

Commonsense Knowledge Aware Concept Selection for Diverse and Informative Visual Storytelling

Hong Chen^{1,3*}, Yifei Huang¹, Hiroya Takamura^{2,3}, Hideki Nakayama^{1,3}

¹The University of Tokyo, Japan

²Tokyo Institute of Technology, Japan

³National Institute of Advanced Industrial Science and Technology, Japan

{chen, nakayama}@nlab.ci.i.u-tokyo.ac.jp, hyf@iis.u-tokyo.ac.jp, takamura.hiroya@aist.go.jp

Abstract

Visual storytelling is a task of generating relevant and interesting stories for given image sequences. In this work we aim at increasing the diversity of the generated stories while preserving the informative content from the images. We propose to foster the diversity and informativeness of a generated story by using a concept selection module that suggests a set of concept candidates. Then, we utilize a large scale pre-trained model to convert concepts and images into full stories. To enrich the candidate concepts, a commonsense knowledge graph is created for each image sequence from which the concept candidates are proposed. To obtain appropriate concepts from the graph, we propose two novel modules that consider the correlation among candidate concepts and the image-concept correlation. Extensive automatic and human evaluation results demonstrate that our model can produce reasonable concepts. This enables our model to outperform the previous models by a large margin on the diversity and informativeness of the story, while retaining the relevance of the story to the image sequence.

Introduction

Telling a story based on a sequence of images is a natural task for humans and a fundamental problem for machine intelligence for various scenarios such as assisting the visually impaired people. Also known as visual storytelling (VST), the task has raised extensive research attention, since VST requires the model to not only understand the complex content within one image but also reason about the event across images as they occur and change. Since image sequences contain rich and diverse information, it is especially difficult for a model to tell a relevant story that is both informative of the image content and diverse in story style.

Most previous works on VST construct end-to-end frameworks (Yang et al. 2019; Wang et al. 2018; Jung et al. 2020; Yu, Bansal, and Berg 2017). However, although these methods can produce legitimate stories with high score in automatic metrics like BLEU (Papineni et al. 2002), it is shown that the stories tend to be monotonous which contains



Figure 1: Stories generated by the state-of-the-art model (Jung et al. 2020) and our proposed model using concept selection (red). The state-of-the-art model tends to generate similar stories (blue) for different inputs. Compared with it, our model can generate more informative and diverse stories.

limited lexical diversity and knowledge (Hsu et al. 2019a) (see the example in Figure 1). Recently, two-stage generation methods, also known as plan-write strategy, aroused much research attention in story generation tasks (Yao et al. 2019; Martin et al. 2017; Ammanabrolu et al. 2020). When adopted to the task of VST, Hsu et al. (2019a) shows that this strategy is capable of generating more diverse stories compared with end-to-end methods. However, their method directly generates concepts from the images using sequence-to-sequence models. Since the concept is selected from the full vocabulary, this kind of direct generation often produces concepts of low quality which affects the informativeness of the story.

In this work we aim to generate stories that are both diverse and informative for a given input image sequence. Taking the advantage of the previous two-stage models, we detect image concepts and construct concept graphs for proposing a set of concept candidates, and propose two

*With help from the AAAI Publications Committee.

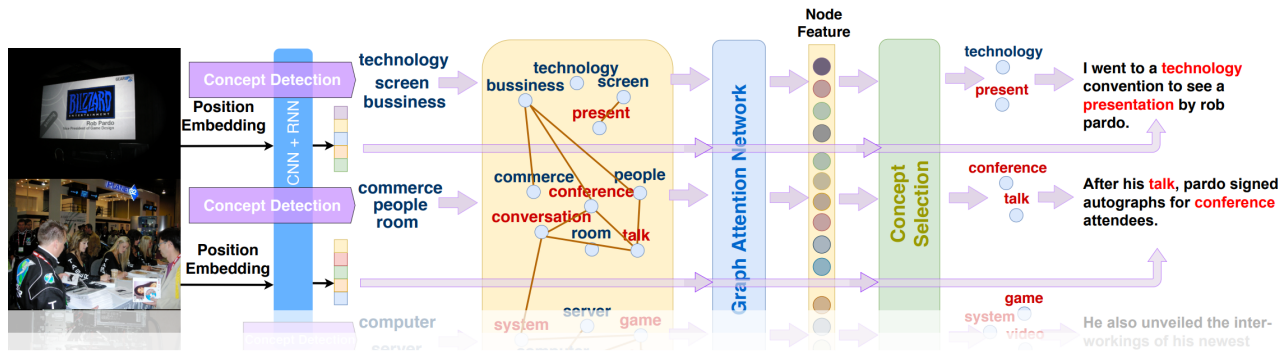


Figure 2: An overview of our visual storytelling model. The image features are obtained by a pretrained CNN combined with a bi-LSTM layer. The concepts are obtained from a concept detection model and enriched by ConceptNet (Liu and Singh 2004). These concepts from the nodes in a graph and are connected according to the relationship in the knowledge base. Initialized by the word embedding vector, the concept features are then updated by a Graph Attention Network. Our proposed concept selection module is then applied to select exact concept words using the image features and the concept features. Finally, both image features and concept words are used to generate a full story.

novel methods for better selecting the appropriate concept for the second generation stage. After detecting the concept in each input image, we first extend the concepts into a larger commonsense graph using ConceptNet (Liu and Singh 2004). This extension step increases the informativeness of generated stories. Since selecting appropriate candidates from the concept graph is critical for generating stories of good quality, a natural way is to use a graph attention network (Veličković et al. 2017) to refine the node features. Using Graph Attention Network can allow message passing along the graph structure, so that information of related concepts can be updated and integrated. This would allow us to get a better pool of concept candidates.

