enjoyed being there . some of the soldiers started singing .					graduation ceremony
VST	we were a lot of people there . he was very happy to be there . he gave a speech to the audience . he was very happy to be there . we had a great time at the end of the day .					---
AREL	we had a lot of fun at the meeting . he was very proud of his accomplishments . the president of the class gave a speech to the audience . the president of the school was very proud of his accomplishments . at the end of the ceremony , the family was happy to be together .					---
TAVST	yesterday was the day of the military graduation ceremony . the speaker gave a speech to the audience . they were very happy to be there . the soldier was very proud of his award . we had a great time at the end of the day .					the army ceremony

Figure 4: Story samples generated by different models and the ground truth. The words in red, blue and yellow color represent the topic, subject, and emotion, respectively. (Best viewed in color)

generated as its final score. Figure 3 presents the results from different models. Results illustrate that our method can generate stories with higher in-story sentiment consistency.

Topic-Story Sentiment Consistency Considering that albums related to some events might express a tendency to a certain polarity. For example, the sentiment of stories about *new year's eve* are more likely to be positive while the sentiment of stories about *breaking up* are more likely to be negative. We enumerate albums with different event types to see if the model has the ability to generate stories with sentiment consistent with the type of events. For each story, we add all the sentiment scores of sentences as its final score. The higher score a story obtain, the more positive the story is. Four types of events are considered. Results are shown in Table 5. In general, all automatic models tend to generate stories with higher sentiment scores compared to human-written stories. This is because a large portion of albums in the dataset are related to positive events. Both *VST* and *AREL* generate stories with similar sentiment scores for both types of events. This indicates that they are not able to distinguish positive and negative events. With the guidance of topic description, our model *TAVST* is able to distinguish events with different sentiment tendency.

Method	Positive		Negative	
	new year's eve	sporting event	breaking up	car accident
GT	1.20	1.25	0.48	0.38
TAVST	1.31	1.17	0.53	0.42
AREL	1.22	1.15	0.95	0.81
VST	1.65	1.82	1.66	1.84

Table 5: Sentiment scores corresponding to different types of events. Note that higher score indicates more positive polarity.

3.7 Case Study

Figure 4 shows an example of the ground-truth story and stories generated automatically by different models. The words in red, blue and yellow color represent the topic, subject, and emotion, respectively. Our model shows promising results according to topic consistency, which further confirms that our model can extract appropriate topic which serves as the guidance of generating a topic-consistent story.

4 Related work

This paper is related to the fields of image captioning, visual storytelling and multi-task learning.

Image Captioning In early works (Yang et al., 2011; Elliott and Keller, 2013), image captioning is treated as a ranking problem, which is based on retrieval models to identify similar captions from the database. Later, the end-to-end frameworks based on the CNN and RNN are adopted by researchers (Xu et al., 2015; Karpathy and Fei-Fei, 2015; Vinyals et al., 2017; Dai et al., 2017). Such works focus on the literal description of image content, while the generated texts is limited in a single sentence.

Visual Storytelling Visual storytelling is the task of generating a narrative paragraph for an image stream. Huang et al. (2016) introduce the first dataset (VIST) for visual storytelling, and establishes some baseline approaches. An attention-based RNN with a skip gated recurrent unit (Liu et al., 2017) is designed to maintain longer range information. In addition, Yu et al. (2017a) design a hierarchically-attentive RNN structure. More recently, Wang et al. (2018a) and Wang et al. (2018b) propose to utilize reinforcement learning frameworks for this task. Wang et al. (2020) propose to translate images to graph-based semantic representations to benefit representing images, while it is not fair to compare with our method because they introduce information from other datasets. The most similar work to ours is from Huang et al. (2019). They propose to generate a local semantic concept for each image in the sequence and generate a sentence for each image using a semantic compositional network in a fashion of hierarchical reinforcement learning. Although both of us consider topic information to facilitate the story generation. Our model is different from three aspects. First, the concepts of topic are different. We treat topic as the global semantic context of the album while topic represents local semantic information in their case adhering to each single image. Second, our modeling topic is more interpretable. We generate topic description directly instead of producing latent representation and this provides more insights for further improving the performance. Third, the communication framework is compatible with any RL based training methods. Experiment results also show that with RL, our framework can outperform their model.

Multi-Task Learning Collobert and Weston (2008) first propose a method for processing NLP tasks in a deep learning framework using multi-task learning. Jing et al. (2018) build a multi-task learning framework which jointly performs the prediction of tags and the generation of paragraphs. These multi-task learning methods share a certain network structure, and at the output layer design a specific network structure for different tasks, improving the performance of different tasks. However, unlike these multi-task learning methods, we use another multi-agent method (Sukhbaatar et al., 2016; Wang et al., 2019b). In this work, we define two kinds of agents for two generation tasks which can interact and share useful information. We also notice that in other areas, there are also some works (Xing et al., 2017; Wang et al., 2019c) consider incorporating topic information.

