

# R<sup>3</sup>Net: Relation-embedded Representation Reconstruction Network for Change Captioning

Yunbin Tu<sup>1\*</sup>, Liang Li<sup>2†</sup>, Chenggang Yan<sup>3</sup>, Shengxiang Gao<sup>1</sup>, Zhengtao Yu<sup>1†</sup>

<sup>1</sup>Yunnan Key Laboratory of Artificial Intelligence,  
Kunming University of Science and Technology

<sup>2</sup>Key Lab of Intell. Info. Process., Inst. of Comput. Tech., Chinese Academy of Sciences

<sup>3</sup>Intelligent Information Processing Laboratory, Hangzhou Dianzi University

tuyunbin1995@foxmail.com, liang.li@ict.ac.cn

## Abstract

Change captioning is to use a natural language sentence to describe the fine-grained disagreement between two similar images. Viewpoint change is the most typical distractor in this task, because it changes the scale and location of the objects and overwhelms the representation of real change. In this paper, we propose a Relation-embedded Representation Reconstruction Network (R<sup>3</sup>Net) to explicitly distinguish the real change from the large amount of clutter and irrelevant changes. Specifically, a relation-embedded module is first devised to explore potential changed objects in the large amount of clutter. Then, based on the semantic similarities of corresponding locations in the two images, a representation reconstruction module (RRM) is designed to learn the reconstruction representation and further model the difference representation. Besides, we introduce a syntactic skeleton predictor (SSP) to enhance the semantic interaction between change localization and caption generation. Extensive experiments show that the proposed method achieves the state-of-the-art results on two public datasets <sup>1</sup>.

## 1 Introduction

Change captioning aims to generate a natural language sentence to detail what has changed in a pair of similar images. It has many practical applications, such as assisted surveillance, medical imaging, and computer assisted tracking of changes in media assets (Jhamtani and Berg-Kirkpatrick, 2018; Tu et al., 2021).

Different from single-image captioning (Kim et al., 2019; Jiang et al., 2019; Fisch et al., 2020), change captioning addresses two-image captioning, which requires not only to understand both

\* This work was done at VIPL research group, CAS.

† Corresponding author

<sup>1</sup>The code of this paper has been made publicly available at <https://github.com/tuyunbin/R3Net>.

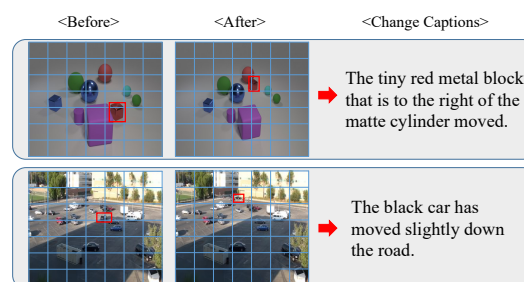


Figure 1: Two examples of change captioning about an object move. The first example shows that the viewpoint changes the scale and location of the objects in the “after” image; the second example shows mostly well-aligned a pair of images with underlying illumination changes from surveillance cameras.

image content, but also to describe their disagreement. As the pioneer work, Jhamtani *et al.* (Jhamtani and Berg-Kirkpatrick, 2018) described semantic changes between mostly well-aligned image pairs with underlying illumination changes from surveillance cameras. However, they did not consider viewpoint changes that often happen in a dynamic world, and image pairs cannot be mostly well aligned in this case. Hence, feature shift between two unaligned images will adversely affect the learning of difference representation. To make this task more practical, recent works (Park et al., 2019; Shi et al., 2020) proposed to address change captioning in the presence of viewpoint changes.

Despite the progress, there are some limitations for the above state-of-the-art methods when modeling the difference representation. First, the object information of each image is only learned at feature-level, and this is difficult to discriminate fine-grained difference when changed object is too tiny and surrounded by the large amount of clutter, as shown in Figure 1. Actually, when an object moved, its semantic relations with surrounding objects would change as well, and this can help explore the fine-grained change. Thus, it is im-

portant to model the difference representation at both feature and relation levels. Second, directly applying subtraction between a pair of unaligned images (Park et al., 2019) may learn the difference representation with much noise, because viewpoint changes the scale and location of the objects. However, we can observe that those unchanged objects are still in the approximate locations. Hence, it is beneficial to reveal the unchanged representation and further model the difference representation based on the semantic similarities in the corresponding locations of two images.

In this paper, we propose a Relation-embedded Representation Reconstruction Network (R<sup>3</sup>Net) to handle viewpoint changes and model the fine-grained difference representation between two images in the process of representation reconstruction. Concretely, for “before” and “after” images, the relation-embedded module respectively performs semantic relation reasoning among object features via the self-attention mechanism. This can enhance the fine-grained representation ability of original object features. To model the difference representation, a representation reconstruction module (RRM) is designed, where a “shadow” representation (“after” or “before”) is used to reconstruct a “source” representation (“before” or “after”). The RRM first leverages every location in the “source” to stimulate the corresponding locations in the “shadow” to judge their semantic similarities, *i.e.*, “response signals”. Further, under the guidance of the signals, the RRM picks out the unchanged features from the “shadow” as the “reconstruction” representation. The “difference” representation is computed with the changed features between the “source” and “reconstruction”. Next, a dual change localizer is devised to use the representation of difference as the query to localize the changed object feature on the “before” and “after”, respectively. Finally, the localized features are fed into an attention-based caption decoder for caption generation.