For selecting the most adequate concept from the candidates as the input to the second stage of our model, two novel modules are proposed in this work. The first one, named Sequential Selection Module (SSM), operates in a straightforward manner that uses an encoder-decoder for selecting concepts for each image. Differently from SSM, the second module called Maximal Clique Selection Module (MCSM) processes the concept graph as a whole. It learns a probability for each concept in the training phase, and during inference it finds a maximal clique using the Bron Kerbosch algorithm (Bron and Kerbosch 1973). The concepts within the clique are used for the next story generation step. Our experiments show that improved quality of concept selection can greatly help to increase the diversity of the generated stories while keeping the relevance with the input images.

The second stage of our model generates a story with the image features and the selected concepts. Other than using the same module for fair comparison with existing works, we also propose to modify the large scale pre-trained model BART (Lewis et al. 2019) to input the images and concepts and output the full stories.

We conduct extensive experiments on the public VIST dataset (Huang et al. 2016). Our experiments demonstrate that using our proposed concept selection modules, our generated stories can achieve better performance on both auto-

matic metric and multiple human evaluation metrics using the same generation module. When equipped with BART, the quality of the stories can be remarkably improved, with the generated story diversity similar to human writing.

In summary, our main contributions are listed as follows:

- We propose two novel modules SSM and MCSM to select concepts from the given candidates concepts under a plan-write two-stage visual storytelling system. The experiments show that our proposed methods can output more appropriate concepts than the previous work.
- We exploit modified BART as our story generation module to mitigate the problem caused by limited vocabulary and knowledge in the dataset. To the best of our knowledge, this is the first work to use a large scale pre-trained language model in a visual storytelling task.
- Large scale experiments using automatic metrics and human evaluation show that our model can outperform previous models by a large margin in both diversity and informativeness, while retaining the relevance and logicity as the previous work.

Related Work

Visual storytelling aims at generating stories for image sequences. Many existing methods focused on generating relevant and coherent stories (Hu et al. 2020; Wang et al. 2020; Chandu et al. 2019; Li et al. 2019; Hsu et al. 2019b). These works can be separated into two lines: one line is to construct end-to-end models to generate the stories directly from the images. The other line is to build a two-stage generation model that first outputs a mid-level abstraction and then generates the full story.

End-to-End Methods

Wang et al. (2018) proposed a visual storytelling framework which is widely used as a base model in the coming-up studies. This framework uses an end-to-end structure that first

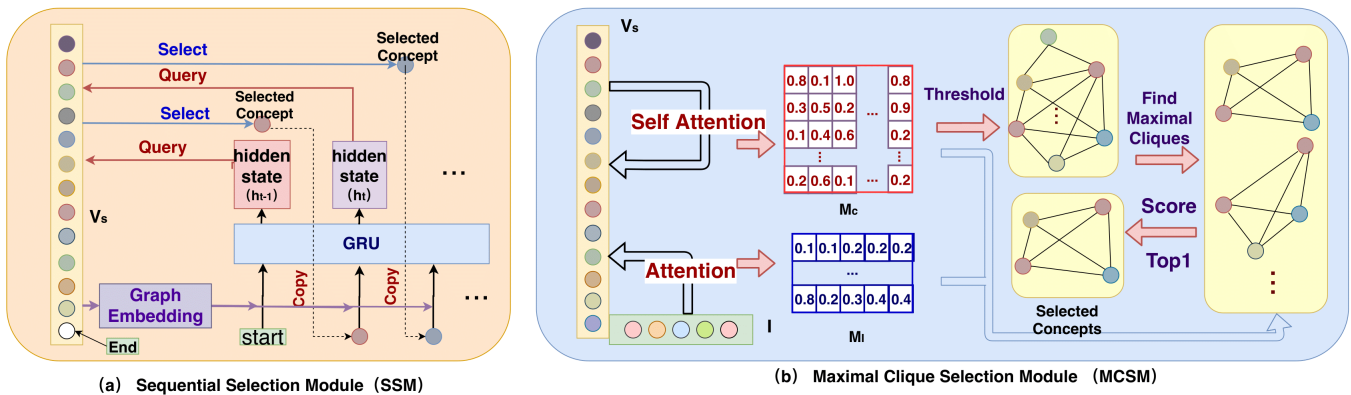


Figure 3: (a) Sequential Selection Module: this module sequentially use current hidden state as a query to select the concept from the commonsense graph. (b) Maximal Clique Selection Module: this module calculates concept-to-concept and image-to-concept correlation maps. The correlation maps can be viewed as a fully connected graph of concepts. We set a threshold to prune the edges. Then, a maximal clique algorithm is applied to find the maximal cliques in the remaining graph. Finally, those cliques are scored and the one with highest score will be selected and the concepts to be used in the next generation stage.

convert the image into features and then transfer its information to the adjacent images by a BiLSTM layer. Finally, a decoder decodes the features separately and merge the sentences into a story. While many succeeding works (Huang et al. 2019; Jung et al. 2020) can achieve high automatic scores, the story may not be interesting and informative (Hu et al. 2020) for human as they often contain repetitive texts and limited information. One of the main reasons for the low diversity and informativeness is that these model are trained end-to-end under the maximum likelihood estimation (Yang et al. 2019).

