

# Distractors-Immune Representation Learning with Cross-modal Contrastive Regularization for Change Captioning

Yunbin Tu<sup>1</sup>, Liang Li<sup>2\*</sup>, Li Su<sup>1\*</sup>, Chenggang Yan<sup>3</sup>, and Qingming Huang<sup>1</sup>

<sup>1</sup> University of Chinese Academy of Sciences, Beijing, China

tuyunbin22@mailsucas.ac.cn, {suli,qmhuang}@ucas.ac.cn

<sup>2</sup> Key Laboratory of AI Safety of CAS, Institute of Computing Technology, Chinese Academy of Sciences (CAS), Beijing, China liang.li@ict.ac.cn

<sup>3</sup> Hangzhou Dianzi University (HDU) & Lishui Institute of HDU, China

**Abstract.** Change captioning aims to succinctly describe the semantic change between a pair of similar images, while being immune to distractors (illumination and viewpoint changes). Under these distractors, unchanged objects often appear pseudo changes about location and scale, and certain objects might overlap others, resulting in perturbational and discrimination-degraded features between two images. However, most existing methods directly capture the difference between them, which risk obtaining error-prone difference features. In this paper, we propose a distractors-immune representation learning network that correlates the corresponding channels of two image representations and decorrelates different ones in a self-supervised manner, thus attaining a pair of stable image representations under distractors. Then, the model can better interact them to capture the reliable difference features for caption generation. To yield words based on the most related difference features, we further design a cross-modal contrastive regularization, which regularizes the cross-modal alignment by maximizing the contrastive alignment between the attended difference features and generated words. Extensive experiments show that our method outperforms the state-of-the-art methods on four public datasets. The code is available at <https://github.com/tuyunbin/DIRL>.

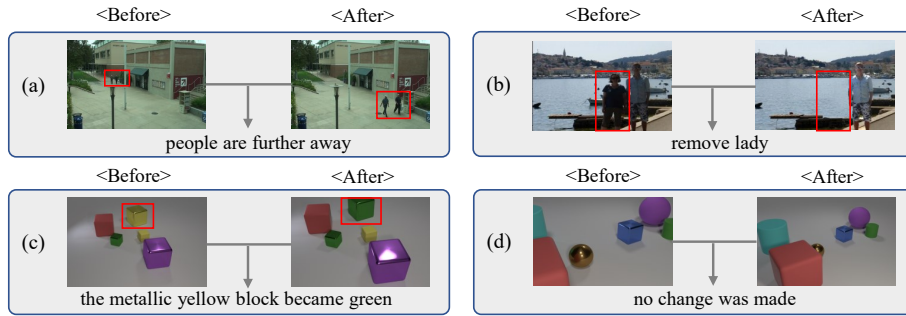
**Keywords:** Change Captioning · Distractors-Immune Representation Learning · Cross-modal Contrastive Regularization

## 1 Introduction

Recently, researchers have witnessed the great success toward vision-language understanding and generation [21, 30, 43, 44]. As an emerging task, change captioning [27, 36, 46] is to describe what has semantically changed between two similar images in natural language (Fig. 1 (a)-(d)). This task has a wide range of practical applications, such as generating reports for surveillance area changes [10]

---

\* Corresponding authors



**Fig. 1:** The examples of change captioning under different scenarios. The first and second cases show that object moving and dropping, respectively. The third one shows the color change under distractors (viewpoint and illumination changes), where the real change is overwhelmed by pseudo changes. The last one shows with only distractors. Changed objects are shown in red boxes.

and pathological changes between medical images [17], as well as providing visually-impaired users with the explanations of image editing effects [29].

This task also poses a formidable challenge: models should be powerful enough to not only understand the contents of two images, but also describe the semantic change between them, while being immune to distractors (viewpoint and illumination changes). In a dynamic environment, two images of the same scene are usually obtained in the presence of distractors. In this situation, unchanged objects between two images often appear obviously pseudo changes about the scale and location (Fig. 1 (c) (d)), where the features of same objects within the image pair might be perturbational. Especially, as a viewpoint drastically changes, certain objects might partially overlap the others, *e.g.*, the brown ball is partially occluded by the red block in the “after” image of Fig. 1 (d). Accordingly, the feature discrimination might be weakened under this viewpoint. In short, there are perturbational and discrimination-degraded features between two image representations under distractors.

Most previous methods directly subtract between such two image representations [9, 24, 28, 38] or compute their feature similarity [25, 26, 45, 47], which risk capturing error-prone difference features in between. The latest works SCORER [35] and SMART [34] overcome the distractors by implementing contrastive learning between similar/dissimilar image pairs. By maximizing the alignment of similar ones, both methods make the features of unchanged objects non-perturbational under distractors. However, both methods disregard the feature discrimination-degraded problem, leaving the influence of distractors upon a pair of image representations remained.

In this paper, we propose a **Distractors-Immune Representation Learning** (DIRL) network to attain a pair of stable image representations that are non-perturbational and discriminative under distractors, for robust change captioning. Concretely, given two raw image representations, DIRL first computes a

channel correlation matrix between them. Next, DIRL implements cross-channel decorrelation to optimize this matrix to be close to the identity matrix, *i.e.*, making the corresponding channels of two image representations have similar semantics, while enforcing different channels to be independent. In this self-supervised fashion, DIRL not only alleviates feature perturbation between two images, but also enhances feature discrimination for each image. As such, the model can sufficiently interact the two stable image representations to infer their reliable difference features, which are then translated into a linguistic sentence via a transformer decoder.