Besides, we introduce a Syntactic Skeleton Predictor (SSP) to enhance the semantic interaction between change localization and caption generation. As observed in Figure 1, a caption mainly consists of a set of nouns, adjectives, and verbs. These words can convey main information of the changed object and its surrounding references, called *syntactic skeletons*. The skeletons, which are predicted based on a global semantic representation derived

from the R<sup>3</sup>Net, can supervise the modeling of difference representation and provide the decoder with high-level semantic cues about change type. This makes the learned difference representation more relevant to the target words and enhances the quality of generated sentences.

The main contributions of this paper are as follows: (1) We propose the R<sup>3</sup>Net to learn the fine-grained change from the large amount of clutter and overcome viewpoint changes by embedding semantic relations into object features and performing representation reconstruction with respect to the two images. (2) The SSP is introduced to enhance the semantic interaction between change localization and caption generation via predicting a set of syntactic skeletons based on a global semantic representation derived from the R<sup>3</sup>Net. (3) Extensive experiments show that the proposed method outperforms the state-of-the-art approaches by a large margin on two public datasets.

## 2 Related Work

**Change Captioning.** Captioning the change in the existence of viewpoint changes is a novel task in the vision-language community (Zhang et al., 2017; Tu et al., 2017; Deng et al., 2021; Li et al., 2020; Liu et al., 2020). As the first work, DUDA (Park et al., 2019) directly applied subtraction between two images to capture their semantic difference. However, due to viewpoint changes, direct subtraction between an unaligned image pair cannot reliably model the correct change (Shi et al., 2020). Later, M-VAM (Shi et al., 2020) proposed to measure the feature similarity across different regions in an image pair and find the most matched regions as unchanged parts. However, since there are a lot of similar objects, cross-region searching will face the risk of matching the query region with a similar but incorrect region, impacting subsequent change localization. In contrast, in our representation reconstruction network, the prediction of unchanged and changed features are based on the semantic similarities of the corresponding locations in two images. This can avoid the risk of reconstructing “source” with incorrect parts from “shadow”.

**Skeleton Prediction in Captioning.** Syntactic skeletons can provide the high-level semantic cues (e.g., attribute, class) about objects, so they are widely used in image/video captioning works. These methods either used skeletons as main information to generate captions (Fang et al., 2015;




Figure 2: The architecture of the proposed method, consisting of a relation-embedded representation reconstruction network, a syntactic skeleton predictor, a dual change localizer and an attention-based caption decoder.

Gan et al., 2017; Dai et al., 2018) or leveraged them to bridge the semantic gap between vision and language (Gao et al., 2020; Tu et al., 2020). Although the skeletons played different roles in the above methods, the common point was that they only represent basic information of objects in images or videos. Different from them, besides basic information, we try to use skeletons to capture the changed information among objects.

### 3 Methodology

As shown in Figure 2, the architecture of our method consists of four main parts: (1) a relation-embedded representation reconstruction network ( $R^3$ Net) to learn the fine-grained change in the presence of viewpoint changes; (2) a dual change localizer to focus on the specific change in a pair of images; (3) a syntactic skeleton predictor (SSP) to learn syntactic skeletons based on a global semantic representation derived from the  $R^3$ Net; (4) an attention-based caption decoder to describe the change under the guidance of the learned skeletons.

#### 3.1 Relation-embedded Representation Reconstruction Network

##### 3.1.1 Relation-embedded Module

We first exploit a pre-trained CNN model to extract the object-level features  $X_{bef}$  and  $X_{aft}$  for a

pair of “before” and “after” images, where  $X_i \in \mathbb{R}^{C \times H \times W}$  and  $C, H, W$  indicate the number of channels, height, and width. However, only utilizing these independent features is difficult to distinguish fine-grained change from the large amount of clutter (similar objects). And related works (Wu et al., 2019; Huang et al., 2020; Yin et al., 2020) have shown that capturing semantic relations among objects is useful for a thorough understanding of an image.

Motivated by this, we devise a relation-embedded ( $R_{emb}$ ) module based on the self-attention mechanism (Vaswani et al., 2017) to implicitly learn semantic relations among objects in each image. Specifically, we first reshape  $X_i \in \mathbb{R}^{C \times H \times W}$  to  $X_i \in \mathbb{R}^{N \times C}$  ( $N = HW$ ), where  $i \in (bef, aft)$ . Then, the semantic relations are embedded into independent object features of each image based on the scaled dot-product attention:

$$R_{emb}(Q, K, V) = \text{softmax} \left( \frac{QK^T}{\sqrt{d_k}} \right) V, \quad (1)$$

where the queries, keys and values are the projections of the object features in  $X_i$  and  $i \in (bef, aft)$ :

$$(Q, K, V) = (X_i W_i^Q, X_i W_i^K, X_i W_i^V). \quad (2)$$

Thus,  $X_{bef}$  and  $X_{aft}$  are updated to  $\hat{X}_{bef}$  and  $\hat{X}_{aft}$ ,

respectively:

$$\begin{aligned}\hat{X}_{bef} &= \text{R}_{emb}(X_{bef}, X_{bef}, X_{bef}), \\ \hat{X}_{aft} &= \text{R}_{emb}(X_{aft}, X_{aft}, X_{aft}).\end{aligned}\quad (3)$$

When the model fully understands each image content, it can better capture the fine-grained difference between the image pair in the subsequent representation reconstruction.

### 3.1.2 Representation Reconstruction Module

The state-of-the-art method (Park et al., 2019) applied direct subtraction between a pair of unaligned images, which is prone to capture the difference with noise in the presence of viewpoint changes.