Two-Stage Methods

To alleviate the low diversity problem, Hsu et al. (2019a) proposed to generate several concepts before outputting the full stories. The discrete concept words can guide the decoder to produce more diverse stories. This plan-and-write strategy (Yao et al. 2019) can substantially increase the diversity of the stories. During the planning, to enhance the concepts that models can obtain, some researchers (Hsu et al. 2019a; Yang et al. 2019) introduce external commonsense knowledge database such as OpenIE (Angeli, Premkumar, and Manning 2015), Visual Genome (Krishna et al. 2017) or ConceptNet in the VST task. Their results show that using external knowledge base helps to generate more informative sentences. These works also show that the concepts are critical for the quality of generated stories because they can control the story flow.

In this work we aim to improve the concept selection for increasing the diversity and informativeness of VST. We propose two concept selection modules that carefully selects concept from a well-designed pool of concept candidates. The stories generated using our selected concepts thus become more diverse and informative. We further introduce to modify the pretrained model BART and use it to generate even better stories.

Method

Figure 2 depicts an overview of our proposed model. Given a sequence of N images as input, our model 1) encode image features, 2) construct a large commonsense graph for the images, 3) update concept feature in the graph, 4) select the concepts from the graph and 5), send concepts and image features into the decoder to output the full story. The details of each step are as follows:

Image Feature Encoding

We send the images into ResNet152 (He et al. 2016) to obtain image features $I = \{I_1, \dots, I_N\}$. Following Wang et al. (2018), a bidirectional GRU further encodes the high-level visual features to capture sequential information. The encoded feature for each image contains both the information of itself and the information from the adjacent images. Note that position embedding is applied before sending the image features into GRU to identify the order of the images.

Commonsense Graph Construction

To build our commonsense knowledge graph for image sequences, we need some seed concepts. Following Yang et al. (2019), we use clarifai¹ to obtain the top 10 seed concepts from each image. Each concept is used as a query to select relative commonsense concepts in the ConceptNet (Liu and Singh 2004). Since the number of the commonsense concepts is usually very large (> 500), we make several rules to filter some concepts which are less useful:

- Remove the commonsense concepts that appear less than 5 times in the whole training corpus.
- Remove the commonsense concepts that do not co-occur with the seed concepts either in the sentence or in the training corpus.
- If the concept number is still larger than K for one image. We simply randomly sample K words from it.

¹www.clarifai.com

After the filtering process, each image contains K concepts. Note that while different images can obtain the same concept in one image sequence, they can represent different semantic meanings in different positions of the story. Each concept forms a node in the commonsense graph. An undirected edge is established between concepts if they are related in ConceptNet. Also, a concept in one image will connect to the related concepts in the adjacent images to allow information flow between images. Like (Yang et al. 2019), we do not use the specific relation (*e.g.*, isA, has) between concepts. Till now, we build a graph which is a graph structure that is both connected within an image and between images.

Concept Features Update

The concept features are initialized with word embedding vectors. To incorporate the visual information into the concepts, we also connect the image feature to its corresponding concept features in the graph. These features are updated by a two-layer Graph Attention Network, which passes information between connected concepts and image using attention mechanism.

Concept Selection Module

We propose two methods to select concepts given the concept features and the image features.

To better formalize the procedure in the methods, we denote $c^{i,j}$ as the j -th concept of the i -th ($1 \leq i \leq N$) image. we let $\mathcal{C}_S = \{c_S^{1,1}, \dots, c_S^{N,K}\}$ and $\mathcal{C}_G = \{c_G^{1,1}, \dots\}$ denote the concepts set in the source candidate concepts and the full word set in the gold story, respectively. The target concepts are their intersection: $\mathcal{C}_T = \mathcal{C}_S \cap \mathcal{C}_G$.

Sequential Selection Module (SSM) One straightforward way of selecting concepts is to adopt an encoder-decoder model where we can forward the updated concept features into the encoder, and the decoder will output the selected concepts. Inspired by the Copy Mechanism (Gu et al. 2016), instead of generating a probability distribution with vocabulary size in each step, the SSM outputs are directly chosen from the inputs \mathcal{C}_S . As shown in Figure 3(a), we use a GRU (Cho et al. 2014) to first encode the concept embedding feature v_S^{t-1} and the hidden state into a new hidden state h^t . We then use h^t to query all the concepts in \mathcal{C}_S to get a probability p_S for each concept in the source set. Finally the concept with the highest probability is selected as the output concept, while its feature is directly copied for the generation of the next step:

$$\begin{aligned} h^t &= GRU(h^{t-1}, v_S^{t-1}) \\ p_S &= softmax((W_h h^t)^T W_c V_S) \\ c_S^t &= argmax(p_S) \end{aligned} \quad (1)$$

Here W_h and W_c are trainable projection matrices. The objective function is to maximize the probability score of the concepts which locate in \mathcal{C}_T .

$$\mathcal{L}_{ssm} = -\sum y_{S,T} \log(p_S), \quad (2)$$

where $y_{S,T}$ is an indicator of whether a concept in \mathcal{C}_S is in \mathcal{C}_T . The sequence selection step stops when the module generates <end> token. This <end> token is added to the set of candidate concepts with a uniform random initialized feature without any update during the training phase. The same procedure is done to the <start> token except that it is not involved in the candidates.

