

# Dependent Multi-Task Learning with Causal Intervention for Image Captioning

Wenqing Chen<sup>1,2</sup>, Jidong Tian<sup>1,2</sup>, Caoyun Fan<sup>1,2</sup>, Hao He<sup>1,2\*</sup> and Yaohui Jin<sup>1,2\*</sup>

<sup>1</sup>MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

<sup>2</sup>State Key Lab of Advanced Optical Communication System and Network,  
Shanghai Jiao Tong University

{wenqingchen, frank92, fcy3649, hehao, jinyh}@sjtu.edu.cn

## Abstract

Recent work for image captioning mainly followed an extract-then-generate paradigm, pre-extracting a sequence of object-based features and then formulating image captioning as a single sequence-to-sequence task. Although promising, we observed two problems in generated captions: 1) content inconsistency where models would generate contradicting facts; 2) not informative enough where models would miss parts of important information. From a causal perspective, the reason is that models have captured spurious statistical correlations between visual features and certain expressions (e.g., visual features of "long hair" and "woman"). In this paper, we propose a dependent multi-task learning framework with the causal intervention (DMTCI). Firstly, we involve an intermediate task, bag-of-categories generation, before the final task, image captioning. The intermediate task would help the model better understand the visual features and thus alleviate the content inconsistency problem. Secondly, we apply Pearl's do-calculus on the model, cutting off the link between the visual features and possible confounders and thus letting models focus on the causal visual features. Specifically, the high-frequency concept set is considered as the proxy confounders where the real confounders are inferred in the continuous space. Finally, we use a multi-agent reinforcement learning (MARL) strategy to enable end-to-end training and reduce the inter-task error accumulations. The extensive experiments show that our model outperforms the baseline models and achieves competitive performance with state-of-the-art models.

## 1 Introduction

Image captioning is a typical vision-to-language problem, which aims to automatically generate descriptions for images [Chen *et al.*, 2015]. Recently, a line of studies used an extract-then-generate paradigm with a preliminary task, object detection [Anderson *et al.*, 2018]. These studies first extracted

\*Corresponding authors

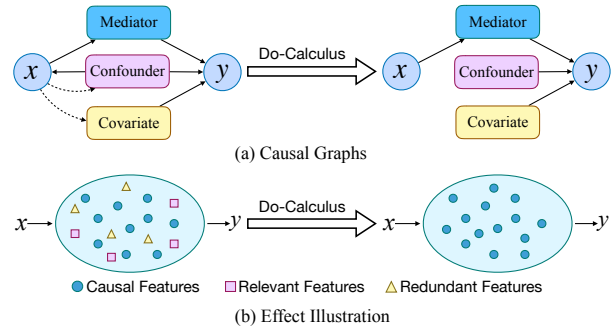


Figure 1: The motivation for our work. (a) A typical causal graph including the mediator, the confounder, and the covariate. According to Pearl's do-calculus on the input  $x$ , the causal graph will be intervened to the right causal graph, which is a robust prediction model. The dashed lines represent the spurious or biased correlations that should not be used for prediction. (b) Illustration of the change of the feature space.

a sequence of object-based region features from the images via a pre-trained detection network, and then used LSTM-based [Hochreiter and Schmidhuber, 1997] or Transformer-based [Vaswani *et al.*, 2017] sequence-to-sequence (seq2seq) models to generate captions. Empirically, the pre-extracted features greatly help models produce high-fidelity captions [Anderson *et al.*, 2018].

Although the paradigm achieved promising performance, we have observed two problems in the generated captions: 1) content inconsistency; 2) not informative enough. Regarding the first problem, given an image containing a "man" with long hair, we found that models often generated "woman" instead, which led to a contradicting fact<sup>1</sup>. Another example for the second problem is that, given an image with the subject-object-predicate triple ("man", "playing", "Wii game"), we found models ignoring "Wii game" and using the quite general words "video game" instead. The generated caption would not be judged wrong but not informative enough.

From a causal perspective, the reason is that the models are inclined to build spurious connections between certain attributes of images and high-frequency categories or pred-

<sup>1</sup>The mentioned image and the tested model can be found in Section 4.4 and Section 4.1, respectively.

icates. For instance, models are possible to capture the spurious statistical correlation between the visual features of "long hair" and "woman" through the confounder "female". Once "long hair" is detected, models may incorrectly predict "man with long hair" as "woman". To further explain this phenomenon, we show a typical causal graph in Figure 1. Let  $x$  and  $y$  be the given image and one ground-truth caption, respectively. We show the directed acyclic graph (DAG) including the mediator, the confounder, and the covariate. A robust model is supposed to utilize the mediator for prediction, which has causal effects on  $y$ , rather than utilizing the dashed lines for prediction.