To distinguish semantic change from viewpoint changes, a representation reconstruction module (RRM) is proposed, where the inputs are a ‘‘source’’ representation  $\hat{X}_p \in \mathbb{R}^{N \times C}$  and a ‘‘shadow’’ representation  $\hat{X}_s \in \mathbb{R}^{N \times C}$ . Concretely, first, we exploit each location of  $\hat{X}_p$  to stimulate the corresponding location of  $\hat{X}_s$ . The response degrees of all locations in  $\hat{X}_s$  are regarded as the response signals  $\alpha$  that measure the semantic similarities between corresponding locations in two images :

$$\alpha = \text{Sigmoid}(\hat{X}_p W_p + \hat{X}_s W_s + b_s), \quad (4)$$

where  $W_p, W_s \in \mathbb{R}^{C \times C}$  and  $b_s \in \mathbb{R}^C$ . Second, we use  $\hat{X}_s$  to reconstruct  $\hat{X}_p$  under the guidance of the response signals  $\alpha$ :

$$\tilde{X}_p = \alpha \odot \hat{X}_s, \quad (5)$$

where  $\tilde{X}_p \in \mathbb{R}^{N \times C}$  is the ‘‘reconstruction’’ representation, which represent unchanged features with respect to ‘‘source’’. Finally, the ‘‘difference’’ representation is captured by subtracting ‘‘reconstruction’’  $\tilde{X}_p$  from ‘‘source’’  $\hat{X}_p$ :

$$\hat{X}_{diff} = \hat{X}_p - \tilde{X}_p. \quad (6)$$

Since the predicted unchanged and changed features in uni-directional reconstruction are only with respect to one kind of ‘‘source’’ representation (e.g., ‘‘before’’), the model cannot predict the changed feature when it is not in the ‘‘source’’. For an efficient model, it should capture all underlying changes with respect to both images. To this end, we extend the RRM from uni-direction to bi-direction. Specifically, we first use the ‘‘before’’ as ‘‘source’’ to predict unchanged and changed features, and then use the ‘‘after’’ as ‘‘source’’ to do so. Thus, the

‘‘reconstruction’’ and ‘‘difference’’ w.r.t. the ‘‘before’’ and ‘‘after’’ are formulated as:

$$\begin{aligned}\tilde{X}_p^{bef}, \hat{X}_{diff}^{bef} &= \text{RRM}(\hat{X}_p^{bef}, \hat{X}_s^{aft}), \\ \tilde{X}_p^{aft}, \hat{X}_{diff}^{aft} &= \text{RRM}(\hat{X}_p^{aft}, \hat{X}_s^{bef}).\end{aligned}\quad (7)$$

Finally, we obtain a bi-directional difference representation by a fully-connected layer:

$$\hat{X}_{diff} = \text{ReLU}\left(\left[\hat{X}_{diff}^{bef}; \hat{X}_{diff}^{aft}\right] W_f + b_f\right). \quad (8)$$

### 3.2 Dual Change localizer

When the bi-directional difference representation  $\hat{X}_{diff}$  is computed, we exploit it as the query to localize the changed feature in  $\hat{X}_{bef}$  and  $\hat{X}_{aft}$ , respectively. Specifically, the dual change localizer first predicts two separate attention maps  $a_{bef}$  and  $a_{aft}$ :

$$\begin{aligned}X'_{bef} &= [\hat{X}_{bef}; \hat{X}_{diff}], X'_{aft} = [\hat{X}_{aft}; \hat{X}_{diff}], \\ a_{bef} &= \sigma(\text{conv}_2(\text{ReLU}(\text{conv}_1(X'_{bef}))))), \\ a_{aft} &= \sigma(\text{conv}_2(\text{ReLU}(\text{conv}_1(X'_{aft}))))),\end{aligned}\quad (9)$$

where  $[\cdot]$ ,  $\text{conv}$ , and  $\sigma$  denote concatenation, convolutional layer, and sigmoid activation function, respectively. Then, the changed features  $l_{bef}$  and  $l_{aft}$  are localized via applying  $a_{bef}$  and  $a_{aft}$  to  $\hat{X}_{bef}$  and  $\hat{X}_{aft}$ :

$$\begin{aligned}l_{bef} &= \sum_{H,W} a_{bef} \odot \hat{X}_{bef}, l_{bef} \in \mathbb{R}^C, \\ l_{aft} &= \sum_{H,W} a_{aft} \odot \hat{X}_{aft}, l_{aft} \in \mathbb{R}^C.\end{aligned}\quad (10)$$

Finally, we compute the local difference feature w.r.t. both  $l_{bef}$  and  $l_{aft}$  from two directions:

$$\begin{aligned}l_{diff}^{b \rightarrow a} &= l_{bef} - l_{aft}, \quad l_{diff}^{a \rightarrow b} = l_{aft} - l_{bef}, \\ l_{diff}^{b \leftrightarrow a} &= \text{ReLU}\left(W_d \left[ l_{diff}^{b \rightarrow a}; l_{diff}^{a \rightarrow b} \right] + b_d\right).\end{aligned}\quad (11)$$

### 3.3 Syntactic Skeleton Predictor

A syntactic skeleton predictor (SSP) is introduced to learn a set of syntactic skeletons based on the outputs derived from the R<sup>3</sup>Net. The predicted skeletons can provide the caption decoder with high-level semantic cues about changed objects and supervise the modeling of difference representation. This aims to enhance the semantic interaction between change localization and caption generation.