Maximal Clique Selection Module (MCSM) Different from SSM, this method aims to calculate the co-occurrence probability of all candidate concepts c_s in the graph. An illustration of MCSM is shown in Figure 3(b). In the beginning, we calculate self-attention to compute a correlation matrix $M_C \in (NK \times NK)$ which contains the correlation between each pair of nodes. We also calculate another correlation matrix for each image $M_I \in (N \times K)$ indicating the correlation between the concept embedding feature (v_S) and image features (I).

$$\begin{aligned} M_C &= \sigma(v_S^T W_a^T W_b v_S) \\ M_I &= \sigma(I^T W_c^T W_d v_S) \end{aligned} \quad (3)$$

Here, W_a, W_b, W_c, W_d is trainable weights, σ denotes sigmoid activation function. Intuitively, the concepts that appear in a gold story should own high correlations with each other, and the image should be highly related to the gold concepts to describe it. Thus, our target correlation maps can be written as follow:

$$\begin{aligned} \hat{M}_C[i, j] &= \begin{cases} 1, & c_i \in \mathcal{C}_T \wedge c_j \in \mathcal{C}_T \\ 0, & \text{otherwise} \end{cases} \\ \hat{M}_I[i, j] &= \begin{cases} 1, & c_j \in \mathcal{C}_T^i \\ 0, & \text{otherwise} \end{cases} \end{aligned} \quad (4)$$

Then, the objective is to minimize the difference between predicted and target correlation maps:

$$\mathcal{L}_{mcsm} = \lambda_1 \|M_C - \hat{M}_C\|_2^2 + \lambda_2 \|M_I - \hat{M}_I\|_2^2 \quad (5)$$

In testing phase, M_C can be viewed as a fully connected graph in which the edge weights correspond to the values in the matrix. Therefore, a low edge weight means less co-occurrence probability between two concepts. Based on this assumption, we set a threshold τ to remove the edges whose weight is less than τ . Then we apply Bron Kerbosch algorithm (Bron and Kerbosch 1973) to find all maximal cliques from the remaining sub-graph. Finally, we score each of them with equation 6 and select a clique with maximum score s . The output concepts are the nodes within the selected cliques.

$$\begin{aligned} s &= s_C + s_I \\ s_C &= \frac{1}{(\|\mathcal{C}_P\| - 1) \|\mathcal{C}_P\|} \sum_i \sum_{j \neq i} \log(M_C[i, j]) \\ s_I &= \frac{1}{\|\mathcal{C}_P\|} \sum_{i=1}^N \sum_{c_j \in \mathcal{C}_P^i} \log(M_I[i, j]). \end{aligned} \quad (6)$$

where \mathcal{C}_P denotes the concepts in a clique, and \mathcal{C}_P^i denotes the concept of the i -th image in the clique.

Concept to Story Module

The selected concepts are assigned to its corresponding image to generate the sentences. We tried two kinds of encoder-decoder to decode the story: 1) a simple encoder-decoder module that uses multi-head pooling to encode the concept embeddings, and decode the sentences with a RNN decoder. 2) a large scale encoder-decoder which both can encode the input and output the sentences.

RNN Decoder decodes sentences separately for each image and then concatenates into a story, while BART accepts all the images and concepts at once to output a full story.

RNN To get the concept feature from the concept words, we apply Multi-head Pooling (Liu and Lapata 2019) to calculate multi head self-attention score on the input concept embedding and do weighted summation on the concept embedding. Each image and corresponding concepts are decoded separately with the same decoder as Jung et al. (2020). The decoder accepts the image features I_i and the pooled concept feature v_i as input. Formally, the generation process can be written as:

$$h_i^t = RNN\left(h_i^{t-1}, [w_{t-1}^i; I_i; v_i]\right) \quad (7)$$

$$\pi^t = \text{softmax}\left(W_s h_i^t\right) \quad (8)$$

Here W_s is a projection matrix for outputting the probability distribution π^t on the full vocabulary. “;” denotes the channel-wise concatenation operation. Finally, all of the corresponding words are merged into a complete story.

BART Although there exists many large scale pretrained language models, none of them are used on VST in previous works. In this paper, we propose to use a modified version of BART as our base decoder model. BART is pretrained on a large dataset and their arbitrary noise function helps to mitigate the exposure bias. For inputting the image features into BART, We add a new randomly initialized projection matrix to transform the size of image features to fit BART’s input. For the concepts we simply use the pretrained BART word embedding. Since BART is powerful enough to generate a long story. Different with RNN decoder, we input N images and all concepts into the BART together and use a $\langle SEP \rangle$ token to separate the images and concepts, and to separate the concepts from different images. The transformed image feature and concept embedding are sent to BART and generate the final stories.

Experiment

Dataset

We conduct experiments on the widely used VIST dataset (Huang et al. 2016), which consists of 40101 samples for training 5050 test samples. Each sample contains 5 images and one corresponding gold story. For fair comparison, we follow the same experiment setting as (Jung et al. 2020) except that we set the vocabulary size to 28000. All models use the same fixed random seed.

Concept Detector

In this paper, we use the Concept Detection API produced by clarifai. This Detector can predict 11,000 concepts which is larger than any other detection model. This powerful pre-trained detector helps us to precisely find out the concepts inside the images, so that the word in the gold sentences can be easily involved in our knowledge enhanced candidate concepts.