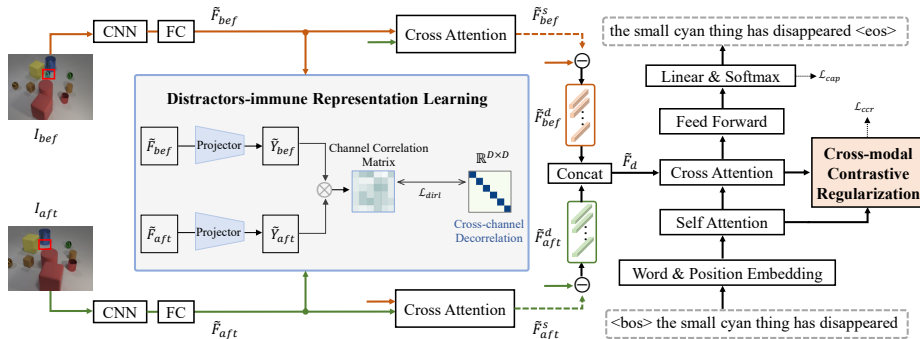
During sentence generation, as the changed object typically occurs in a local region with weak feature, traditional attention mechanisms used in mainstream methods [14, 25, 31] may lead to less satisfactory cross-modal alignment. To facilitate correct alignment, we further design a **Cross-modal Contrastive Regularization (CCR)**, which is built upon the cross-attention module of transformer decoder. When the cross-attention module yields the attended difference features, CCR further regularizes them by maximizing the contrastive alignment between them and generated words, which helps the decoder generate the change caption based on the most related difference features.

**Our key contributions are:** **(1)** We propose the novel DIRL network to learn a pair of stable image representations under distractors by correlating their corresponding channels and decorrelating different channels. Based on the two distractors-immune representations, the model is able to learn the robust difference features between them for caption generation. **(2)** We design the CCR to regularize the cross-modal alignment by maximizing the contrastive alignment between the features of attended difference and generated words. This helps the decoder generate words based on the most related difference features. **(3)** Our method performs favourably against state-of-the-art methods on four public datasets with different change scenarios.

## 2 Related Work

Change captioning is a new task in the vision-and-language area [4, 6, 12, 16, 32, 39]. The pioneer work [13] releases the dataset about changes in surveillance scenarios with potential illumination changes. Afterwards, Tan *et al.* [29] collect a dataset about image editing scenes. Both works [14, 24] develop two datasets for simulating distractors (viewpoint and illumination changes), where the unchanged objects' scale and location would illustrate obviously pseudo changes.

To learn the difference features under distractors, prior works [9, 20, 24] directly subtract between two image representations, which are hard to generalize to unaligned image pairs. SGCC [18] introduces 3D information of object depths to overcome viewpoint changes. Recent works [14, 26, 31, 47] first measure the feature similarity to summarize the shared features and then remove them to obtain difference features. On the other hand, the works [7, 25, 45] correlate the similar features of two images to implicitly deduce the difference features. However, the above methods do not well address the influence of distractors upon image rep-



**Fig. 2:** The framework of our method, where the core blocks are **distractors-immune representation learning** network and **cross-modal contrastive regularization**. FC and Concat are short for the fully-connect layer and concatenation operation.

representations. The latest works SCORER [35] and SMART [34] try contrastive learning between similar/dissimilar image pairs, which model the alignment between similar ones to handle distractors. There are two differences between both works and our DIRL. First, SCORER and SMART capture relations between all paired/unpaired images in a batch, whereas DIRL computes the correlation between the feature channels of each image pair. Second, SCORER and SMART only make the features between two images non-perturbational under distractors, while DIRL additionally considers enhancing feature discrimination for each image to better recognize semantic changes.

During caption generation, since the changed object typically appears in a local region with weak feature, it is difficult to learn reliable alignment via current attention modules in prevalent methods. To address this issue, SMART [34] employs part-of-speech to guide the decoder to dynamically use visual information. These works [9, 14, 35] propose a cycle consistence module to enforce correct cross-modal alignment. The other works [3, 7, 45] design a pre-training and fine-tuning strategy to strengthen the alignment. In this paper, we regularize cross-modal alignment by proposing CCR, which maximizes the contrastive alignment between the generated words and attended difference features. We find that two works in single-image captioning [6, 44] tried the similar idea. Both works model the coarse-grained cross-modal contrastive regularization between the input image and generated caption. Instead, our CCR reuses the attended difference features; establishes the fine-grained contrastive regularization between them and generated words during decoding. In this way, the model can use the most relevant difference features to predict next words.

### 3 Methodology

As shown in Fig. 2, the overall framework of our method includes following parts. First, a pre-trained CNN model extracts the raw representations of an

image pair, which are then transformed into a low-dimensional space. Next, the two image representations are fed into the proposed DIRL to make them non-perturbational and discriminative under distractors. Subsequently, a cross-attention module interacts the two stable representations to mine the shared features, which are removed to learn the difference features for decoding into a sentence. During sentence generation, the designed CCR regularizes the cross-modal alignment, so as to help the decoder generate the sentence based on the most related difference features.

### 3.1 Image Pair Feature Extraction

Given a pair of images “before”  $I_{bef}$  and “after”  $I_{aft}$ , we first use a pre-trained CNN to extract their feature representations. They are denoted as  $F_{bef}$  and  $F_{aft}$ , where  $F_o \in \mathbb{R}^{C \times H \times W}$ .  $C$ ,  $H$ ,  $W$  indicate the number of channels, height, and width. Then, we project them into a low-dimensional space of  $\mathbb{R}^D$  by a 2D-convolution:

$$\tilde{F}_o = \text{conv}_2(F_o) + \text{pos}(F_o), \quad (1)$$

where  $o \in (bef, aft)$ .  $\text{pos}$  is the learnable position encoding function.