Inspired by (Gao et al., 2020; Gan et al., 2017), we treat this problem as a multi-label classification task. Suppose there are  $N$  training image pairs, and  $y_j = [y_{j1}, \dots, y_{jK}] \in \{0, 1\}^K$  is the label vector of the  $j$ -th image pair, where  $y_{jk} = 1$  if the image pair is annotated with the skeleton  $k$ , and  $y_{jk} = 0$  otherwise.

Specifically, first, we apply a mean-pooling layer over the concatenated semantic representations of  $\hat{X}_{bef}$ ,  $\hat{X}_{aft}$ , and  $\hat{X}_{diff}$  to obtain a global semantic representation  $S_j$ :

$$S_j = \frac{1}{H, W} \sum_{H, W} [\hat{X}_{bef}; \hat{X}_{aft}; \hat{X}_{diff}]. \quad (12)$$

Then, the probability scores  $p_j$  of all syntactic skeletons for  $j$ -th image pair is computed by:

$$p_j = \text{sigmoid}(U_g \text{ReLU}(W_g S_j) + b_g), \quad (13)$$

where  $p_j = [p_{j1}, \dots, p_{jK}]$  denotes the probability scores of  $K$  skeletons in  $j$ -th image pair. To maximize the probability scores of syntactic skeletons, we use the multi-label loss to optimize the SSP. It can be formulated as:

$$L_s = -\frac{1}{N} \sum_{j=1}^N \sum_{k=1}^K (y_{jk} \log p_{jk} + (1 - y_{jk}) \log (1 - p_{jk})), \quad (14)$$

where  $N$  and  $K$  indicate the number of all training samples and annotated skeletons of an image pair. The loss can be considered as the supervision signal to regularize the learning of difference representation in the R<sup>3</sup>Net.

### 3.4 Skeleton-guided Caption generation

Since the predicted skeletons are the explicit semantic concepts of the changed object and its surrounding references, the captions are generated under the guidance of them. Specifically, first, the predicted probability scores  $p_j$  are embedded as a skeleton feature  $E[p_j]$ :

$$E[p_j] = \text{ReLU}(W_q(E_q p_j) + b_q), \quad (15)$$

where  $E_q \in \mathbb{R}^{k \times M}$  is a skeleton embedding matrix and  $M$  is the dimension of the skeleton feature.  $W_q \in \mathbb{R}^{M \times M}$  and  $b_q \in \mathbb{R}^M$  are the parameters to be learned. Then, we exploit a semantic attention module to focus on the key semantic feature from  $l_{bef}$ ,  $l_{aft}$ , and  $l_{diff}^{b \leftrightarrow a}$ , which is relevant to the target word:

$$l_{\text{dyn}}^{(t)} = \sum_i \beta_i^{(t)} l_i, \quad (16)$$

where  $i \in (bef, aft, diff)$ .  $\beta_i^{(t)}$  is computed by an attention LSTM<sub>a</sub> under the guidance of the predicted skeleton feature  $E[p_j]$ :

$$\begin{aligned} v &= \text{ReLU}\left(W_{a_1} \left[ l_{bef}; l_{diff}^{b \leftrightarrow a}; l_{aft} \right] + b_{a_1}\right), \\ u^{(t)} &= \left[ v; E[p_j]; h_c^{(t-1)} \right], \\ h_a^{(t)} &= \text{LSTM}_a \left( h_a^{(t)} \mid u^{(t)}, h_a^{(0:t-1)} \right), \\ \beta^{(t)} &\sim \text{Softmax}\left(W_{a_2} h_a^{(t)} + b_{a_2}\right). \end{aligned} \quad (17)$$

where  $W_{a_1}$ ,  $b_{a_1}$ ,  $W_{a_2}$ , and  $b_{a_2}$  are learnable parameters.  $h_a^{(*)}$  and  $h_c^{(*)}$  are hidden states of the attention module LSTM<sub>a</sub> and the caption decoder LSTM<sub>c</sub>, respectively.

Finally, the caption generation process is also guided by the predicted skeleton feature. We feed it, attended visual feature, and the previous word embedding to the caption decoder LSTM<sub>c</sub> to predict a series of distributions over the next word:

$$\begin{aligned} c^{(t)} &= \left[ E[w_{t-1}]; E[p_j]; l_{\text{dyn}}^{(t)} \right], \\ h_c^{(t)} &= \text{LSTM}_c \left( h_c^{(t)} \mid c^{(t)}, h_c^{(0:t-1)} \right), \\ w_t &\sim \text{Softmax}\left(W_c h_c^{(t)} + b_c\right), \end{aligned} \quad (18)$$

where  $E$  is a word embedding matrix;  $W_c$  and  $b_c$  are learnable parameters.

### 3.5 Joint Training

We jointly train the caption decoder and SSP in an end-to-end manner. For the SSP, the multi-label loss is minimized by the Eq. (14). For the decoder, given the target ground-truth words  $(w_1, \dots, w_m)$ , we minimize its negative log-likelihood loss:

$$L_{\text{cap}}(\theta_c) = -\sum_{t=1}^m \log p(w_t \mid w_{<t}; \theta_c), \quad (19)$$

where  $\theta_c$  are the parameters of the decoder and  $m$  is the length of the caption. The final loss function is optimized as follows:

$$L(\theta) = L_{\text{cap}} + \lambda * L_s, \quad (20)$$

where the hyper-parameter  $\lambda$  is to seek a trade-off between the decoder and SSP.

