

Title	Relational Future Captioning Model for explaining likely collisions in daily tasks
Sub Title	
Author	神原, 元就(Kambara, Motonari) 杉浦, 孔明(Sugiura, Komei)
Publisher	慶應義塾大学AI・高度プログラミングコンソーシアム
Publication year	2023
Jtitle	AICカンファレンス予稿集 (2023.) ,p.50- 54
JaLC DOI	
Abstract	Domestic service robots that support daily tasks are a promising solution for elderly or disabled people. It is crucial for domestic service robots to explain the collision risk before they perform actions. In this paper, our aim is to generate a caption about a future event. We propose the Relational Future Captioning Model (RFCM), a crossmodal language generation model for the future captioning task. The RFCM has the Relational Self-Attention Encoder to extract the relationships between events more effectively than the conventional self-attention in transformers. We conducted comparison experiments, and the results show the RFCM outperforms a baseline method on two datasets.
Notes	会議名：AICカンファレンス2023 開催地：慶應義塾大学日吉キャンパス 日時：2023年3月4日 第3章既発表セッション要旨 既発表要旨-3
Genre	Conference Paper
URL	https://koara.lib.keio.ac.jp/xoonips/modules/xoonips/detail.php?koara_id=KO11003001-20230304-0050

慶應義塾大学学術情報リポジトリ(KOARA)に掲載されているコンテンツの著作権は、それぞれの著作者、学会または出版社/発行者に帰属し、その権利は著作権法によって保護されています。引用にあたっては、著作権法を遵守してご利用ください。

The copyrights of content available on the Keio Associated Repository of Academic resources (KOARA) belong to the respective authors, academic societies, or publishers/issuers, and these rights are protected by the Japanese Copyright Act. When quoting the content, please follow the Japanese copyright act.

RELATIONAL FUTURE CAPTIONING MODEL FOR EXPLAINING LIKELY COLLISIONS IN DAILY TASKS

Motonari Kambara and Komei Sugiura

Keio University, Japan

ABSTRACT

Domestic service robots that support daily tasks are a promising solution for elderly or disabled people. It is crucial for domestic service robots to explain the collision risk before they perform actions. In this paper, our aim is to generate a caption about a future event. We propose the Relational Future Captioning Model (RFCM), a crossmodal language generation model for the future captioning task. The RFCM has the Relational Self-Attention Encoder to extract the relationships between events more effectively than the conventional self-attention in transformers. We conducted comparison experiments, and the results show the RFCM outperforms a baseline method on two datasets.

Index Terms— Future Captioning, Domestic Service Robots, Relational Self-Attention

1. INTRODUCTION

Domestic service robots (DSRs) that naturally communicate with users to support household tasks are a promising solution for elderly or disabled people. DSRs are expected to perform most tasks autonomously, and so they could damage objects and themselves. It therefore would be useful if they could explain the potential risks associated with their actions through natural language. However, a DSR’s ability to generate natural language explanations is still insufficient.

Given this background, we focus on future captioning for daily tasks [1]. Fig. 1 shows a typical use case where a DSR puts a plastic bottle on a table. In this situation, it would be desirable to tell the user that “the hand may contact the white bottle, which may cause the bottle to further contact the apple next to it, causing the apple to fall.” This task is difficult in that models need to predict future events and generate captions. In fact, there is a big gap in the quality of the reference and the generated sentences by typical video captioning models, as shown in Sec. 5.

Although there are many existing video captioning models, those models cannot generate appropriate future captions. This is because many existing methods insufficiently model the temporal relationship between the visual features and sentences. The transformer self-attention [2] used in those meth-

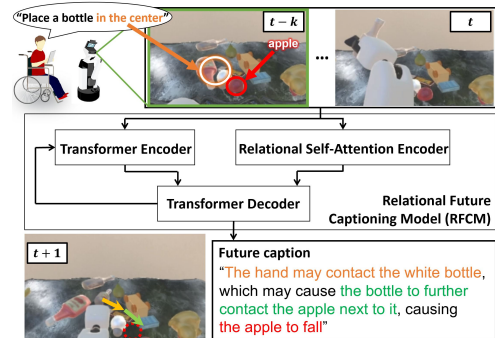


Fig. 1. Overview of RFCM: RFCM generates captions of future events from past events.

ods typically models the relationship between the visual features of the event and the sentence at time t . In addition, many video captioning methods are inappropriate for future captioning because those models use the captions after the next timestep to generate a caption for the next timestep.