### 3.2 Distractors-immune Representation Learning

After obtaining the two position-embedded image representations, we propose a DIRL network to make them non-perturbational and discriminative under distractors in a self-supervised manner. Concretely, we first project  $\tilde{F}_{bef}$  and  $\tilde{F}_{aft}$  into a common embedding space with shared parameters:

$$\tilde{Y}_o = \text{MLP}(\tilde{F}_o), o \in (bef, aft), \quad (2)$$

where MLP is a two-layer multi-layer perceptron with the ReLU activation function in between. Then, we compute a channel correlation matrix between  $\tilde{Y}_{bef}$  and  $\tilde{Y}_{aft}$  along the batch dimension:

$$\mathcal{C}_{ij} = \frac{\sum_b \tilde{y}_{b,i}^{bef} \tilde{y}_{b,j}^{aft}}{\sqrt{\sum_b (\tilde{y}_{b,i}^{bef})^2} \sqrt{\sum_b (\tilde{y}_{b,j}^{aft})^2}}, \quad (3)$$

where  $b$  indexes batch samples and  $i, j$  index the channel dimension of features.  $\mathcal{C} \in \mathbb{R}^{D \times D}$  is a channel correlation matrix. Inspired by the recent self-supervised learning methods [42, 48], we perform cross-channel decorrelation to enforce this matrix to be close to the identity matrix, which is implemented by using the  $\ell_2$ -norm minimization:

$$\mathcal{L}_{dirl} = \sum_i (1 - \mathcal{C}_{ii})^2 + \alpha \sum_i \sum_{j \neq i} \mathcal{C}_{ij}^2. \quad (4)$$

The first term equates the diagonal of  $\mathcal{C}$  to one, *i.e.*, the corresponding channels of two image representations will be correlated and thus have similar semantics

under distractors. The second term equates the off diagonal of  $\mathcal{C}$  to zero, *i.e.*, the different channels will be decorrelated. This enhances the discrimination of each image representation.  $\alpha$  is a trade-off parameter to balance the importance between two terms, which is discussed in the supplementary material.

### 3.3 Difference Representation Learning

The difference features are modeled based on the two distractors-immune image representations. As shown in Fig. 1, most objects are identical between two images. Hence, it is natural to extract their shared features to infer the difference features. For a powerful model, it should capture all potential changes w.r.t. both images. To this end, we first learn the difference features for each image and then combine them to construct the difference features between two images.

Specifically, we first compute the shared features on each image by the multi-head cross-attention (MHCA) mechanism [40]:

$$\begin{aligned}\tilde{F}_{bef}^s &= \text{MHCA} \left( \tilde{F}_{bef}, \tilde{F}_{aft}, \tilde{F}_{aft} \right), \\ \tilde{F}_{aft}^s &= \text{MHCA} \left( \tilde{F}_{aft}, \tilde{F}_{bef}, \tilde{F}_{bef} \right).\end{aligned}\tag{5}$$

Subsequently, we subtract each from the corresponding image representation to compute the difference features on each image, respectively:

$$\begin{aligned}\tilde{F}_{bef}^d &= \tilde{F}_{bef} - \tilde{F}_{bef}^s, \\ \tilde{F}_{aft}^d &= \tilde{F}_{aft} - \tilde{F}_{aft}^s.\end{aligned}\tag{6}$$

Both  $\tilde{F}_{bef}^d$  and  $\tilde{F}_{aft}^d$  are then concatenated as the omni-representation of difference features between two images, which is implemented by a fully-connected layer with the ReLU activation function:

$$\tilde{F}_d = \text{ReLU} \left( \left[ \tilde{F}_{bef}^d; \tilde{F}_{aft}^d \right] W_c + b_c \right),\tag{7}$$

where  $[\cdot]$  is a concatenation operation.

### 3.4 Caption Generation

After learning  $\tilde{F}_d \in \mathbb{R}^{HW \times D}$ , we use a standard transformer decoder [40] to translate it into a sentence. First, we obtain the embedding features of  $m$  words (ground-truth words during training, predicted words during inference). Then, we use the multi-head self-attention (MHSA) layer to model relationships among these word embedding features  $E[W] = \{E[w_1], \dots, E[w_m]\}$ :

$$\hat{E}[W] = \text{MHSA} (E[W], E[W], E[W]).\tag{8}$$

Subsequently, we use the relation-embedded word features  $\hat{E}[W]$  to attend to the related difference features from  $\tilde{F}_d$  based on the MHCA layer:

$$\hat{V} = \text{MHCA} (E[\hat{W}], \tilde{F}_d, \tilde{F}_d).\tag{9}$$

Next, the  $\hat{V} \in \mathbb{R}^{HW \times D}$  is passed to a feed-forward network (FFN) to obtain the enhanced difference features:

$$\hat{V}' = \text{LayerNorm}(\hat{V} + \text{FFN}(\hat{V})). \quad (10)$$

Finally, the probability distributions of target words are calculated via a single hidden layer:

$$W = \text{Softmax}(\hat{V}'W_h + b_h), \quad (11)$$

where  $W_h \in \mathbb{R}^{D \times \nu}$  and  $b_h \in \mathbb{R}^\nu$  are the learnable parameters.  $\nu$  is the dimension of vocabulary size.

### 3.5 Cross-modal Contrastive Regularization

It is often the case that the changed object appears in a local area with weak feature, so it is difficult for the cross-attention module to directly align the captured difference features with word features. To this end, we devise the CCR to regularize the cross-modal alignment, which reuses the attended difference features  $\hat{V}$  computed in Eq. (9) and maximizes the contrastive alignment between them and generated words  $\hat{E}[W]$ . Specifically, we first compute the global representation for  $\hat{E}[W]$  and  $\hat{V}$  via mean-pooling operation, respectively:

$$\begin{aligned} \tilde{E}[W] &= \frac{1}{m} \hat{E}[W], \\ \tilde{V} &= \frac{1}{HW} \hat{V}, \end{aligned} \quad (12)$$

where  $m$  and  $HW$  are the length of  $\hat{E}[W]$  and  $\hat{V}$ . Then, given a training batch, we sample  $B$  pairs of global representations of generated words and attended difference features. For  $k$ -th global word representation  $\tilde{E}[W]_k$ ,  $k$ -th global difference representation  $\tilde{V}_k$  is its positive, while the others will be the negatives in this batch. The similarity between two global representations is computed by dot-product:

$$\text{sim}(\tilde{E}[W], \tilde{V}) = \tilde{E}[W] \cdot \tilde{V}^\top, \quad (13)$$

Next, we leverage the InfoNCE loss [22] to pull semantically close pairs of  $\tilde{E}[W]_k$  and  $\tilde{V}_k$  together and push away non-related pairs:

$$\begin{aligned} \mathcal{L}_{t2v} &= -\frac{1}{B} \sum_k \log \frac{e^{(\text{sim}(\tilde{E}[W]_k, \tilde{V}_k)/\tau)}}{\sum_r e^{(\text{sim}(\tilde{E}[W]_k, \tilde{V}_r)/\tau)}}, \\ \mathcal{L}_{v2t} &= -\frac{1}{B} \sum_k \log \frac{e^{(\text{sim}(\tilde{V}_k, \tilde{E}[W]_k)/\tau)}}{\sum_r e^{(\text{sim}(\tilde{V}_k, \tilde{E}[W]_r)/\tau)}}, \\ \mathcal{L}_{ccr} &= \frac{1}{2}(\mathcal{L}_{t2v} + \mathcal{L}_{v2t}), \end{aligned} \quad (14)$$

where  $\tau$  is the temperature hyper-parameter. This loss formulates a self-supervisory signal to regularize the cross-modal alignment during caption generation, so as to improve the quality of generated captions.