In this paper, we propose a dependent multi-task framework with the causal intervention (DMTCI) to eliminate the spurious correlations. Firstly, we formulate the image caption as a dependent multi-task (DMT) learning problem [Chen *et al.*, 2020], which includes two tasks: bag-of-categories (BOC) generation and caption generation. The outputs of the first task are considered as the mediators, which are already annotated in certain datasets like MSCOCO and also utilized by recent work [Wang *et al.*, 2020; Li *et al.*, 2020]. Secondly, in terms of the confounders, we are probably to have no access to the real confounders. For example, "female" is not annotated. Fortunately, an available approach is using so-called "proxy confounders" [Louizos *et al.*, 2017]. Since the spurious correlations problem we observe is that the models over-generating high-frequency categories or predicates, we let the proxy confounders be sampled from a set of concepts including high-frequency predicates and fine-grained object categories. Then we apply the variational inference [Kingma and Welling, 2014] to estimate the real confounders and let it sampled from a prior distribution. As a side advantage, our model can also infer covariates from the captions. Finally, we apply a multi-agent reinforcement learning strategy (MARL) to enable end-to-end training and reduce the inter-task error accumulations. An extensive set of experiments demonstrates that our model outperforms the baseline models and achieves competitive performance with those meticulously designed Transformers on the MSCOCO dataset. The main contributions of this work can be summarized as follows:

- We propose a dependent multi-task framework for image captioning, considering the BOC generation as an intermediate task to alleviate the inconsistency problem.
- We further introduce a latent de-confounding approach, which reduces the spurious correlations and estimates richer confounders through the variational inference.
- We cooperate with a multi-agent reinforcement learning approach to better utilize the task dependencies based on our multi-task framework and reduce the inter-task error accumulations.

## 2 Related Work

We categorize recent related work on image captioning into three groups according to the different focus of the work.

**Architectural Innovation.** The first category of work mainly focused on developing novel networks with powerful

representation abilities. Some early studies encoded each image as a single feature vector through a pre-trained convolutional neural network (CNN), and then fed it to recurrent neural networks (RNNs) based on long short-term memory units [Vinyals *et al.*, 2015; Donahue *et al.*, 2015]. Afterward, Xu *et al.* [2015] explored the attention mechanisms for image captioning, which could dynamically focus on salient features at each decoding step. Another important line of work cooperated the task "object detection" for image captioning [Karpathy and Li, 2015; Fang *et al.*, 2015]. Anderson *et al.* [2018] extracted object-level features by a pre-trained Faster R-CNN model, and combined them with the attention mechanism, which achieved significantly higher performance than object-free models. Follow-up researches further captured the relationships among objects by adding self-attention on object-level features [Huang *et al.*, 2019], or using graph convolutional networks [Yao *et al.*, 2018; Yang *et al.*, 2019]. More recently, Transformer-based [Vaswani *et al.*, 2017] models were proposed with spatial position embeddings and better normalization layers [Guo *et al.*, 2020], meshed attentions [Cornia *et al.*, 2020] or higher-order attentions [Pan *et al.*, 2020]. It is worthwhile to mention that almost all the state-of-the-art (SOTA) models utilized the self-critical reinforcement (SCRL) strategy [Rennie *et al.*, 2017].

**Knowledge Enhancement.** The second category of work involved prior knowledge. Kim *et al.* [2019] proposed a relationship-based captioning model that considered part-of-speech (POS, i.e. subject-object-predicate categories) tags of captions, and therefore introduced beneficial inductive bias. Shi *et al.* further explored the semantics conveyed in the subject-object-predicate triples to enhance image caption generation [Shi *et al.*, 2020]. The subject-object-predicate triples are extracted from captions by a textual scene graph parser introduced in [Schuster *et al.*, 2015].

**Causal Intervention.** The third category of work is quite related to our work, which mainly focused on reducing the spurious correlations with causal intervention. Wang *et al.* used causal intervention to better learn the region features of the object detection task [Wang *et al.*, 2020]. When trying to learn the conditional probability between two objects like "toilet" and "person", they assumed other objects as confounders and learned the objective of  $p(\text{toilet}|\text{do}(\text{person}))$  by cutting off the link between other co-occurred objects and the object "person". Yang *et al.* assumed the dataset as the confounder and adjust it as concepts, such as "apple" and "green". Then they cut off the link between these concepts and image features.

## 3 Methodology

In this section, we first formulate the image captioning task as a dependent multi-task learning problem in Section 3.1. Then we illustrate our mediation and de-confounding approach in Section 3.2 and 3.3, with the model details in Section 3.4. Finally, we introduce the multi-agent reinforcement learning approach in Section 3.5.