Implementation Details

When training SSM, since we assume that there is no order relationship between the concepts in one image, so during the training phase, we randomly shuffle the target concept in one image. When training BART, we conduct a two-stage training: 1) freeze all BART parameters except for the image projection matrix. 2) finetune all parameters. We use Adam (Kingma and Ba 2014) optimizer with an initial learning rate of $4e-4$ in the first stage, then the learning rate is decreased to $4e-5$ in the fine tuning stage. Each stage is trained for 5 epochs. All the other parts of our model are trained with an Adam Optimizer with learning rate $4e-4$. During training, we follow Jung et al. (2020) to blind one of the images starting from the 50-th epoch and increase the blinding into two images from epoch 80. The training stops at epoch 100. Our model uses gold concepts extracted from gold stories to train the concept to story model. This step is similar to the common auto-regressive models that use the target token as the input to generate the next token. As has been discussed in the previous sections, this kind of generation often meets with the problem caused by the train-test discrepancy that we cannot see the gold concepts in the testing phase. To alleviate, a simple and effective way is to add noise to the inputs. In this work, we add two kinds of noise into the inputs in the story generation module: masking and random replacement. We mask 30% concepts and replace 20% of them into other similar words in training.

τ Search in Maximal Clique Selection Module

τ is set as the threshold in pruning edges for the maximal clique selection. Larger τ leads to fewer concepts that can be selected and will further result in the lack of imagination (diversity) in the generated story. However, smaller τ would lead to too many concepts that may mislead the model to generate irrelevant stories. To make a trade-off, we initialize τ as 0.3 and continual decreasing the number until Bron Kerbosch algorithm (Bron and Kerbosch 1973) can produce at least 5 candidate cliques that contains 7 to 15 candidate concepts in each clique.

Decoding Strategy

During model inference, usually beam search is used in decoding the sentences from the decoder (Wang et al. 2018; Jung et al. 2020; Yu, Bansal, and Berg 2017). However, it has been proved that beam search can result in bland, incoherent stories, or gets stuck in repetitive outputs (Holtzman et al. 2019). In this paper, to further improve diversity in generated stories, we apply nucleus sampling (Holtzman et al. 2019)

Method	Precision	Recall	F measure
Rand	2.68	2.45	2.56
C_Attn	30.38	43.37	35.86
I2C	31.32	20.75	24.96
SSM	40.43	40.30	40.36
MCSM	45.30	40.90	42.99

Table 1: Concept selection performance of different methods. The results show that our MCSM achieved the best f-score among all methods.

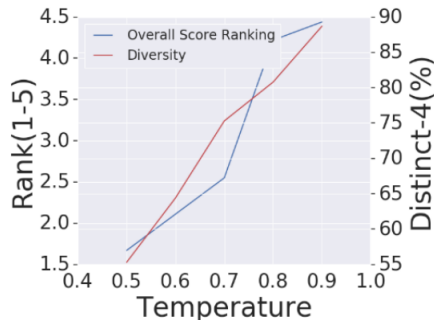


Figure 4: The y axis shows the ranking score (lower the better) and Distinct-4 score (higher the better) while we change the temperature in nucleus sampling with the model trained from scratch on VIST dataset. We ask the workers to rank the overall score for five stories generated by 5 different temperatures. As we can see, with the increasing of the temperature, the stories become more diverse, however, the quality of them become lower.

which outperforms top-k sampling method. However, nucleus sampling does not perform well without a well-trained language model. We show this by an experiment where we train our model with RNN as the story generation module on the VIST dataset and decode using nucleus sampling. Note the VIST dataset is relatively small compared with other natural language datasets like Book Corpus (Zhu et al. 2015). We conduct human evaluation and let the annotators to rank the overall story qualities (lower the better) generated with different softmax temperatures [0.5,0.6,0.7,0.8,0.9] for the same image sequence. We randomly pick 100 samples for this experiment. Figure 4 shows that with the temperature increases, the story quality would drop. Especially, when the temperature is 0.9 which is recommended in Holtzman et al. (2019), the generated sentences are incomprehensible and full of grammatical errors. This result motivates us to use a large scale pre-trained language model in our experiment.

Thus, for fair comparison with previous works, the same RNN with beam search is used as the story generation module to validate the effectiveness of our concept planning model. While, we use nucleus sampling (temperature=0.9, p=0.9) when BART (Lewis et al. 2019) is applied to demonstrate a higher upper bound of our model’s ability.

Method	Dist-2	Dist-3	Dist-4
INet \star	8.36	18.92	31.02
KS \star	10.84	22.90	36.51
KG-Story \dagger	18.73	38.65	57.22
Our(MCSM)	13.98	34.01	54.11
Image+BART \dagger	21.63	46.23	67.57
Our(MCSM)+BART \dagger	34.95	69.88	88.74
Gold	47.76	82.27	95.05

Table 2: Diversity of generated stories by different methods. Two-stage generation methods can produce more diverse stories. Using BART, we can achieve the diversity close the human writing, while achieving same level story quality in other aspect. \dagger denotes the story generation module is pre-trained with other dataset. \star denotes end-to-end methods.

Experiment on Concept Selection

We first test the ability to select appropriate concepts for different models. Since each image sequence in the test sample corresponds to 3 to 5 gold stories, we report the highest performance of the output selection result with respect to all gold stories in each sample. Similar as Keyphrase generation tasks, we apply precision, recall and f measure to evaluate the efficiency of concept selection methods. The precision is defined as the number of correct concepts selected over the target concepts, and recall is defined as the number of correct selections over all candidates.