### 3.6 Joint Training

We train the overall architecture end-to-end by maximizing the likelihood of the observed word sequence. Given the ground-truth words  $(w_1^*, \dots, w_m^*)$ , we minimize the negative log-likelihood loss:

$$\mathcal{L}_{cap}(\theta) = - \sum_{t=1}^m \log p_{\theta}(w_t^* | w_{<t}^*), \quad (15)$$

where  $p_{\theta}(w_t^* | w_{<t}^*)$  is computed by Eq. (11), and  $\theta$  are all the learnable parameters. In addition, our method is self-supervised by the losses of DURL and CCR. Thus, the total loss is defined as:

$$\mathcal{L} = \mathcal{L}_{cap} + \lambda_d \mathcal{L}_{dirl} + \lambda_c \mathcal{L}_{ccr}, \quad (16)$$

where  $\lambda_d$  and  $\lambda_c$  are the trade-off parameters that are discussed in Sec. 4.6.

## 4 Experiments

### 4.1 Datasets and Evaluation Metrics

**Spot-the-Diff** [13] has 13,192 image pairs from surveillance cameras, where each image pair includes an underlying illumination change. According to the official split, we split it into training, validation, and testing with a ratio of 8:1:1.

**CLEVR-Change** [24] has 79,606 image pairs and 493,735 captions. It contains five change types, *i.e.*, “Color”, “Texture”, “Add”, “Drop”, and “Move”, as well as moderate distractors. We use the official split of 67,660 for training, 3,976 for validation, and 7,970 for testing, respectively.

**CLEVR-DC** [14] has 48,000 image pairs with the same change types as CLEVR-Change, but this dataset includes drastic distractors. We use the official split with 85% for training, 5% for validation, and 10% for testing.

**Image Editing Request** [29] is comprised of 3,939 image pairs with 5,695 editing instructions. We use the official split with 3,061 image pairs for training, 383 for validation, and 495 for testing, respectively.

**Evaluation Metrics:** We follow the state-of-the-art methods to use the following five metrics for evaluating the quality of generated captions: BLEU-4 [23], METEOR [2], ROUGE-L [19], CIDEr [41] and SPICE [1]. We compute all the results by the Microsoft COCO evaluation server [5].

### 4.2 Implementation Details

For fair-comparison, we follow current SOTA methods to utilize a pre-trained ResNet-101 [8] to extract the features of a pair of images, with the dimension of  $1024 \times 14 \times 14$ . We project them into a lower dimension of 512. We set the hidden size of the model and word embedding size as 512 and 300. Temperature  $\tau$  in Eq. (14) is set to 0.5. We train the model to converge with 10K iterations in total. Adam optimizer [15] is employed to minimize the negative log-likelihood loss of Eq. (16). During inference, we use the greedy decoding strategy to generate captions. More details are shown in the supplementary material.



**Table 1:** Comparison with SOTA methods on the Spot-the-Diff dataset.

Model	BLEU-4	METEOR	ROUGE-L	CIDEr	SPICE
M-VAM [26] (ECCV 2020)	10.1	12.4	31.3	38.1	14.0
DUDA+TIRG [9] (CVPR 2021)	8.1	12.5	29.9	34.5	-
R <sup>3</sup> Net+SSP [37] (EMNLP 2021)	-	13.1	<u>32.6</u>	36.6	18.8
VACC [14] (ICCV 2021)	9.7	12.6	32.1	41.5	-
MCCFormers-D [25] (ICCV 2021)	10.0	12.4	-	<b>43.1</b>	18.3
IFDC [11] (TMM 2022)	8.7	11.7	30.2	37.0	-
I3N-TD [47] (TMM 2023)	<b>10.3</b>	13.0	31.5	<u>42.7</u>	18.6
VAR-Trans [31] (TIP 2023)	-	12.5	29.3	30.3	17.3
SCORER+CBR [35] (ICCV 2023)	<u>10.2</u>	13.2	-	38.9	18.4
SMART [34] (TPAMI 2024)	-	<u>13.5</u>	31.6	39.4	<u>19.0</u>
<b>DIRL+CCR (Ours)</b>	<b>10.3</b>	<b>13.8</b>	<b>32.8</b>	40.9	<b>19.9</b>

**Table 2:** Comparison with SOTA methods on the CLEVR-Change dataset.

Model	BLEU-4	METEOR	ROUGE-L	CIDEr	SPICE
DUDA [24] (ICCV 2019)	42.9	29.7	-	94.6	19.9
DUDA+TIRG [9] (CVPR 2021)	49.9	34.3	65.4	101.3	27.9
MCCFormers-D [25] (CVPR 2021)	53.3	37.1	70.8	119.1	30.4
R <sup>3</sup> Net+SSP [37] (EMNLP 2021)	52.7	36.2	69.8	116.6	30.3
IFDC [11] (TMM 2022)	47.2	29.3	63.7	105.4	-
I3N [47] (TMM 2023)	53.1	37.0	70.8	117.0	32.1
NCT [33] (TMM 2023)	53.1	36.5	70.7	118.4	30.9
VAR-Trans [31] (TIP 2023)	53.6	36.7	71.0	119.1	30.5
SCORER+CBR [35] (ICCV 2023)	<u>54.4</u>	<u>37.6</u>	71.7	<u>122.4</u>	31.6
SMART [34] (TPAMI 2024)	54.3	37.4	<u>71.8</u>	<b>123.6</b>	<b>32.0</b>
<b>DIRL+CCR (Ours)</b>	<b>54.6</b>	<b>38.1</b>	<b>71.9</b>	<b>123.6</b>	<u>31.8</u>

### 4.3 Performance Comparison

**Results on the Spot-the-Diff Dataset.** This dataset contains image pairs from surveillance cameras with underlying illumination changes. We compare DIRL+CCR with the state-of-the-art methods: M-VAM [26], DUDA+TIRG [9], R<sup>3</sup>Net+SSP [37], VACC [14], MCCFormers-D [25], IFDC [11], I3N-TD [47], VAR-Trans [31], SCORER+CBR [35], and SMART [34].