### 3.1 Dependent Multi-Task Learning

Given the image and one corresponding ground-truth caption  $(\mathbf{x}, \mathbf{y})$ , we use a pre-trained Faster R-CNN network [Anderson *et al.*, 2018] to extract a sequence of features  $\mathbf{R} \in \mathbb{R}^{L_r \times d}$  from  $\mathbf{x}$  where  $L_r$  and  $d$  denote the number of the detected objects and the dimension of features, respectively. The pre-trained Faster R-CNN network is fixed during further training. Recent work mainly formulated image captioning as a single-task learning problem, directly maximizing  $p(\mathbf{y}|\mathbf{R})$  while we formulate image captioning as a DMT problem, which maximizes the log-probability:

$$\log p(\mathbf{y}, \mathbf{m}|\mathbf{R}) = \log p(\mathbf{m}|\mathbf{R}) + \log p(\mathbf{y}|\mathbf{R}, \mathbf{m}) \quad (1)$$

where  $\mathbf{m}$  is an intermediate variable conveying important information to generate final captions. We choose the BOC as  $\mathbf{m}$  which are already annotated in MSCOCO dataset and also utilized in recent work [Wang *et al.*, 2020; Li *et al.*, 2020].

### 3.2 Mediation with DMT

There is a train-test discrepancy in Equation 1. In the training stage, we can use ground-truth object categories to supervise the learning objective  $\log p(\mathbf{m}|\mathbf{R})$  while in the inference stage,  $\mathbf{m}$  is not available. To reduce the train-test discrepancy, we include  $\tilde{\mathbf{m}}$ , the predicted BOC, for further training after pre-training with Equation 1. The learning objective of caption generation becomes:

$$\begin{aligned} p(\mathbf{y}|\mathbf{R}) &= \sum_{\mathbf{m}} p(\mathbf{y}|\mathbf{R}, \mathbf{m})p(\mathbf{m}|\mathbf{R}) \\ &= \mathbb{E}_{[\tilde{\mathbf{m}} \sim p(\mathbf{m}|\mathbf{R})]} [p(\mathbf{y}|\mathbf{R}, \tilde{\mathbf{m}})] \end{aligned} \quad (2)$$

where  $\mathbb{E}$  denotes the expectation operation. Then the multi-task objective is maximizing:

$$\begin{aligned} &\log p(\mathbf{m}|\mathbf{R}) + \log p(\mathbf{y}|\mathbf{R}) \\ &= \log p(\mathbf{m}|\mathbf{R}) + \mathbb{E}_{[\tilde{\mathbf{m}} \sim p(\mathbf{m}|\mathbf{R})]} [\log p(\mathbf{y}|\mathbf{R}, \tilde{\mathbf{m}})] \end{aligned} \quad (3)$$

which utilizes the dependency between tasks and meanwhile reduces the train-test discrepancy. In the period of maximum likelihood estimation (MLE), we use the Gumbel-softmax approach [Jang *et al.*, 2017] when sampling  $\tilde{\mathbf{m}} \sim p(\mathbf{m}|\mathbf{R})$  which can back-propagate gradients between tasks (details can be found in Appendix A in the longer version on Arxiv<sup>2</sup>).

We view Equation 2 from the perspective of causal intervention. A typical causal graph including  $\mathbf{m}$  as a mediator variable is shown in Figure 2. The approach of causal intervention relies on the do-calculus [Pearl, 2010]. Specifically, when applying the do( $\cdot$ ) operation on the input variable  $X$ , we are trying to intervene  $X$  to the observed value  $x$  denoted by do( $X = x$ ). After the do-calculus,  $x$  (represented by  $\mathbf{R}$  in practice) is not determined by other variables. In other words, all the arrows pointing to  $x$  will be cut off while the arrows from  $x$  to other variables remain. It means that, for the mediation causal graph, causal intervention will let  $p(\mathbf{y}|\text{do}(\mathbf{R})) = p(\mathbf{y}|\mathbf{R})$  which does not modify Equation 2.

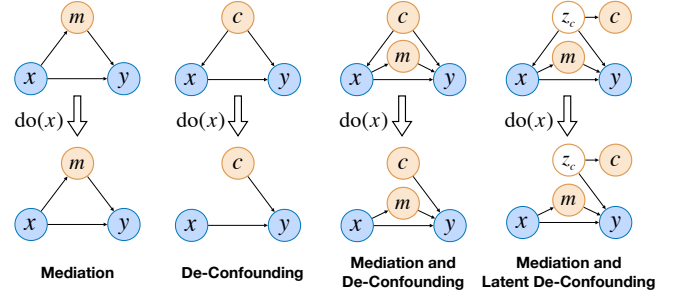


Figure 2: Illustration of the change of causal graphs after the do-calculus, which include the mediator  $\mathbf{m}$ , the proxy confounder  $\mathbf{c}$ , and the latent confounder  $\mathbf{z}_c$ .