In this paper, we propose the Relational Future Captioning Model (RFCM). The RFCM can generate captions that take into account the relationship between past events. This is because it has a source–target attention structure that generates appropriate captions for future events from the relationships between events. In this structure, the features derived from past clips are used as a source, and the features derived from both past clips and captions are used as the target.

Fig. 1 shows an overview of the RFCM, which consists of three modules: a Relational Self-Attention (RSA [3]) Encoder, transformer encoder, and transformer decoder. The difference with respect to existing methods is that our method includes the RSA Encoder to more effectively extract the relationships between events than the conventional self-attention in transformers. Our code is available at this URL¹.

The main contributions of this paper are as follows:

- We propose the RFCM, a crossmodal language generation model for the future captioning task.
- The RSA Encoder is introduced to extract the relationships between events more effectively than the conventional self-attention in transformers. The output of the RSA Encoder is used to compute a source–target attention to obtain fine-grained caption representations.

This material is based upon work supported by JSPS KAKENHI Grant Number 20H04269, JST Moonshot, and NEDO.

¹<https://github.com/keio-smilab22/RelationalFutureCaptioningModel>

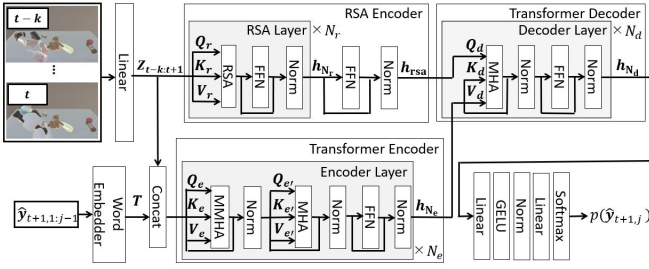


Fig. 2. The framework of RFCM. In the figure, Norm denotes layer normalization.

2. RELATED WORK

There have been many studies in the field of caption generation [1, 4–9]. This field includes image captioning and video captioning. [10] is a survey paper in the image captioning field. [11] is a survey paper in the video captioning field.

The field of video captioning can be divided into several subtasks such as a video captioning task, dense video captioning task, and future captioning task. The video captioning task involves generating a description about an input event. Many video captioning models have been proposed [4, 7, 12–14]. [13, 14] are typical examples of the numerous pre-training models. The dense video captioning task involves the generation of detailed captions. [5, 15, 16] are representative dense video captioning models. The future captioning is the task dealt with in this paper. [1, 17, 18] are representative future captioning models. [17] generates explanatory sentences of near future events using in-vehicle camera images and vehicle motion information observed from past to present. [1] and [17] used the CNN and LSTM encoder-decoder, while our method used the transformer encoder-decoder. [19] and [20] proposed event and video prediction models, respectively.

3. PROBLEM STATEMENT

Our target is to generate a sentence describing the next event in the video, which we call the future captioning task [1]. Fig. 1 shows a typical scene. In this task, the desired behavior is to generate a description by inferring the situation at time $t + 1$ given the clips up to time t .

The task is characterized by the following:

- **Input:** video clips up to time t . We define the term “clip” as a sequence of frames representing an event.
- **Output:** a sentence describing an event that is expected to occur at time $t + 1$.

We assume that pretraining on other tasks is not allowed because knowledge transfer is outside the scope of this study.

4. PROPOSED METHOD

Fig. 2 shows the framework of the RFCM. The input to the RFCM $\mathbf{x} \in \mathbb{R}^{(k+1) \times d_{in}}$ is defined as $\mathbf{x} = \{\mathbf{x}_{t-k}, \dots, \mathbf{x}_t\}$, where $\mathbf{x}_t \in \mathbb{R}^{d_{in}}$ and $k \in \mathbb{N}$ denote the clip at time t and the number of input events that occurred before t , respectively.