The results are shown in Table 1. Our DIRL+CCR obtains the best results on most metrics. Compared with the latest works SCORER+CBR and SMART, DIRL+CCR surpasses them on all metrics, especially with increases of 8.2% and 4.7% on SPICE. As the image pairs on this dataset are from surveillance cameras, most objects have no good postures, which makes the model difficult to locate and describe changed objects. In this situation, our DIRL+CCR also achieves an encouraging performance, which shows its robustness.

**Results on the CLEVR-Change Dataset.** This dataset contains basic geometric objects and moderate distractors. We evaluate DIRL+CCR under

**Table 3:** Evaluation on CLEVR-Change with varied change categories by METEOR.

Model	METEOR				
	Color	Texture	Add	Drop	Move
DUDA [24] (ICCV 2019)	32.8	27.3	33.4	31.4	23.5
DUDA+TIRG [9] (CVPR 2021)	36.1	30.4	37.8	36.7	27.0
R <sup>3</sup> Net+SSP [37] (EMNLP 2021)	38.9	35.5	38.0	37.5	30.9
IFDC [11] (TMM 2022)	33.1	27.9	36.2	31.4	31.2
I3N [47] (TMM 2023)	39.9	36.7	<u>39.9</u>	<b>38.1</b>	30.6
NCT [33] (TMM 2023)	39.1	36.3	39.0	37.2	30.5
SMART [34] (TPAMI 2024)	<u>40.2</u>	<u>37.8</u>	39.3	<b>38.1</b>	<u>31.5</u>
<b>DIRL+CCR (Ours)</b>	<b>40.7</b>	<b>38.2</b>	<b>40.0</b>	<u>37.9</u>	<b>33.5</b>

**Table 4:** Comparison with SOTA methods on the CLEVR-DC dataset.

Model	BLEU-4	METEOR	ROUGE-L	CIDEr	SPICE
DUDA [24] (ICCV 2019)	40.3	27.1	-	56.7	16.1
M-VAM [26] (ECCV 2020)	40.9	27.1	-	60.1	15.8
VACC [14] (ICCV 2021)	45.0	29.3	-	71.7	<b>17.6</b>
NCT [33] (TMM 2023)	47.5	<u>32.5</u>	65.1	76.9	15.6
VAR-Trans [31] (TIP 2023)	48.3	32.4	-	77.6	15.4
SCORER+CBR [35] (ICCV 2023)	<u>49.4</u>	<b>33.4</b>	<u>66.1</u>	<u>83.7</u>	16.2
<b>DIRL+CCR (Ours)</b>	<b>51.4</b>	32.3	<b>66.3</b>	<b>84.1</b>	<u>16.8</u>

the setting of semantic change & distractors and detailed change categories. The comparison state-of-the-art methods are: DUDA [24], DUDA+TIRG [9], R<sup>3</sup>Net+SSP [37], MCCFormers-D [25], IFDC [11], I3N [47], VAR-Trans [31], SCORER+CBR [35], and SMART [34].

The results are shown in Table 2-3. In the both tables, our DIRL+CCR gains the superior results under the setting of semantic change & distractors, and obtains best results on most change categories. The recent work NCT [33] is a match-based method under the architecture of transformer, which directly correlates two image representations to extract the shared features and then removes them to compute the difference features. We find that DIRL+CCR outperforms it on all metrics. This comparison validates the effectiveness of performing distractors-immune representation learning before change localization, and cross-modal alignment regularizing during caption generation.

**Results on the CLEVR-DC Dataset.** The experiment is conducted on CLEVR-DC with extreme distractors. We compare with the following state-of-the-art methods: DUDA [24], M-VAM [26], VACC [14], NCT [33], VAR-Trans [31], and SCORER+CBR [35]. The results are shown in Table 4.

We note that the overall performance of our DIRL+CCR is better. Specifically, DIRL+CCR obtains the best results on BLEU-4, ROUGE-L, and CIDEr metrics, while achieving comparable performance against the state-of-the-art

**Table 5:** Comparison with the SOTA methods on Image Editing Request. “\*” indicates that the model is trained by pre-training and fine-tuning strategy.

Model	BLEU-4	METEOR	ROUGE-L	CIDEr
VIXEN-QFormer* [3] (AAAI 2024)	7.9	14.4	33.5	35.4
VIXEN-CLIP* [3] (AAAI 2024)	8.6	<b>15.4</b>	<b>42.5</b>	<b>38.1</b>
DUDA [24] (ICCV 2019)	6.5	12.4	37.3	22.8
Dyn rel-att [29] (ACL 2019)	6.7	12.8	37.5	26.4
BDLSCR [45] (IJIS 2022)	6.9	14.6	38.5	27.7
NCT [33] (TMM 2023)	8.1	<u>15.0</u>	38.8	34.2
VARD-Trans [31] (TIP 2023)	10.0	14.8	39.0	<u>35.7</u>
SCORER+CBR [35] (ICCV 2023)	10.0	<u>15.0</u>	<u>39.6</u>	33.4
SMART [34] (TPAMI 2024)	<u>10.5</u>	<b>15.2</b>	39.1	<b>37.8</b>
<b>DIRL+CCR (Ours)</b>	<b>10.9</b>	<u>15.0</u>	<b>41.0</b>	34.1

methods on the other metrics. Especially, DIRL+CCR obtains relative improvement of 8.2%, 6.4% and 4.0% on BLEU-4 against the recent works NCT, VARD-Trans and SCORER+CBR, respectively. The results validate the robustness of our method under the disturbance of drastic distractors.