## 4 Experiments

### 4.1 Datasets and Evaluation Metrics

**CLEVR-Change** dataset (Park et al., 2019) is a large-scale dataset with a set of basic geometry

Table 1: Ablation studies on CLEVR-Change in terms of total performance.

Method	Total				
	BLEU-4	METEOR	ROUGE-L	CIDEr	SPICE
Baseline	53.1	37.6	70.8	115.6	31.3
RRM	53.5	39.2	72.3	119.2	32.5
R <sup>3</sup> Net	54.2	39.4	72.7	122.3	<b>32.6</b>
R <sup>3</sup> Net+SSP	<b>54.7</b>	<b>39.8</b>	<b>73.1</b>	<b>123.0</b>	<b>32.6</b>

Table 2: Ablation studies on CLEVR-Change in terms of two settings, where B-4, M, R, C, and S are short for BLEU-4, METEOR, ROUGE-L, CIDEr, and SPICE, respectively.

Method	Scene Change					None-scene Change				
	B-4	M	R	C	S	B-4	M	R	C	S
Baseline	51.0	33.3	65.7	102.4	28.0	61.0	49.9	75.8	114.3	34.5
RRM	51.8	35.7	69.0	110.1	30.4	60.0	49.6	75.6	115.0	34.5
R <sup>3</sup> Net	52.5	36.0	69.5	114.8	<b>30.5</b>	<b>62.0</b>	50.0	75.9	116.3	<b>34.8</b>
R <sup>3</sup> Net+SSP	<b>52.7</b>	<b>36.2</b>	<b>69.8</b>	<b>116.6</b>	30.3	61.9	<b>50.5</b>	<b>76.4</b>	<b>116.4</b>	<b>34.8</b>

objects, which consists of 79,606 image pairs and 493,735 captions. The change types can be categorized into six cases, *i.e.*, “Color”, “Texture”, “Add”, “Drop”, “Move” and “Distractors (e.g., viewpoint change)”. We use the official split with 67,660 for training, 3,976 for validation and 7,970 for testing.

**Spot-the-Diff** dataset (Jhamtani and Berg-Kirkpatrick, 2018) contains 13,192 well-aligned image pairs from surveillance cameras. Based on the official split, the dataset is split into training, validation, and testing with a ratio of 8:1:1.

Following the state-of-the-art methods (Park et al., 2019; Shi et al., 2020; Tu et al., 2021), we use five standard metrics to evaluate the quality of generated sentences, *i.e.*, BLEU-4 (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), ROUGE-L (Lin, 2004), CIDEr (Vedantam et al., 2015) and SPICE (Anderson et al., 2016). We get all the results based on the Microsoft COCO evaluation server (Chen et al., 2015).

## 4.2 Implementation Details

We use ResNet-101 (He et al., 2016) pre-trained on the Imagenet dataset (Russakovsky et al., 2015) to extract object features, with the dimension of  $1024 \times 14 \times 14$ . We project these features into a lower dimension of 256. The hidden size of overall model is set to 512 and the number of attention heads in relation-embedded module is set to 4. The number of skeletons in an image pair is set to 50. The dimension of words is set to 300. For the hyper-parameter  $\lambda$ , we empirically set it as 0.1. In the training phase, we use Adam

optimizer (Kingma and Ba, 2014) with the learning rate of  $1 \times 10^{-3}$ , and set the mini-batch size as 128 and 64 on CLEVR-Change and Spot-the-Diff. At inference, for fair comparison, we follow the pioneer works (Park et al., 2019; Jhamtani and Berg-Kirkpatrick, 2018) in the two datasets to use greedy decoding strategy for caption generation. Both training and inference are implemented with PyTorch (Paszke et al., 2019) on a Tesla P100 GPU.

## 4.3 Ablation Studies

To figure out the contribution of each module of the proposed network, we conduct the following ablation studies on CLEVR-Change: (1) Baseline which is based on DUDA (Park et al., 2019); (2) RRM which is the representation reconstruction module; (3) R<sup>3</sup>Net which augments the RRM with a relation-embedded module; (4) R<sup>3</sup>Net+SSP which augments the R<sup>3</sup>Net with a syntactic skeleton predictor.

**The evaluation on Total Performance.** Total performance is to simultaneously evaluate the model under both scene change and none-scene change. Experimental results are shown in Table 1. We can observe that each module and the full method improve the total performance of Baseline. This indicates that our method not only can correctly judge whether there is an semantic change between a pair of images, but also can describe the change in an accurate natural language sentence.

**The evaluation on the settings of Scene Change and None-scene Change.** In the setting of scene change, both object and viewpoint changes

Table 3: Comparing with state-of-the-art methods on CLEVR-Change in Total Performance. RL refers to the training strategy of reinforcement learning.

Method	RL	Total				
		B-4	M	R	C	S
Capt-Dual (Park et al., 2019)	×	43.5	32.7	-	108.5	23.4
DUDA (Park et al., 2019)	×	47.3	33.9	-	112.3	24.5
M-VAM (Shi et al., 2020)	×	50.3	37.0	69.7	114.9	30.5
M-VAM+RAF (Shi et al., 2020)	✓	51.3	37.8	70.4	115.8	30.7
R <sup>3</sup> Net+SSP	×	<b>54.7</b>	<b>39.8</b>	<b>73.1</b>	<b>123.0</b>	<b>32.6</b>

Table 4: Comparing with state-of-the-art methods on CLEVR-Change in terms of two settings.