From \mathbf{x} , $\mathbf{Z}_{t-k:t+1} \in \mathbb{R}^{(k+2) \times d_{in}}$ is obtained as follows:

$$\mathbf{z}_\tau = \begin{cases} f_z(\mathbf{x}_\tau) & (\tau > t) \\ f_z(\mathbf{x}_\tau, \mathbf{x}_t) & (\tau \leq t), \end{cases}$$

$$\mathbf{Z}_{t-k:t+1} = \{\mathbf{z}_{t-k}; \dots; \mathbf{z}_{t+1}\},$$

where $\tau = t - k, \dots, t + 1$, $f_z(\cdot)$ denotes a linear transformation. $\mathbf{Z}_{t-k:t+1}$ contains information on an event sequence.

4.1. RSA Encoder

The RSA Encoder uses RSA [3], which uses correlations between query and key, and self-correlations among values as the relational context. Although simple source-target attention does not use self-correlations among values, it is important to consider their correlations because video is a time-series data.

The RSA Encoder consists of N_r RSA layers. In the first layer, positional encoding using trigonometric functions is performed for $\mathbf{Z}_{t-k:t+1}$ using the procedure described in [2].

Then, query $\mathbf{Q}_r \in \mathbb{R}^{d_{rsa}}$, key $\mathbf{K}_r \in \mathbb{R}^{(k+2) \times d_{rsa}}$, and value $\mathbf{V}_r \in \mathbb{R}^{(k+2) \times d_{rsa}}$ are obtained as follows:

$$\mathbf{Q}_r = \mathbf{z}_t, \mathbf{K}_r = \mathbf{V}_r = \mathbf{Z}_{t-k:t+1},$$

where d_{rsa} denotes the size of each layer. Next, the basic kernel $\varphi_p \in \mathbb{R}^{k+2}$ and relational kernel $\varphi_h \in \mathbb{R}^{k+2}$ are obtained as follows:

$$\begin{aligned} \varphi_p &= \mathbf{W}_p \mathbf{Q}_r, \\ \mathbf{Q}_{r'} &= \{\mathbf{Q}_r; \dots; \mathbf{Q}_r\} \in \mathbb{R}^{(k+2) \times d_{rsa}}, \\ \varphi_h &= \mathbf{W}_h f_{\text{flatten}}(\mathbf{Q}_{r'} \odot \mathbf{K}_r), \end{aligned}$$

where $f_{\text{flatten}}(\cdot)$ denotes flattening. Next, relational context $\Phi_g \in \mathbb{R}^{(k+2) \times d_{rsa}}$ obtained as follows:

$$\Phi_g = \mathbf{V}_r + \mathbf{W}_g \mathbf{V}_r^\top \mathbf{V}_r.$$

Finally, RSA $\varphi \in \mathbb{R}^{d_{rsa}}$ is obtained as follows:

$$\varphi = (\varphi_p + \varphi_h)^\top \Phi_g.$$

RSA φ contains information about the relationships between input events. To use φ as a latent feature at time t , we replace \mathbf{z}_t with φ . From this, we obtain $\mathbf{h}_r = \{\mathbf{z}_{t-k}; \dots; \mathbf{z}_{t-1}; \varphi; \mathbf{z}_{t+1}\} \in \mathbb{R}^{(k+2) \times d_{rsa}}$. The output of layer \mathbf{h}_{N_r} ($n_r = 1, \dots, N_r$) is obtained by applying a feedforward network (FFN) and layer normalization (LN) layers to \mathbf{h}_r .

The output of the encoder $\mathbf{h}_{rsa} \in \mathbb{R}^{(k+2) \times d_{rsa}}$ is given by $\mathbf{h}_{rsa} = f_{LN}(f_{FFN}(\mathbf{h}_{N_r}))$, where $f_{FFN}(\cdot)$ and $f_{LN}(\cdot)$ denote the FFN and LN layers, respectively.

4.2. Transformer encoder/decoder

The transformer encoder consists of N_e encoder layers. Each layer consists of a Masked Multi-Head Attention (MMHA), Multi-Head Attention (MHA), and FFN layers.

The input to the encoder is $\mathbf{h}_c = \{\mathbf{Z}_{t-k:t+1}; \mathbf{T}\} \in \mathbb{R}^{(k+I+2) \times d_{in}}$, where I denotes the maximum length of sentences. During training, $\mathbf{T} \in \mathbb{R}^{I \times d_{in}}$ denotes the text features obtained by embedding a reference sentence at time $t + 1$, \mathbf{y}_{t+1} with the BERT embedder [21]. During inference, \mathbf{T} denotes the text features for the generated words $\hat{\mathbf{y}}_{t+1,1:j-1}$ when the j -th word is generated.