We compare among several methods:

- **Rand**: A simple baseline where we randomly pick 3 concepts from the candidates for each image. On average, each image contains 2.65 gold concepts.
- **C_Attn**: We extract the attended concepts where the attention score is larger than a threshold from the model of Yang et al. (2019). This is an end-to-end model with sigmoid attention on concept words. We choose 0.8 as the threshold since this contributes the best f-score.
- **Image to concept(I2C)**: This is a straightforward version of concept selection where the concepts are directly generated from the images. We simply add a projection layer on each hidden state to predict the concept words from the vocabulary size of the concepts, which is very similar to the model of Hsu et al. (2019a).
- **SSM**: Our proposal which uses a copy mechanism in each step of selection.
- **MCSM**: Our proposal which calculates the correlation score for concept-concept and image-concept and uses maximal clique selection.

Qualitative results are shown in Table 1. We can see that our proposed SSM and MCSM can achieve significantly higher f-score than other methods. This helps our model to keep the story relevance to the input images while generating diverse stories.

Experiment on Visual Storytelling

Here we show the results of the visual storytelling. We use the following baselines for comparison:

Method	B-3	B-4	R	M	C
INet \star	23.5	14.4	29.7	35.3	9.5
KS \star	24.7	15.0	31.0	35.0	9.8
Rand+RNN	13.3	6.1	27.2	31.1	2.2
C_Attn+RNN	20.7	11.2	29.7	34.5	7.8
SSM+RNN	22.1	12.0	30.0	35.4	10.5
MCSM+RNN	23.1	13.0	30.7	36.1	11.0

Table 3: Automatic metric in story quality. We report BLEU (B), METEOR (M), ROUGH-L (R), and CIDEr (C) scores. The two-stage generation can achieve higher METEOR and CIDEr scores.

INet (Jung et al. 2020) This is a recent work which uses a “hide-and-tell” strategy to train an end-to-end model. In this method no concept is used.

KS (Yang et al. 2019) This method uses sigmoid attention to incorporate concept features into the model. We change the structure of the visual encoder and decoder the same as **INet** for fair comparison.

KG-Story \dagger (Hsu et al. 2019a) is a strong baseline that use two stage plan-write strategy and pretrain the decoder on RocStories Corpora (Mostafazadeh et al. 2017). \dagger indicates the model introduces external large-scale pretraining data.

Image+BART \dagger is an end-to-end baseline that uses BART on top of image features to directly generate the story. This baseline is one-stage that does not generate concepts.

We also change the concept selection module and story generation module in our model to validate the effectiveness of each component. Specifically, we compare: **Rand+RNN**, **C_Attn+RNN**, **SSM+RNN**, **MCSM+RNN**, and **MCSM+BART \dagger** .

Comparison on Diversity We first compare the ability of generating diverse stories of different models. Quantitative comparison is shown in Table 2. We report Distinct-n (Dist-n) scores (Li et al. 2015) that calculate the percentage of unique n-gram in all generated stories in the test dataset. Higher score means less inter-story repetition. From the table, two stage generation methods (KG-Story and ours) can achieve significantly higher diversity scores. Our MCSM can generate the most diverse stories among all the methods without using external pretrained models. When equipped with BART, we can even achieve diversity close to human writing. We show in the following that the increased diversity also improves the overall quality of the story.

Automatic Evaluation In Table 3, for comparing the quality of generated stories, we use automatic metrics BLEU (B) (Papineni et al. 2002), METEOR (M) (Banerjee and Lavie 2005), ROUGH-L (R) (Lin 2004), and CIDEr (C) (Vedantam, Lawrence Zitnick, and Parikh 2015). Note that it remains tricky for automatic scores to appropriately evaluate story qualities. Using concept, KS can achieve better performance than INet that does not use concept. From the comparison of the variants of our model, we can see that better concept selection can lead to better automatic scores.




Figure 5: We calculate the Pearson Correlation Coefficient for each criteria in human evaluation. R, I, L, O denotes Relevance, Informativeness, Logicality and Overall, respectively. We can see that Informativeness is almost independent to Relevance and Logicality, while is highly correlated to Overall score.

With reasonable concept selection, our SSM and MCSM can achieve highest METEOR and CIDEr scores.

Human Evaluation To better evaluate the quality of generated stories, we conduct human evaluation to compare pair-wise outputs with several models via the Amazon Mechanical Turk (AMT). We randomly sample 200 image sequences from the test set and generate stories using each model. For each sample pair, two annotators participate in the judgement and decide their preference on either story (or tie) in terms of Relevance, Informativeness, Logicality and Overall. **Relevance** evaluates how relevant the stories and the images are. **Informativeness** assesses how much information can be achieved in the generated stories, and this score from one side reflects the diversity of stories. **Logicality** evaluates the logical coherence in the stories. **Overall** is a subjective criterion that shows the preference of workers.

Table 4 shows the human evaluation result. Since there exists randomness in human evaluation, we compute the Cohen’s Kappa coefficient and found that all evaluations are in Moderate Agreement and Fair agreement, which indicates the evaluation result is reasonable as good inner agreement between evaluators is reached. We also conduct a Sign test to illustrate the significance of the evaluation difference: if p is below 0.05 it would indicate that the two compared models have a significant performance difference. From the comparison between MCSM and INet and the comparison between MCSM and KS, we can see that our two-stage planning method greatly outperforms the end-to-end models, especially in the informativeness score. The MCSM module also outperforms the SSM module, which indicates positive correlation between the quality of concept selection and the overall quality of generated stories. Finally, using BART with MCSM can help to achieve further informativeness and generate even better stories.