Method	RL	Scene Change				None-scene Change			
		B-4	M	C	S	B-4	M	C	S
Capt-Dual (Park et al., 2019)	×	38.5	28.5	89.8	18.2	56.3	44.0	108.9	28.7
DUDA (Park et al., 2019)	×	42.9	29.7	94.6	19.9	59.8	45.2	110.8	29.1
M-VAM+RAF (Shi et al., 2020)	✓	-	-	-	-	-	<b>66.4</b>	<b>122.6</b>	33.4
R <sup>3</sup> Net+SSP	×	<b>52.7</b>	<b>36.2</b>	<b>116.6</b>	<b>30.3</b>	<b>61.9</b>	50.5	116.4	<b>34.8</b>

happen. In the setting of none-scene change, there are only distractors, such as viewpoint change and illumination change. The experimental results are shown in Table 2. Under the setting of scene change, we can observe that 1) the RRM, R<sup>3</sup>Net, and R<sup>3</sup>Net+SSP all significantly improve the Baseline; 2) the R<sup>3</sup>Net is much better than the RRM; 3) the best performance is achieved when augmenting the R<sup>3</sup>Net with the SSP. The above observations indicate that 1) compared to direct subtraction between a pair of unaligned images, it is effective to capture difference representation via the R<sup>3</sup>Net, because it can overcome the distraction of viewpoint change; 2) learning semantic relations among object features is important, because these relations can enrich the raw object features, helpful for exploring fine-grained changes; 3) the SSP can enhance the semantic interaction between change localization and caption generation, and thus further improve the quality of generated sentences.

Besides, under the setting of non-scene change, we can observe that the RRM is worse than the Baseline on some metrics. Our conjecture is that, on one hand, due to the large amount of clutter and only representing the image pair at feature-level, the RRM cannot learn the exact semantic similarities of corresponding locations in the two images, performing worse on some metrics. On the other hand, the Baseline learns a coarse difference representation between two unaligned images by a direct subtraction, so it is prone to learn a wrong change

type or simply judge nothing has changed. This leads to the results that it performs worse than the RRM with the total performance and scene change, but achieves higher scores than the RRM on some metrics with none-scene change. In fact, when embedding semantic relations among object features, the R<sup>3</sup>Net outperforms the Baseline in the both settings. This further indicates that it is beneficial to thoroughly understand image content via modeling semantic relations among object features.

#### 4.4 Performance Comparison

##### 4.4.1 Results on CLEVR-Change Dataset

In this dataset, we compare with four state-of-the-art methods, Capt-Dual (Park et al., 2019), DUDA (Park et al., 2019), M-VAM (Shi et al., 2020), and M-VAM+RAF (Shi et al., 2020), in four settings: 1) scene change; 2) none-scene change; 3) total (scene and none-scene changes); 4) specific types of scene change.

From Table 3 and Table 4, under two kinds of settings and total performance, we can observe that our method surpasses Capt-Dual and DUDA with a large margin. Compared to M-VAM+RAF, the total performance of our method is much better than it, which indicates our method is more robust. As shown in Table 4, under the setting of none-scene change, it outperforms our method on METEOR and CIDEr. This could be a benefit of the reinforcement learning, while also sharply increasing the training time and computation complexity.

Table 5: A Detailed breakdown of Change Captioning evaluation on CLEVR-Change by different change types: “Color” (C), “Texture” (T), “Add” (A), “Drop” (D), and “Move” (M).

Method	RL	Metrics	C	T	A	D	M
Capt-Dual (Park et al., 2019)	×	CIDEr	115.8	82.7	85.7	103.0	52.6
DUDA (Park et al., 2019)	×	CIDEr	120.4	86.7	108.3	103.4	56.4
M-VAM+RAF (Shi et al., 2020)	✓	CIDEr	122.1	98.7	<b>126.3</b>	115.8	82.0
R <sup>3</sup> Net+SSP	×	CIDEr	<b>139.2</b>	<b>123.5</b>	122.7	<b>121.9</b>	<b>88.1</b>
Capt-Dual (Park et al., 2019)	×	METEOR	32.1	26.7	29.5	31.7	22.4
DUDA (Park et al., 2019)	×	METEOR	32.8	27.3	33.4	31.4	23.5
M-VAM+RAF (Shi et al., 2020)	✓	METEOR	35.8	32.3	37.8	36.2	27.9
R <sup>3</sup> Net+SSP	×	METEOR	<b>38.9</b>	<b>35.5</b>	<b>38.0</b>	<b>37.5</b>	<b>30.9</b>
Capt-Dual (Park et al., 2019)	×	SPICE	19.8	17.6	16.9	21.9	14.7
DUDA (Park et al., 2019)	×	SPICE	21.2	18.3	22.4	22.2	15.4
M-VAM+RAF (Shi et al., 2020)	✓	SPICE	28.0	26.7	30.8	<b>32.3</b>	22.5
R <sup>3</sup> Net+SSP	×	SPICE	<b>31.6</b>	<b>30.8</b>	<b>32.3</b>	31.7	<b>25.4</b>

Table 6: Comparing with state-of-the-art methods on Spot-the-Diff.

Method	RL	M	R	C	S
DDLA	×	12.0	28.6	32.8	-
DUDA	×	11.8	29.1	32.5	-
SDCM	×	12.7	29.7	36.3	-
FCC	×	12.9	29.9	36.8	-
static rel-att	×	13.0	28.3	34.0	-
dynamic rel-att	×	12.2	31.4	35.3	-
M-VAM	×	12.4	31.3	38.1	14.0
M-VAM+RAF	✓	12.9	<b>33.2</b>	<b>42.5</b>	17.1
R <sup>3</sup> Net+SSP	×	<b>13.1</b>	32.6	36.6	<b>18.8</b>

Table 5 is the specific change types. Among five changes, the most challenging types are “Texture” and “Move”, because they are always confused with irrelevant illumination or viewpoint changes. Compared to the SOTA methods, our method achieves excellent performances under both change types. This shows that our method can better distinguish the attribute change or movement of objects from the illumination or viewpoint change.