The input to the n_e+1 -th layer is \mathbf{h}_{n_e} ($n_e = 0, \dots, N_e-1$). We set $\mathbf{h}_0 = \mathbf{h}_c$. In the MMHA layer, query $\mathbf{Q}_e \in \mathbb{R}^{(k+I+2) \times d_e}$, key $\mathbf{K}_e \in \mathbb{R}^{(k+I+2) \times d_e}$, and value $\mathbf{V}_e \in \mathbb{R}^{(k+I+2) \times d_e}$ ($d_e = d_{enc}/N_h$) are obtained as $\mathbf{Q}_e = \mathbf{W}_q^{(e)} \mathbf{h}_{n_e}$, $\mathbf{K}_e = \mathbf{W}_k^{(e)} \mathbf{h}_{n_e}$, and $\mathbf{V}_e = \mathbf{W}_v^{(e)} \mathbf{h}_{n_e}$ ($e = 1, \dots, N_h$), where

N_h and d_{enc} denote the number of attention heads and the size of each layer, respectively. During training, in the MMHA layer, we mask the m -th and subsequent word tokens to prevent the encoder from using the information of the words in the m -th word prediction (teacher forcing). To obtain attention $\mathbf{A}_{\text{MMHA}} \in \mathbb{R}^{(k+I+2) \times d_{\text{enc}}}$ from $\mathbf{Q}_e, \mathbf{K}_e$, and \mathbf{V}_e , we used the computation shown in [2]. Then, in the MHA layer, attention \mathbf{A}_{MHA} is obtained using the same computation used in the MMHA layer, where $\mathbf{h}_{\text{enc}}, \mathbf{Q}_e, \mathbf{K}_e$, and \mathbf{V}_e are replaced by $\mathbf{A}_{\text{MMHA}}, \mathbf{Q}_{e'}, \mathbf{K}_{e'}$, and $\mathbf{V}_{e'}$, respectively. The output of the $n_e + 1$ -th encoder layer, $\mathbf{h}_{n_e+1} \in \mathbb{R}^{(k+I+2) \times d_{\text{enc}}}$, is given by $\mathbf{h}_{n_e+1} = f_{\text{LN}}(f_{\text{FFN}}(\mathbf{A}_{\text{MHA}}))$. The output of the encoder is \mathbf{h}_{N_e} .

The transformer decoder takes as inputs \mathbf{h}_{rsa} and \mathbf{h}_{N_e} . This module consists of N_d decoder layers. The structure of each layer is similar to that of the MHA and FFN layers in the transformer encoder. However, there is a difference in that each layer has the source–target attention structure. Therefore, the query is created based on \mathbf{h}_{rsa} , whereas the key and value are created based on \mathbf{h}_{N_e} . The output is $\mathbf{h}_{N_d} \in \mathbb{R}^{(k+I+2) \times d_{\text{dec}}}$, where d_{dec} denotes the size of each layer.

Finally, the prediction probability of the generated word $p(\hat{\mathbf{y}}_{t+1,j}) \in \mathbb{R}^{N_v}$ is given by $p(\hat{\mathbf{y}}_{t+1,j}) = \text{softmax}(f_{\text{gen}}(\mathbf{h}_{N_d}))$, where N_v denotes a vocabulary size and $f_{\text{gen}}(\cdot)$ denotes the computation by the first fully connected, GELU, LN, and final fully connected layers.