Importance of Informativeness in Story Quality

We calculate the Pearson Correlation Coefficient on four criteria in human evaluation. In Figure 5, R, I, L, O denotes Relevance, Informativeness, Logicality and Overall, respec-

Choices(%)	MCSM vs INet		MCSM vs KS		MCSM vs SSM		MCSM+BART [†] vs KS		MCSM+BART [†] vs MCSM	
	MCSM	INet	MCSM	KS	MCSM	SSM	MCSM+BART	KS	MCSM+BART	MCSM
Relevance	47.4	35.6	26.3	31.6	50.5	40.0	28.8	33.6	35.2	35.2
Informativeness	51.0*	31.6	46.3*	28.9	44.7	41.2	62.5**	18.8	58.8**	23.5
Logicity	35.5	34.3	34.2	29.0	32.9	42.3	35.3	33.3	40.2	37.5
Overall	55.0**	30.0	44.7	34.2	48.3	37.1	43.5**	23.0	47.0*	31.6

Table 4: Human evaluation. Numbers indicate the percentage of annotators believe that a model outperforms its opponent. Methods without (+BART) means using RNN as the story generation module. Cohen’s Kappa coefficients (κ) for all evaluations are in Moderate (0.4-0.6) or Fair (0.2-0.4) agreement, which ensures inter-annotator agreement. We also conduct a sign test to check the significance of the differences. The scores marked with * denotes $p < 0.05$ and ** indicates $p < 0.01$ in sign test.

tively. Ranged from -1 to 1, the Informativeness score has low correlation score with Logicity (0.011) and Relevance (0.08), while a high correlation score with Overall (0.47). This indicates that Informativeness is almost independent on Relevance and Logicity, but highly dependent on the Overall score. This suggests that humans tend to choose stories with more interesting information. This phenomenon proves the significance of informativeness and diversity in visual storytelling.

Case Study

We show a qualitative result of a random test sample in Figure 6. This is a hard example because the last three images are very similar and the objects in all images are hard to recognize. We can see that INet generates monotonous and even irrelevant sentences. KS can generate better sentences but still low in lexical diversity. For the stories generated by two-stage strategy with RNN (SSM+RNN, MCSM+RNN), we can see that the story follows the selected concepts and the stories seem more reasonable than that of end-to-end training methods. When using BART, we compare three methods that represent no concept selection (Image+BART), bad concept selection (Rand+BART) and ours concept selection (MCSM+BART). We can see that without using concepts or using randomly selected concepts, the generated stories are of low quality and to a certain extent irrelevant to the images. However, when guided by the selected concept, the story becomes vivid, relevant and logical.

Conclusion

In this work we exploit concept selection for improving the diversity and informativeness of stories generated from image sequences. By constructing a commonsense graph and two novel modules for concept selection, our proposed model outperforms all previous works in diversity by a large margin while still preserving the relevance and logical consistency on the VIST dataset. Our future direction aims to increase the relevance of the generated story by better leveraging the visual information.

Acknowledgements

We thank the anonymous reviewers for the useful comments. This paper was based on results obtained from a project, JPNP20006, commissioned by the New Energy and Industrial Technology Development Organization (NEDO)

INet	This is a picture of a bird. There was a lot of birds on the ground. There was a bird in the water. The birds were in the water. It was a great day for a swim.
KS	It was a nice day for a walk on the water. There were many birds in the sky. There were some pretty white ones. the water was very calm. There were also a lot of birds in the water.
SSM+RNN	Concepts: 1 day 2 ∅ 3 bird 4 water time lot 5 bird frozen back This is a picture of a day. The water was covered in snow. This is a picture of a bird. The bird spent was a lot of time in the water. The bird is frozen and he could not go back.
MCSM+RNN	Concepts: 1 day winter cold 2 spot 3 bird 4 water fish 5 lake It was a cold day for winter. We spotted some of the trees we saw. I saw a bird. The bird was fishing in some water. We had a great time at the lake.
Image+BART	We went to the beach. There was a lot of interesting things to see. There was a bird in the sky. We had a great time. It was very cold at night. Concepts: 1 home debating 2 dock alarm 3 bucket cap 4 oyster photo 5 slush americana
Rand+BART	There was a debate at home before heading out to the dock. Our alarm went off and we kept watch over the dock. Then they noticed the bird with its cap in the air and watched the bird fly away from the safety of the bucket. He jumped up to take a photo of this magnificent oyster. We ended our day by the water feeding with american birds. Concepts: 1 day winter cold 2 spot 3 bird 4 water fish 5 lake
MCSM+BART	A cold winter day at the lake. [male] spotted an animal on the lake and thought that it was a strange bird he could spot. The bird seemed peaceful, and seemed to be enjoying the cold. It then tried to fish in the water. [male] kept a close eye on the bird, but had to keep his distance from the lake.
Gold	Deer looking for food. And disappeared after hearing noises. Bird waiting patiently. Perching in the cold. About to take a flight.

Figure 6: The examples of generated stories by different methods. Our MCSM and SSM can generate better stories compared with other baselines that do not use BART. When using the pretrained BART, the concept selection with MCSM can produce vivid and informative story.

and was also supported by JSPS KAKENHI Grant Number JP19H04166.