Hence, compared to the current SOTA methods from different dimensions, the generalization ability of our method is much better. This benefits from the merits that 1) the R<sup>3</sup>Net can learn the fine-grained change and overcome viewpoint changes in the process of representation reconstruction; 2) the SSP can enhance the semantic interactions between change localization and caption generation.

#### 4.4.2 Results on Spot-the-Diff Dataset

The image pairs in this dataset are mostly well aligned. We compare with eight SOTA meth-

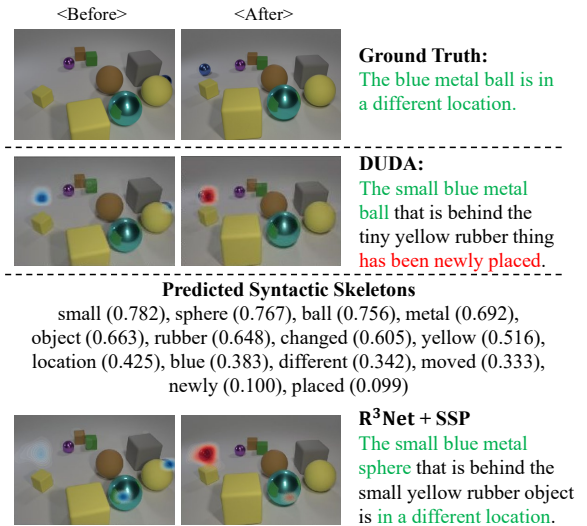


Figure 3: A example about “Move” case from the test set of CLEVR-Change, which involves the caption generated by humans (Ground Truth), DUDA (current SOTA method) and R<sup>3</sup>Net+SSP. We also visualize the predicted syntactic skeletons and the localization results on the “before” (blue) and “after” (red).

ods and most of them cannot consider handling viewpoint changes: DDLA (Jhamtani and Berg-Kirkpatrick, 2018), DUDA (Park et al., 2019), SDCM (Oluwasanmi et al., 2019a), FCC (Oluwasanmi et al., 2019b), static rel-att / dynamic rel-att (Tan et al., 2019), and M-VAM / M-VAM+RAF (Shi et al., 2020).

The results are shown in Table 6. We can observe that when training without reinforcement learning, our method achieves the best performances on METEOR, ROUGE-L and SPICE. Compared to M-VAM+RAF trained by the reinforcement learn-






Figure 4: Qualitative examples of R<sup>3</sup>Net+SSP. The left is a successful case that R<sup>3</sup>Net+SSP localizes the accurate changed object and generates a correct sentence to describe the change. The right is a failure case that a slight movement of the object is not correctly described.

ing strategy, our method still outperforms them on METEOR and SPICE. Since there is no viewpoint change in this dataset, the superiority mainly results from that the relation-embedded module can enhance the fine-grained representation ability of object features, and the syntactic skeleton predictor can enhance the semantic interaction between change localization and caption generation.

#### 4.5 Qualitative Analysis

Figure 3 shows an example about the case of “Move” from the test set of CLEVR-Change. We can observe that DUDA localizes a wrong region on the “before” and thus misidentifies “Move” as “Add”. By contrast, the R<sup>3</sup>Net+SSP can accurately locate the moved object on the “before” and “after” images, which benefits from two merits. First, the R<sup>3</sup>Net is able to localize the fine-grained change in the presence of viewpoint changes. Second, the SSP can predict the key skeletons based on the representations of image pair and their difference learned from the R<sup>3</sup>Net. For instance, the skeletons of “changed” and “location” has the higher probability scores than “newly” and “placed”. This can provide the decoder with high-level semantic cues to generate the correct sentence.

Figure 4 illustrates two cases about “Move”. In the left example, the R<sup>3</sup>Net+SSP successfully distinguishes the changed object (i.e., small grey cylinder) and predicts accurate skeletons with high probability scores. The right example is a failure case. In general, we can observe that the grey cylinder is localized and the main skeletons are predicted,

which indicates that the R<sup>3</sup>Net learns a reliable representation of difference. However, the decoder still generates the wrong sentence. The reason behind the failure may be that the movement of this cylinder is very slight and the decoder receives the weak information of change (including skeletons). In our opinion, a possible solution for this challenge is to model position information for object features, which would enhance their position representation ability and help localize the slight movements.

## 5 Conclusion

In this paper, we propose a relation-embedded representation reconstruction network (R<sup>3</sup>Net) and a syntactic skeleton predictor (SSP) to address change captioning in the presence of viewpoint changes, where the R<sup>3</sup>Net can explicitly distinguish semantic changes from viewpoint changes and the SSP is to enhance the semantic interaction between change localization and caption generation. Extensive experiments show that the state-of-the-art results are achieved on two public datasets, CLEVR-Change and Spot-the-Diff.

## Acknowledgements

The work was supported by National Natural Science Foundation of China (No. 61761026, 61972186, 61732005, 61762056, 61771457, and 61732007), in part by Youth Innovation Promotion Association of Chinese Academy of Sciences (No. 2020108), and CCF-Baidu Open Fund (No. 2021PP15002000).