The global loss function \mathcal{L} is defined as:

$$\mathcal{L} = \lambda_{\text{ce}} \mathcal{L}_{\text{CE}}(y_{t+1}, p(\hat{\mathbf{y}}_{t+1})) + \lambda_{\text{iwp}} \mathcal{L}_{\text{iwp}}(\mathbf{y}_{t+1,1}, p(\hat{\mathbf{y}}_{t+1,1})) + \lambda_{\text{corr}} \mathcal{L}_{\text{corr}} + \lambda_{\text{mse}} \mathcal{L}_{\text{MSE}}(\mathbf{x}_{t+1}, \mathbf{z}_{t+1}),$$

where $\mathcal{L}_{\text{MSE}}(\cdot, \cdot)$ and $\mathcal{L}_{\text{CE}}(\cdot, \cdot)$ denote the mean square error and the cross entropy loss functions, respectively. Here, $\lambda_{\text{ce}}, \lambda_{\text{corr}}, \lambda_{\text{mse}}$ and λ_{iwp} are hyperparameters. $\mathcal{L}_{\text{corr}}$ penalizes the case where $\hat{\mathbf{y}}_{t+1}$ describes an event that occurs before or after t . \mathcal{L}_{iwp} penalizes the case where the predicted word $\mathbf{y}_{t+1,1}^*$ is incorrect. This is defined as $\mathcal{L}_{\text{iwp}}(\mathbf{y}_{t+1,1}, p(\hat{\mathbf{y}}_{t+1,1})) = \gamma_{\text{iwp}} \mathcal{L}_{\text{CE}}(y_{t+1,1}, p(\hat{\mathbf{y}}_{t+1,1}))$, where $\gamma_{\text{iwp}} = 1/W$. The parameter W denotes the number of appearances of the words in the training set.

5. EXPERIMENTS

5.1. Dataset

In the experiment, we evaluated our model on the YouCook2-FC and BILA-caption datasets. The YouCook2-FC dataset is a dataset for the future captioning task. Generally, in cooking, the next procedure is determined based on the previous procedure. We built the YouCook2-FC dataset based on the YouCook2 dataset [22]. We set the number of samples included in the training, validation, and test sets to 7435, 1569, and 3035, respectively.

The BILA-caption dataset was newly built to evaluate future captioning models that describe likely collisions in object placement tasks. Fig. 1 shows a typical sample from of the BILA-caption dataset. To build the dataset, we extended SIGVerse [23], which was used in the World Robot Summit 2018

Partner Robot Challenge/Virtual Space competition [24]. In the simulator, a DSR placed a randomly selected everyday object (e.g., a bottle or can) in the center of one of five types of furniture (e.g., a table or shelf). Each sample was annotated with a statement explaining the situation that occurred as a result of the robot placement action. The dataset consists of 1K videos and 1K english captions. The total and average lengths of the videos are 2.2 h and 8 s, respectively. In the dataset, each clip was given a sentence explaining dangerous events (e.g., collision events) and their causes. The vocabulary size is 245 words. We also set the number of samples included in the training, validation, and test sets to 800, 100, and 100, respectively.

5.2. Experimental setup

We preprocessed clips in the datasets as follows. For the YouCook2-FC dataset, we trimmed the videos using the given start and end times. For the BILA-caption dataset, we trimmed the videos so that the clip starts when the arm of the DSR starts to move and ends when one of the following two conditions have been met: a collision event has occurred and more than 4 s have passed since the start time. Each clip was first converted to 0.6 and 8 fps for the YouCook2-FC and BILA-caption datasets, respectively. For the YouCook2-FC dataset, we used the procedure shown in [25]. For the BILA-caption dataset, we used S3D [26] pretrained on the Howto100m dataset [27] to obtain 512-dimensional features. Then, we used a fully connected layer to obtain 384-dimensional features.

The experimental setup was as follows: the optimizer, learning rate, batch size, and number of epochs were Adam (β_1 : 0.9, β_2 : 0.999), $1.0\text{e-}4$, 16, and 25, respectively. For each module, $N_e, d_{\text{enc}}, N_h, N_d, d_{\text{dec}}, N_r$, and d_{rsa} were 3, 384, 12, 3, 384, 2, 384, respectively. For the loss function, $\lambda_{\text{ce}}, \lambda_{\text{iwp}}, \lambda_{\text{corr}}, \lambda_{\text{mse}}$ were 30, 1.0, 0.1, 0.005, and 10, respectively. For \mathcal{L}_{iwp} , we handled the words that appear more than n_{th} times. We set n_{th} to 30. We also sets W for the YouCook2-FC and BILA-caption datasets to 3000 and 1000, respectively. The number of trainable parameters of RFCM and the number of multiply-add operations are 3.1M and 540M, respectively.