**Results on the Image Editing Request Dataset.** To further validate the generalization, we conduct the experiment on the scenario of image editing. The comparison methods are DUDA [24], Dynamic rel-att [29], BDLSCR [45], NCT [33], VARD-Trans [31], SCORER+CBR [35], and SMART [34]. Besides, we compare our method with the latest work VIXEN [3]. VIXEN is first pre-trained on a large-scale image editing dataset (which has not been released yet) and fine-tuned on the Image Editing Request dataset.

The results are shown in Table 5. Compared with the end-to-end training methods, our DIRL+CCR achieves the best results on two metrics and is on par with the SOTA methods on other two metrics. Compared with VIXEN, our DIRL+CCR also achieves a competitive performance. The comparison results validate the effectiveness and generalization of our method.

**Performance Analysis.** In brief, compared to the state-of-the-art methods, our method achieves superior results in different scenarios. This benefits from that 1) DIRL can make the paired image representations non-perturbational and discriminative under distractors. This helps learn robust difference features for caption generation. 2) CCR regularizes the cross-modal alignment by maximizing the contrastive alignment between the features of generated words and attended difference, thus helping improve the quality of yielded captions.

#### 4.4 Ablation Study

**Ablation study of each module.** To figure out the contribution of the proposed DIRL and CCR, we conduct the ablation study on the CLEVR-DC dataset (with extreme distractors). The baseline model is a vanilla Transformer. It directly interacts two image representations to extract share features, which are

**Table 6:** Ablation study about DIRL and CCR on the CLEVR-DC dataset.

Ablation	DIRL	CCR	BLEU-4	ROUGE-L	CIDEr	SPICE
Transformer	×	×	48.9	65.6	79.6	15.7
Transformer	✓	×	50.5	65.8	81.8	16.2
Transformer	×	✓	49.3	65.5	82.7	16.4
Transformer	✓	✓	<b>51.4</b>	<b>66.3</b>	<b>84.1</b>	<b>16.8</b>


**Table 7:** Ablation study about DIRL on the CLEVR-DC dataset.

Ablation	non-perturbation	discrimination	BLEU-4	ROUGE-L	CIDEr	SPICE
Transformer	×	×	48.9	65.6	79.6	15.7
DIRL	✓	×	<b>50.7</b>	65.2	79.3	<b>16.2</b>
DIRL	×	✓	50.1	65.1	79.5	<b>16.2</b>
DIRL	✓	✓	50.5	<b>65.8</b>	<b>81.8</b>	<b>16.2</b>



then removed from image representations to attain the difference features for caption generation. The results shown in Table 6.

When we augment the vanilla Transformer with DIRL, it yields a performance boost. This validates DIRL can learn the two image representations that are non-perturbational and discriminative under the extreme distractors. As such, the model can better mine the shared features to learn reliable difference features. Besides, the performance of vanilla Transformer is enhanced by equipping it with CCR, which verifies that CCR can regularize the cross-modal alignment. By using both DIRL and CCR, the vanilla Transformer achieves best performance, in particular with the improvements of 5.7% and 7.0% on CIDEr and SPICE metrics. This shows that each module not only plays its unique role, but also supplements each other. In Fig. 3, we visualize the change localization and captioning results of DIRL+CCR and vanilla Transformer under distractors, where the unchanged objects appear obviously pseudo changes about scale and location. The visualization indicates that with the help of DIRL and CCR, the model can pinpoint and describe the actually changed object.

**Ablation study of DIRL.** DIRL aims to learn two stable image representations from the perspectives of non-perturbation and discrimination, where the former is achieved by the first term of Eq. (4); the latter is achieved by the second term. To validate our claim, we conduct the ablative study for non-perturbation and discrimination, respectively. The results are shown in Table 7. The performance of DIRL with either characteristic does not obtain much gains against the vanilla Transformer. Our conjecture is that the influence of distractors is not well solved. By jointly modeling the non-perturbation and discrimination, the performance of DIRL is significantly improved, which shows that each characteristic is key to handle distractors. That is, only if the model learns both, can it make the paired image representations stable under distractors. As such, the model can capture the reliable difference features for caption generation.



**Fig. 3:** Visualization of change localization and captioning results of Dirl+CCR and Transformer. GT is short for ground-truth and changed objects are shown in red boxes.



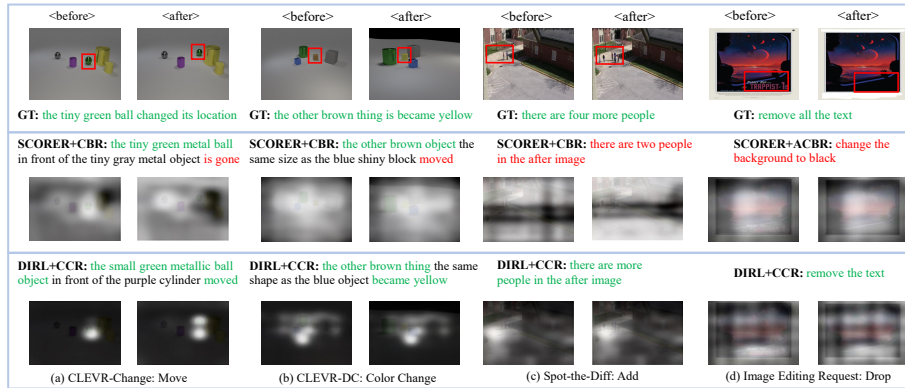
**Fig. 4:** Visualization of captioning performance under varied viewpoint changes. **Fig. 5:** The effects of two trade-off parameters of  $\lambda_d$  and  $\lambda_c$  on CLEVR-DC.