## References

- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. SPICE: semantic propositional image caption evaluation. In *ECCV*, pages 382–398.
- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL*, pages 65–72.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Bo Dai, Sanja Fidler, and Dahua Lin. 2018. A neural compositional paradigm for image captioning. *NeurIPS*, 31:658–668.
- Jincan Deng, Liang Li, Beichen Zhang, Shuhui Wang, Zhengjun Zha, and Qingming Huang. 2021. Syntax-guided hierarchical attention network for video captioning. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, et al. 2015. From captions to visual concepts and back. In *CVPR*, pages 1473–1482.
- Adam Fisch, Kenton Lee, Ming-Wei Chang, Jonathan H Clark, and Regina Barzilay. 2020. Capwap: Captioning with a purpose. In *EMNLP*, pages 8755–8768.
- Zhe Gan, Chuang Gan, Xiaodong He, Yunchen Pu, Kenneth Tran, Jianfeng Gao, Lawrence Carin, and Li Deng. 2017. Semantic compositional networks for visual captioning. In *CVPR*, pages 5630–5639.
- Lianli Gao, Xuanhan Wang, Jingkuan Song, and Yang Liu. 2020. Fused gru with semantic-temporal attention for video captioning. *Neurocomputing*, 395:222–228.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*, pages 770–778.
- Qingbao Huang, Jielong Wei, Yi Cai, Changmeng Zheng, Junying Chen, Ho-fung Leung, and Qing Li. 2020. Aligned dual channel graph convolutional network for visual question answering. In *ACL*, pages 7166–7176.
- Harsh Jhamtani and Taylor Berg-Kirkpatrick. 2018. Learning to describe differences between pairs of similar images. In *EMNLP*, pages 4024–4034.
- Ming Jiang, Junjie Hu, Qiuyuan Huang, Lei Zhang, Jana Diesner, and Jianfeng Gao. 2019. Relevance, extraneousness, omission: A fine-grained evaluation for image captioning. In *EMNLP*, pages 1475–1480.
- Dong-Jin Kim, Jinsoo Choi, Tae-Hyun Oh, and In So Kweon. 2019. Image captioning with very scarce supervised data: Adversarial semi-supervised learning approach. In *EMNLP*, pages 2012–2023.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Liang Li, Shijie Yang, Li Su, Shuhui Wang, Chenggang Yan, Zheng-jun Zha, and Qingming Huang. 2020. Diverter-guider recurrent network for diverse poems generation from image. In *ACM MM*, pages 3875–3883.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Zhenhuan Liu, Jincan Deng, Liang Li, Shaofei Cai, Qianqian Xu, Shuhui Wang, and Qingming Huang. 2020. Ir-gan: Image manipulation with linguistic instruction by increment reasoning. In *ACM MM*, pages 322–330.
- Ariyo Oluwasanmi, Muhammad Umar Aftab, Eatedal Alabdulkreem, Bulbula Kumeda, Edward Y Baagyere, and Zhiqiang Qin. 2019a. Captionnet: Automatic end-to-end siamese difference captioning model with attention. *IEEE Access*, 7:106773–106783.
- Ariyo Oluwasanmi, Enoch Frimpong, Muhammad Umar Aftab, Edward Y Baagyere, Zhiqiang Qin, and Kifayat Ullah. 2019b. Fully convolutional captionnet: Siamese difference captioning attention model. *IEEE Access*, 7:175929–175939.
- 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.
- Dong Huk Park, Trevor Darrell, and Anna Rohrbach. 2019. Robust change captioning. In *ICCV*, pages 4623–4632.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pages 8024–8035.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. In *IJCV*, pages 211–252.
- Xiangxi Shi, Xu Yang, Jiuxiang Gu, Shafiq R. Joty, and Jianfei Cai. 2020. Finding it at another side: A viewpoint-adapted matching encode for change captioning. In *ECCV*, pages 574–590.

- Hao Tan, Franck Deroncourt, Zhe Lin, Trung Bui, and Mohit Bansal. 2019. Expressing visual relationships via language. In *ACL*, pages 1873–1883.
- Yunbin Tu, Tingting Yao, Liang Li, Jiedong Lou, Shengxiang Gao, Zhengtao Yu, and Chenggang Yan. 2021. Semantic relation-aware difference representation learning for change captioning. In *Findings of ACL*, pages 63–73.
- Yunbin Tu, Xishan Zhang, Bingtao Liu, and Chenggang Yan. 2017. Video description with spatial-temporal attention. In *ACM MM*, pages 1014–1022.
- Yunbin Tu, Chang Zhou, Junjun Guo, Shengxiang Gao, and Zhengtao Yu. 2020. Enhancing the alignment between target words and corresponding frames for video captioning. *Pattern Recognition*, page 107702.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*, pages 5998–6008.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *CVPR*, pages 4566–4575.
- Aming Wu, Linchao Zhu, Yahong Han, and Yi Yang. 2019. Connective cognition network for directional visual commonsense reasoning. In *NeurIPS*, pages 5669–5679.
- Yongjing Yin, Fandong Meng, Jinsong Su, Chulun Zhou, Zhengyuan Yang, Jie Zhou, and Jiebo Luo. 2020. A novel graph-based multi-modal fusion encoder for neural machine translation. In *ACL*, pages 3025–3035.
- Xishan Zhang, Ke Gao, Yongdong Zhang, Dongming Zhang, Jintao Li, and Qi Tian. 2017. Task-driven dynamic fusion: Reducing ambiguity in video description. In *CVPR*, pages 3713–3721.