The training was conducted on a machine equipped with an NVIDIA Tesla V100 SXM2 with 16 GB of GPU memory, 240 GB RAM, and an Intel Xeon Gold 6148 processor. It took 6.2 and 1.6×10^{-1} h to train our model on the YouCook2-FC and BILA-caption datasets, respectively. Similarly, inference took 1.7×10^{-2} and 4.9×10^{-2} s/sample on the YouCook2-FC and BILA-caption datasets, respectively. As a condition for early stopping, we used the generalization described in [28].

5.3. Quantitative results

We compared RFCM with Memory-Augmented Recurrent Transformer (MART [7]). We selected MART as the baseline because it is a representative method for video captioning tasks and can be applied to future captioning tasks. Table 1 shows the quantitative results on the YouCook2-FC and

Table 1. Quantitative comparison and ablation studies. The best scores are in bold.

Methods	YouCook2-FC				BILA-caption			
	BLEU4 \uparrow	METEOR \uparrow	ROUGE-L \uparrow	CIDEr-D \uparrow	BLEU4 \uparrow	METEOR \uparrow	ROUGE-L \uparrow	CIDEr-D \uparrow
MART [7]	6.85 \pm 0.18	14.24 \pm 0.07	30.80 \pm 0.21	20.86 \pm 1.07	19.01 \pm 0.74	21.02 \pm 0.45	30.30 \pm 0.64	37.33 \pm 4.37
Ours (w/o RSA)	6.70 \pm 0.36	14.14 \pm 0.50	30.18 \pm 0.18	21.26 \pm 2.83	20.37 \pm 0.36	22.04 \pm 0.19	40.67 \pm 0.48	44.65 \pm 4.89
Ours (w/o Transformer Decoder)	6.68 \pm 0.13	14.09 \pm 0.15	30.16 \pm 0.31	19.75 \pm 1.41	21.08 \pm 1.62	22.39 \pm 0.84	40.92 \pm 1.32	45.05 \pm 6.72
Ours (RFCM)	7.03 \pm 0.15	14.53 \pm 0.09	30.49 \pm 0.21	21.32 \pm 1.09	21.74 \pm 1.02	22.74 \pm 0.57	41.44 \pm 0.86	49.61 \pm 8.02

BILA-caption datasets. The mean and standard deviation were computed on five experimental runs.

The evaluation of the generated sentences was based on several standard metrics for video captioning tasks: BLEU4 [29], ROUGE-L [30], METEOR [31], and CIDEr-D [32]. The primary metric was CIDEr-D.

First, we compared the models on the YouCook2-FC dataset. The table shows that the CIDEr-D score was improved by 0.46 points. The BLEU4 and METEOR scores were also improved by 0.18 and 0.29 points, respectively.

Next, we compared the models on the BILA-caption dataset. The table shows that the CIDEr-D score was drastically improved by 12.28 points. The other metrics scores were also improved. These results indicate that the RFCM generated sentences more appropriately than the baseline.

5.4. Qualitative results

Figs. 3-4 show the qualitative results on the YouCook2-FC and BILA-caption datasets. In the figures, the events are shown in chronological order. Some events are omitted due to space limitations.

The sample in Fig. 3 illustrates a successful case on the YouCook2-FC dataset. In this example, the object added to the pan was “chopped tomatoes.” The baseline method incorrectly described it as “the chopped onions and ginger.” In contrast, our method appropriately described it as “tomato puree.” This result indicates that our method was able to appropriately predict the next step and generate a caption.

Similarly, the sample in Fig. 4 illustrates a successful case on the BILA-caption dataset. In this example, the grasped object was “the white bottle” and the collided object was “the camera.” The baseline method incorrectly described the collided object as “a black teapot.” In contrast, our method appropriately described them as “a white jar” and “the camera,” respectively. This result indicates that our method could appropriately describe the characteristics of the objects.

5.5. Ablation study

We conducted ablation studies for each module. Table 1 quantitatively shows the mean and standard deviation of five experimental runs. We investigated which module contributed the most to the performance improvement using two ablation conditions: (a) w/o RSA, where we used the standard MHA layer [2] instead of the RSA layer, and (b) w/o transformer decoder, where we removed the transformer decoder.