5 Conclusions and Future Work

In this paper, we introduce a topic-aware visual storytelling task, which identifies the global semantic context of a given image sequence and then generate the story with the help of such topic information. We propose a multi-agent communication framework that combines two generation tasks namely topic description generation and story generation effectively. In future, we will explore to model topic generation as a keyword extraction task.

Acknowledgements

This work is partially supported by National Natural Science Foundation of China (No. 71991471), Science and Technology Commission of Shanghai Municipality Grant (No.20dz1200600, No.18DZ1201000, 17JC1420200).

References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *ICCV*.
- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL Workshop*, pages 65–72.
- Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, pages 160–167.
- Bo Dai, Sanja Fidler, Raquel Urtasun, and Dahua Lin. 2017. Towards diverse and natural image descriptions via a conditional gan. In *CVPR*, pages 2970–2979.
- Desmond Elliott and Frank Keller. 2013. Image description using visual dependency representations. In *EMNLP*, pages 1292–1302.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*, pages 770–778.
- Ting-Hao Kenneth 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 *NAACL*, pages 1233–1239.
- Qiuyuan Huang, Zhe Gan, Asli Celikyilmaz, Dapeng Wu, Jianfeng Wang, and Xiaodong He. 2019. Hierarchically structured reinforcement learning for topically coherent visual story generation. In *AAAI*, pages 8465–8472.
- Baoyu Jing, Pengtao Xie, and Eric P. Xing. 2018. On the automatic generation of medical imaging reports. In *ACL*, pages 2577–2586.
- Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, pages 3128–3137.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Chin-Yew Lin and Franz Josef Och. 2004. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *ACL*, page 605.
- Yu Liu, Jianlong Fu, Tao Mei, and Chang Wen Chen. 2017. Let your photos talk: Generating narrative paragraph for photo stream via bidirectional attention recurrent neural networks. In *AAAI*, pages 1445–1452, February.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318.
- Zhou Ren, Xiaoyu Wang, Ning Zhang, Xutao Lv, and Li-Jia Li. 2017. Deep reinforcement learning-based image captioning with embedding reward. In *CVPR*, pages 290–298.
- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. Towards vqa models that can read. In *CVPR*.
- Sainbayar Sukhbaatar, Rob Fergus, et al. 2016. Learning multiagent communication with backpropagation. In *NIPS*, pages 2244–2252.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *CVPR*, pages 4566–4575.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2017. Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE transactions on pattern analysis and machine intelligence (PAMI)*, 39(4):652–663.
- Jing Wang, Jianlong Fu, Jinhui Tang, Zechao Li, and Tao Mei. 2018a. Show, reward and tell: Automatic generation of narrative paragraph from photo stream by adversarial training. In *AAAI*, pages 7396–7403, February.
- Xin Wang, Wenhui Chen, Yuan-Fang Wang, and William Yang Wang. 2018b. No Metrics Are Perfect: Adversarial Reward Learning for Visual Storytelling. In *ACL*, pages 899–909, April.
- Bairui Wang, Lin Ma, Wei Zhang, Wenhao Jiang, and Feng Zhang. 2019a. Hierarchical photo-scene encoder for album storytelling. In *AAAI*.

- Siyuan Wang, Zhongyu Wei, Zhihao Fan, Yang Liu, and Xuanjing Huang. 2019b. A multi-agent communication framework for question-worthy phrase extraction and question generation. In *AAAI*.
- Yue Wang, Jing Li, Hou Pong Chan, Irwin King, Michael R. Lyu, and Shuming Shi. 2019c. Topic-aware neural keyphrase generation for social media language. In *ACL*, pages 2516–2526.
- Ruize Wang, Zhongyu Wei, Piji Li, Qi Zhang, and Xuanjing Huang. 2020. Storytelling from an image stream using scene graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9185–9192.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *EMNLP*, pages 347–354.
- Lijun Wu, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2018. A study of reinforcement learning for neural machine translation. In *EMNLP*, pages 3612–3621.
- Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic aware neural response generation. In *AAAI*.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, pages 2048–2057.
- Yezhou Yang, Ching Lik Teo, Hal Daumé, and Yiannis Aloimonos. 2011. Corpus-guided sentence generation of natural images. In *EMNLP*, pages 444–454.
- Licheng Yu, Mohit Bansal, and Tamara Berg. 2017a. Hierarchically-attentive rnn for album summarization and storytelling. In *EMNLP*, pages 966–971.
- Zhou Yu, Jun Yu, Jianping Fan, and Dacheng Tao. 2017b. Multi-modal factorized bilinear pooling with co-attention learning for visual question answering. In *ICCV*, Oct.