#### 4.5 Measuring Captioning Performance under Distractors

To verify whether our method is immune to distractors, we measure the model’s captioning performance under varied viewpoint changes. The varied viewpoints are computed by the IoU of the bounding boxes of the objects of two images, where the more the camera changes its position, the less the bounding boxes overlap. That is, lower IoU means higher difficulty. The results are shown in Fig. 4, where the compared models are Subtraction (direct subtraction between two image representations), vanilla Transformer, DIRL and DIRL+CCR. We find that the performance of Subtraction is the lowest, showing that direct subtraction generalizes poorly to distractors. DIRL apparently surpasses Transformer, validating it does make the representations of image pair immune to distractors. Further, the performance of DIRL further boosts by augmenting it with CCR, which demonstrates that CCR plays a role in enhancing cross-modal alignment and thus improves the captioning quality under distractors.

#### 4.6 Study on the Trade-off Parameters $\lambda_d$ and $\lambda_c$

We discuss the effects of two trade-off parameters  $\lambda_d$  and  $\lambda_c$  in Eq.(16). Fig. 5 shows the results of DIRL and DIRL+CCR with different  $\lambda_d$  and  $\lambda_c$  on the CLEVR-DC dataset. We first discuss  $\lambda_d$ . As its value increases or decreases, the performance of DIRL changes. Based the results, we set  $\lambda_d$  as 0.03. Then, we fix  $\lambda_d$  to discuss  $\lambda_c$  and set it as 0.05. Similarly, we further discuss  $\lambda_d$  and  $\lambda_c$  on the other three datasets, and set them as 0.5 and 0.3 on the CLEVR-Change dataset; 0.5 and 0.004 on the Spot-the-Diff dataset; 0.001 and 0.05 on the Image Editing Request dataset.



**Fig. 6:** Qualitative analysis on four datasets. We visualize the ground-truth captions (GT), generated captions and change localization results of SCORER+CBR [35] and our DIRL+CCR. The ground-truth changes are shown in red boxes.

#### 4.7 Qualitative Analysis

We conduct qualitative analysis about change localization and caption generation on the four datasets, as shown in Fig. 6 (a)-(d). The compared method is SCORER+CBR [35], which first obtains view-invariant representations and then interacts both to compute the difference features, as well as improving captioning quality by cross-model backward reasoning. In these examples, the change localization maps of SCORER+CBR are relatively scattered and cannot focus on genuine changes. Hence, it fails to generate the correct sentences to describe these changes. By contrast, our DIRL+CCR successfully recognizes and captions the changes. This superiority benefits from making two image representations non-perturbational and discriminative under distractors. As such, our method can better pinpoint the actually changed object between them. During caption generation, CCR further regularizes cross-modal alignment, so as to make the target words generated based on the most related difference features. More qualitative examples are shown in the supplementary material.

## 5 Conclusion

In this paper, we propose the DIRL to learn a pair of non-perturbational and discriminative image representations under distractors, by correlating their corresponding channels and decorrelating different channels. As such, the model can sufficiently mine their shared features to learn the robust difference features for caption generation. Further, we design the CCR to regularize the cross-modal alignment by maximizing the contrastive alignment between the attended difference features and generated words, which helps improve captioning performance. Extensive experiments show that our method yields state-of-the-art results on the four public datasets with different change scenes.

## Acknowledgements

This work was supported in part by National Natural Science Foundation of China: 62322211, 61931008, 62236008, 62336008, U21B2038, 62225207, Fundamental Research Funds for the Central Universities (E2ET1104), “Pioneer” and “Leading Goose” R&D Program of Zhejiang Province (2024C01023).

## References

1. Anderson, P., Fernando, B., Johnson, M., Gould, S.: Spice: Semantic propositional image caption evaluation. In: ECCV. pp. 382–398. Springer (2016)
2. Banerjee, S., Lavie, A.: 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. pp. 65–72 (2005)
3. Black, A., Shi, J., Fai, Y., Bui, T., Collomosse, J.: Vixen: Visual text comparison network for image difference captioning. In: AAAI (2024)
4. Chen, J., Li, L., Su, L., Zha, Z.j., Huang, Q.: Prompt-enhanced multiple instance learning for weakly supervised video anomaly detection. In: CVPR. pp. 18319–18329 (2024)
5. Chen, X., Fang, H., Lin, T.Y., Vedantam, R., Gupta, S., Dollár, P., Zitnick, C.L.: Microsoft coco captions: Data collection and evaluation server. arXiv preprint arXiv:1504.00325 (2015)
6. Cho, J., Yoon, S., Kale, A., Deroncourt, F., Bui, T., Bansal, M.: Fine-grained image captioning with clip reward. In: Findings of NAACL. pp. 517–527 (2022)
7. Guo, Z., Wang, T.J., Laaksonen, J.: Clip4idc: Clip for image difference captioning. In: AACL. pp. 33–42 (2022)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
9. Hosseinzadeh, M., Wang, Y.: Image change captioning by learning from an auxiliary task. In: CVPR. pp. 2725–2734 (2021)
10. Hoxha, G., Chouaf, S., Melgani, F., Smara, Y.: Change captioning: A new paradigm for multitemporal remote sensing image analysis. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–14 (2022)
11. Huang, Q., Liang, Y., Wei, J., Cai, Y., Liang, H., Leung, H.f., Li, Q.: Image difference captioning with instance-level fine-grained feature representation. *IEEE Transactions on Multimedia* **24**, 2004–2017 (2022)
12. Islam, M.M., Ho, N., Yang, X., Nagarajan, T., Torresani, L., Bertasius, G.: Video recap: Recursive captioning of hour-long videos. In: CVPR (2024)
13. Jhamtani, H., Berg-Kirkpatrick, T.: Learning to describe differences between pairs of similar images. In: EMNLP. pp. 4024–4034 (2018)
14. Kim, H., Kim, J., Lee, H., Park, H., Kim, G.: Agnostic change captioning with cycle consistency. In: ICCV. pp. 2095–2104 (2021)
15. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
16. Li, L., Gao, X., Deng, J., Tu, Y., Zha, Z.J., Huang, Q.: Long short-term relation transformer with global gating for video captioning. *IEEE Transactions on Image Processing* **31**, 2726–2738 (2022)