Comparing the results obtained under conditions (a) and (b) with the results of the proposed method on the YouCook2-FC dataset, the CIDEr-D score decreased by 0.06 and 1.57

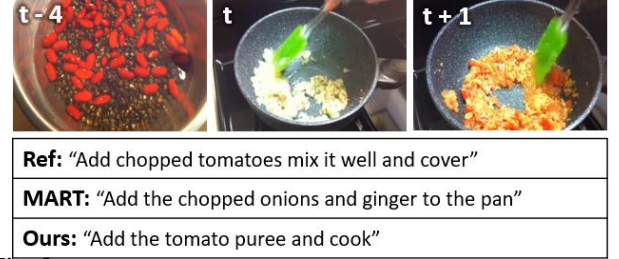


Fig. 3. The successful examples on the YouCook2-FC dataset. Top figures show events. The bottom table shows a reference sentence and sentences generated by the baseline and our methods.

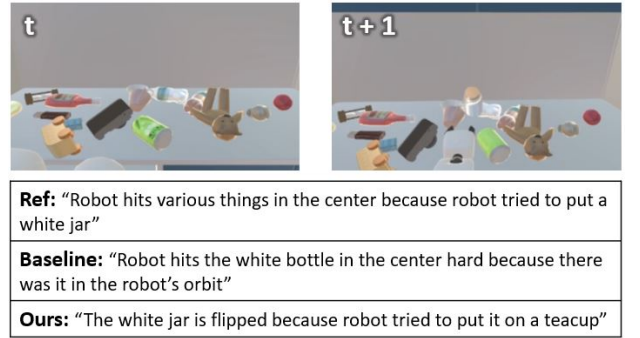


Fig. 4. The successful examples on the BILA-caption dataset.

points, respectively. This indicates that the transformer decoder contributed the most to the performance improvement.

Similarly, comparing results obtained under the conditions (a) and (b) with the results of the proposed method on the BILA-caption dataset, the CIDEr-D score decreased by 4.96 and 4.56 points, respectively. This indicates that the RSA layer contributed the most to the performance improvement.

6. CONCLUSIONS

In this paper, we focused on the future captioning task, which is a task to generate a description about a future event. Specifically, we proposed a future captioning model for daily tasks.

The main contributions of this paper are as follows:

- We proposed the RFCM, a crossmodal language generation model that can generate a description about a future event.
- The RSA Encoder is introduced to extract the relationships between events more effectively than the conventional self-attention in transformers.
- The RFCM outperformed the baseline method on two datasets, BILA-caption and YouCook2-FC.