References

Ammanabrolu, P.; Tien, E.; Cheung, W.; Luo, Z.; Ma, W.; Martin, L. J.; and Riedl, M. O. 2020. Story Realization:

- Expanding Plot Events into Sentences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 7375–7382.
- Angeli, G.; Premkumar, M. J. J.; and Manning, C. D. 2015. Leveraging linguistic structure for open domain information extraction. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 344–354.
- Banerjee, S.; and Lavie, A. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 65–72.
- Bron, C.; and Kerbosch, J. 1973. Algorithm 457: Finding All Cliques of an Undirected Graph. *Commun. ACM* 16(9): 575–577. ISSN 0001-0782. doi:10.1145/362342.362367. URL <https://doi.org/10.1145/362342.362367>.
- Chandu, K.; Prabhumoye, S.; Salakhutdinov, R.; and Black, A. W. 2019. “My Way of Telling a Story”: Persona based Grounded Story Generation. In *Proceedings of the Second Workshop on Storytelling*, 11–21.
- Cho, K.; Van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* .
- Gu, J.; Lu, Z.; Li, H.; and Li, V. O. 2016. Incorporating copying mechanism in sequence-to-sequence learning. *arXiv preprint arXiv:1603.06393* .
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- Holtzman, A.; Buys, J.; Du, L.; Forbes, M.; and Choi, Y. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751* .
- Hsu, C.-C.; Chen, Z.-Y.; Hsu, C.-Y.; Li, C.-C.; Lin, T.-Y.; Huang, T.-H.; and Ku, L.-W. 2019a. Knowledge-Enriched Visual Storytelling. *arXiv preprint arXiv:1912.01496* .
- Hsu, T.-Y.; Huang, C.-Y.; Hsu, Y.-C.; and Huang, T.-H. 2019b. Visual story post-editing. *arXiv preprint arXiv:1906.01764* .
- Hu, J.; Cheng, Y.; Gan, Z.; Liu, J.; Gao, J.; and Neubig, G. 2020. What Makes A Good Story? Designing Composite Rewards for Visual Storytelling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 7969–7976.
- Huang, Q.; Gan, Z.; Celikyilmaz, A.; Wu, D.; Wang, J.; and He, X. 2019. Hierarchically structured reinforcement learning for topically coherent visual story generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 8465–8472.
- Huang, T.-H.; Ferraro, F.; Mostafazadeh, N.; Misra, I.; Agrawal, A.; Devlin, J.; Girshick, R.; He, X.; Kohli, P.; Batra, D.; 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*, 1233–1239.
- Jung, Y.; Kim, D.; Woo, S.; Kim, K.; Kim, S.; and Kweon, I. S. 2020. Hide-and-Tell: Learning to Bridge Photo Streams for Visual Storytelling. *arXiv preprint arXiv:2002.00774* .
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* .
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision* 123(1): 32–73.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461* .
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, B. 2015. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055* .
- Li, N.; Liu, B.; Han, Z.; Liu, Y.-S.; and Fu, J. 2019. Emotion reinforced visual storytelling. In *Proceedings of the 2019 International Conference on Multimedia Retrieval*, 297–305.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Liu, H.; and Singh, P. 2004. ConceptNet—a practical commonsense reasoning tool-kit. *BT technology journal* 22(4): 211–226.
- Liu, Y.; and Lapata, M. 2019. Hierarchical transformers for multi-document summarization. *arXiv preprint arXiv:1905.13164* .
- Martin, L. J.; Ammanabrolu, P.; Wang, X.; Hancock, W.; Singh, S.; Harrison, B.; and Riedl, M. O. 2017. Event representations for automated story generation with deep neural nets. *arXiv preprint arXiv:1706.01331* .
- Mostafazadeh, N.; Roth, M.; Louis, A.; Chambers, N.; and Allen, J. 2017. Lsdsem 2017 shared task: The story cloze test. In *Proceedings of the 2nd Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics*, 46–51.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.
- Vedantam, R.; Lawrence Zitnick, C.; and Parikh, D. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4566–4575.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903* .
- Wang, R.; Wei, Z.; Li, P.; Zhang, Q.; and Huang, X. 2020. Storytelling from an Image Stream Using Scene Graphs. In

Proceedings of the AAAI Conference on Artificial Intelligence, 9185–9192.

Wang, X.; Chen, W.; Wang, Y.-F.; and Wang, W. Y. 2018. No metrics are perfect: Adversarial reward learning for visual storytelling. *arXiv preprint arXiv:1804.09160* .

Yang, P.; Luo, F.; Chen, P.; Li, L.; Yin, Z.; He, X.; and Sun, X. 2019. Knowledgeable Storyteller: A Commonsense-Driven Generative Model for Visual Storytelling. In *IJCAI*, 5356–5362.

Yao, L.; Peng, N.; Weischedel, R.; Knight, K.; Zhao, D.; and Yan, R. 2019. Plan-and-write: Towards better automatic storytelling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7378–7385.

Yu, L.; Bansal, M.; and Berg, T. L. 2017. Hierarchically-attentive rnn for album summarization and storytelling. *arXiv preprint arXiv:1708.02977* .

Zhu, Y.; Kiros, R.; Zemel, R.; Salakhutdinov, R.; Urtasun, R.; Torralba, A.; and Fidler, S. 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)*.