17. Li, M., Lin, B., Chen, Z., Lin, H., Liang, X., Chang, X.: Dynamic graph enhanced contrastive learning for chest x-ray report generation. In: CVPR. pp. 3334–3343 (2023)
18. Liao, Z., Huang, Q., Liang, Y., Fu, M., Cai, Y., Li, Q.: Scene graph with 3d information for change captioning. In: ACM MM. pp. 5074–5082 (2021)
19. Lin, C.Y.: Rouge: A package for automatic evaluation of summaries. In: Text summarization branches out. pp. 74–81 (2004)
20. Liu, C., Zhao, R., Chen, H., Zou, Z., Shi, Z.: Remote sensing image change captioning with dual-branch transformers: A new method and a large scale dataset. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–20 (2022)
21. Liu, X., Li, L., Wang, S., Zha, Z.J., Li, Z., Tian, Q., Huang, Q.: Entity-enhanced adaptive reconstruction network for weakly supervised referring expression grounding. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(3), 3003–3018 (2023)
22. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
23. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: ACL. pp. 311–318 (2002)
24. Park, D.H., Darrell, T., Rohrbach, A.: Robust change captioning. In: ICCV. pp. 4624–4633 (2019)
25. Qiu, Y., Yamamoto, S., Nakashima, K., Suzuki, R., Iwata, K., Kataoka, H., Satoh, Y.: Describing and localizing multiple changes with transformers. In: ICCV. pp. 1971–1980 (2021)
26. Shi, X., Yang, X., Gu, J., Joty, S., Cai, J.: Finding it at another side: A viewpoint-adapted matching encoder for change captioning. In: ECCV. pp. 574–590 (2020)
27. Sun, Y., Qiu, Y., Khan, M., Matsuzawa, F., Iwata, K.: The stvchronos dataset: Towards continuous change recognition in time. In: CVPR. pp. 14111–14120 (2024)
28. Sun, Y., Li, L., Yao, T., Lu, T., Zheng, B., Yan, C., Zhang, H., Bao, Y., Ding, G., Slabaugh, G.: Bidirectional difference locating and semantic consistency reasoning for change captioning. *International Journal of Intelligent Systems* **37**(5), 2969–2987 (2022)
29. Tan, H., Deroncourt, F., Lin, Z., Bui, T., Bansal, M.: Expressing visual relationships via language. In: ACL. pp. 1873–1883 (2019)
30. Tang, W., Li, L., Liu, X., Jin, L., Tang, J., Li, Z.: Context disentangling and prototype inheriting for robust visual grounding. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**(5), 3213–3229 (2024)
31. Tu, Y., Li, L., Su, L., Du, J., Lu, K., Huang, Q.: Viewpoint-adaptive representation disentanglement network for change captioning. *IEEE Transactions on Image Processing* **32**, 2620–2635 (2023)
32. Tu, Y., Li, L., Su, L., Gao, S., Yan, C., Zha, Z.J., Yu, Z., Huang, Q.: I<sup>2</sup>transformer: Intra-and inter-relation embedding transformer for tv show captioning. *IEEE Transactions on Image Processing* **31**, 3565–3577 (2022)
33. Tu, Y., Li, L., Su, L., Lu, K., Huang, Q.: Neighborhood contrastive transformer for change captioning. *IEEE Transactions on Multimedia* **25**, 9518–9529 (2023)
34. Tu, Y., Li, L., Su, L., Zha, Z.J., Huang, Q.: Smart: Syntax-calibrated multi-aspect relation transformer for change captioning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**(7), 4926–4943 (2024)
35. Tu, Y., Li, L., Su, L., Zha, Z.J., Yan, C., Huang, Q.: Self-supervised cross-view representation reconstruction for change captioning. In: ICCV. pp. 2805–2815 (2023)
36. Tu, Y., Li, L., Su, L., Zha, Z.J., Yan, C., Huang, Q.: Context-aware difference distilling for multi-change captioning. In: ACL (2024)



37. Tu, Y., Li, L., Yan, C., Gao, S., Yu, Z.: R<sup>3</sup>Net:relation-embedded representation reconstruction network for change captioning. In: EMNLP. pp. 9319–9329 (2021)
38. Tu, Y., Yao, T., Li, L., Lou, J., Gao, S., Yu, Z., Yan, C.: Semantic relation-aware difference representation learning for change captioning. In: Findings of ACL. pp. 63–73 (2021)
39. Tu, Y., Zhou, C., Guo, J., Li, H., Gao, S., Yu, Z.: Relation-aware attention for video captioning via graph learning. *Pattern Recognition* **136**, 109204 (2023)
40. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: NeurIPS. pp. 5998–6008 (2017)
41. Vedantam, R., Lawrence Zitnick, C., Parikh, D.: Cider: Consensus-based image description evaluation. In: CVPR. pp. 4566–4575 (2015)
42. Wang, Q., Zhang, Y., Zheng, Y., Pan, P., Hua, X.S.: Disentangled representation learning for text-video retrieval. arXiv preprint arXiv:2203.07111 (2022)
43. Xiao, J., Li, L., Lv, H., Wang, S., Huang, Q.: R&b: Region and boundary aware zero-shot grounded text-to-image generation. ICLR (2024)
44. Yang, C.F., Tsai, Y.H.H., Fan, W.C., Salakhutdinov, R.R., Morency, L.P., Wang, F.: Paraphrasing is all you need for novel object captioning. *NeurIPS* **35**, 6492–6504 (2022)
45. Yao, L., Wang, W., Jin, Q.: Image difference captioning with pre-training and contrastive learning. In: AAAI (2022)
46. Yue, S., Tu, Y., Li, L., Gao, S., Yu, Z.: Multi-grained representation aggregating transformer with gating cycle for change captioning. *ACM Trans. Multimedia Comput. Commun. Appl.* (2024)
47. Yue, S., Tu, Y., Li, L., Yang, Y., Gao, S., Yu, Z.: I3n: Intra- and inter-representation interaction network for change captioning. *IEEE Transactions on Multimedia* **25**, 8828–8841 (2023)
48. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow twins: Self-supervised learning via redundancy reduction. In: ICML. pp. 12310–12320 (2021)