7. REFERENCES

- [1] Mehrdad Hosseinzadeh et al., “Video Captioning of Future Frames,” in *WACV*, 2021, pp. 980–989.
- [2] Ashish Vaswani, Noam Shazeer, et al., “Attention Is All You Need,” in *NeurIPS*, 2017, pp. 5998–6008.
- [3] Manjin Kim, Heeseung Kwon, Chunyu Wang, et al., “Relational Self-Attention: What’s Missing in Attention for Video Understanding,” in *NeurIPS*, 2021.
- [4] Xin Wang, Wenhu Chen, Jiawei Wu, Yuan-Fang Wang, et al., “Video Captioning via Hierarchical Reinforcement Learning,” in *CVPR*, 2018, pp. 4213–4222.
- [5] Ranjay Krishna, Kenji Hata, et al., “Dense-Captioning Events in Videos,” in *CVPR*, 2017, pp. 706–715.
- [6] Kelvin Xu, Jimmy Ba, Ryan Kiros, et al., “Show, Attend and Tell: Neural Image Caption Generation with Visual Attention,” in *ICML*, 2015, pp. 2048–2057.
- [7] Jie Lei, Liwei Wang, et al., “MART: Memory-Augmented Recurrent Transformer for Coherent Video Paragraph Captioning,” in *ACL*, 2020, pp. 2603–2614.
- [8] Motonari Kambara and Komei Sugiura, “Case Relation Transformer: A Crossmodal Language Generation Model for Fetching Instructions,” *IEEE RA-L*, vol. 6, no. 4, pp. 8371–8378, 2021.
- [9] Aly Magassouba, Komei Sugiura, et al., “Predicting and attending to damaging collisions for placing everyday objects in photo-realistic simulations,” *Advanced Robotics*, vol. 35, no. 12, pp. 1–13, 2021.
- [10] Zakir Hossain, Ferdous Sohel, Mohd Shiratuddin, et al., “A Comprehensive Survey of Deep Learning for Image Captioning,” *ACM CSUR*, vol. 51, no. 6, pp. 1–36, 2019.
- [11] Nayyer Afaq, Ajmal Mian, et al., “Video Description: A Survey of Methods, Datasets, and Evaluation Metrics,” *ACM CSUR*, vol. 52, no. 6, pp. 1–37, 2019.
- [12] Wentian Zhao, Xinxiao Wu, et al., “Multi-modal Dependency Tree for Video Captioning,” in *NeurIPS*, 2021.
- [13] Chen Sun, Austin Myers, Carl Vondrick, et al., “VideoBERT: A joint model for video and language representation learning,” in *ICCV*, 2019, pp. 7464–7473.
- [14] Huaishao Luo, Lei Ji, Botian Shi, Haoyang Huang, et al., “UniVL: A Unified Video and Language Pre-training Model for Multimodal Understanding and Generation,” *arXiv preprint arXiv:2002.06353*, 2020.
- [15] Chaorui Deng, Shizhe Chen, Da Chen, Yuan He, and Qi Wu, “Sketch, Ground, and Refine: Top-Down Dense Video Captioning,” in *CVPR*, 2021, pp. 234–243.
- [16] Jonghwan Mun, Linjie Yang, et al., “Streamlined Dense Video Captioning,” in *CVPR*, 2019, pp. 6588–6597.
- [17] Yuki Mori, Tsubasa Hirakawa, et al., “Image Captioning in Near Future from Vehicle Camera Images and Motion Information,” in *IEEE IV*, 2021, pp. 1378–1384.
- [18] Tahmida Mahmud, Mohammad Billah, et al., “Prediction and Description of Near-Future Activities in Video,” *CVIU*, vol. 210, pp. 103230, 2021.
- [19] Jie Lei, Licheng Yu, Tamara Berg, et al., “What is More Likely to Happen Next? Video-and-Language Future Event Prediction,” in *EMNLP*, 2020, pp. 8769–8784.
- [20] Bohan Wu, Suraj Nair, et al., “Greedy Hierarchical Variational Autoencoders for Large-Scale Video Prediction,” in *CVPR*, 2021, pp. 2318–2328.
- [21] Jacob Devlin et al., “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in *NAACL-HLT*, 2019, pp. 4171–4186.
- [22] Luowei Zhou, Chenliang Xu, and Jason Corso, “Towards Automatic Learning of Procedures From Web Instructional Videos,” in *AAAI*, 2018, pp. 7590–7598.
- [23] Tetsunari Inamura et al., “Development of Robocup @ Home Simulation towards Long-term Large Scale HRI,” in *Robot Soccer World Cup*, 2013, pp. 672–680.
- [24] “World Robot Summit 2018 Partner robot challenge Virtual Space Rules & Regulations,” 2018.
- [25] Simon Ging et al., “COOT: Cooperative Hierarchical Transformer for Video-Text Representation Learning,” in *NeurIPS*, 2020, pp. 22605–22618.
- [26] Antoine Miech, Jean-Baptiste Alayrac, et al., “End-to-End Learning of Visual Representations from Uncurated Instructional Videos,” in *CVPR*, 2020, pp. 9879–9889.
- [27] Antoine Miech et al., “HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips,” in *ICCV*, 2019, pp. 2630–2640.
- [28] Lutz Prechelt, “Automatic early stopping using cross validation: quantifying the criteria,” *Neural Networks*, vol. 11, no. 4, pp. 761–767, 1998.
- [29] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu, “BLEU: a method for automatic evaluation of machine translation,” in *ACL*, 2002, pp. 311–318.
- [30] Chin-Yew Lin, “ROUGE: A package for automatic evaluation of summaries,” in *Text summarization branches out*, 2004, pp. 74–81.
- [31] Satandeep Banerjee and Alon Lavie, “METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments,” in *the ACL Workshop on IEEM for MTS*, 2005, pp. 65–72.
- [32] Ramakrishna Vedantam et al., “CIDEr: Consensus-based Image Description Evaluation,” in *CVPR*, 2015, pp. 4566–4575.