### 3.3 Latent De-Confounding

However, for the de-confounding causal graph, the causal intervention will cut off the arrows pointing to  $x$ . We denote the proxy confounder and the real confounder by  $\mathbf{c}$  and  $\mathbf{z}_c$ , respectively. We apply the do-calculus on  $x$  and obtain the de-confounding graph shown at the bottom of Figure 2. Firstly, we only consider  $\mathbf{c}$  and obtain the probability distribution:

$$\begin{aligned} p(\mathbf{y}|\text{do}(\mathbf{R})) &= \sum_{\mathbf{c}} p(\mathbf{y}|\mathbf{R}, \mathbf{c})p(\mathbf{c}|\text{do}(\mathbf{R})) \\ &= \sum_{\mathbf{c}} p(\mathbf{y}|\mathbf{R}, \mathbf{c})p(\mathbf{c}) = \mathbb{E}_{[\tilde{\mathbf{c}} \sim p(\mathbf{c})]} [p(\mathbf{y}|\mathbf{R}, \tilde{\mathbf{c}})] \end{aligned} \quad (4)$$

which makes  $p(\mathbf{y}|\text{do}(\mathbf{R})) \neq p(\mathbf{y}|\mathbf{R})$  as  $p(\mathbf{c}|\text{do}(\mathbf{R})) = p(\mathbf{c})$ . It means that only considering the mediators  $\mathbf{m}$  is not optimal during causal intervention. When jointly considering the mediator and the confounder, we do an intervention so that the mediator  $\mathbf{m}$  is no longer affected by the confounder  $\mathbf{c}$ . Then we have:

$$\begin{aligned} p(\mathbf{y}|\text{do}(\mathbf{R})) &= \sum_{\mathbf{c}} \sum_{\mathbf{m}} p(\mathbf{y}|\mathbf{R}, \mathbf{m}, \mathbf{c})p(\mathbf{m}|\mathbf{R})p(\mathbf{c}) \\ &= \mathbb{E}_{[\tilde{\mathbf{c}} \sim p(\mathbf{c}), \tilde{\mathbf{m}} \sim p(\mathbf{m}|\mathbf{R})]} [p(\mathbf{y}|\mathbf{R}, \tilde{\mathbf{m}}, \tilde{\mathbf{c}})] \end{aligned} \quad (5)$$

according to the mediation-confounding graph in Figure 2.

The next important thing is choosing the proxy confounder  $\mathbf{c}$  and estimating the real confounder  $\mathbf{z}_c$ . Generally, confounders will have causal effects on  $\mathbf{x}$  and  $\mathbf{y}$ . For example, "female" may have causal effects on the visual features of "long hair" in  $\mathbf{x}$  and the word "woman" in  $\mathbf{y}$ . Those models modeling a high posterior probability  $q(\text{female}|\text{long hair})$  may incorrectly take a man with long hair as a woman. In this case, "female" may be a workable confounder but not be annotated. Thanks to the so-called proxy confounder, we can estimate  $\mathbf{z}_c$  from proxy confounders  $\mathbf{c}$  (e.g., "girl"). Since we have observed that models over-generated high-frequency expressions, we assume that the proxy confounder  $\mathbf{c}$  can be sampled from a concept set that includes the high-frequency predicates and fine-grained object categories. We use the summarized predicates and fine-grained categories in [Shi *et al.*, 2020] as the concept set. The real confounders  $\mathbf{z}_c$  has causal effects on  $\mathbf{c}$  and we show the final causal graph of mediation and latent de-confounding in Figure 2.

Estimating  $\mathbf{z}_c$  requires the variational inference [Kingma and Welling, 2014], where the probability distribution  $p(\mathbf{z}_c)$

<sup>2</sup><https://arxiv.org/abs/2105.08573>

is assumed to be priorly defined, such as a zero-mean unit-variance Gaussian distribution. Taking the modeling of the prior distribution  $p(\mathbf{c})$  for example, we can use following approximation:

$$\begin{aligned} \log p_{\vartheta, \phi}(\mathbf{c}) &= \log \int_{\mathbf{z}_c} p_{\vartheta}(\mathbf{c}|\mathbf{z}_c)p(\mathbf{z}_c)d\mathbf{z}_c \\ &\geq \mathbb{E}_{[\mathbf{z}_c \sim q_{\phi}(\mathbf{z}_c|\mathbf{c})]} [\log p_{\vartheta}(\mathbf{c}|\mathbf{z}_c)] - \text{KL}(q_{\phi}(\mathbf{z}_c|\mathbf{c})\|p(\mathbf{z}_c)) \end{aligned} \quad (6)$$

where the integral is approximated by maximizing the evidence lower bound (ELBO), and  $\vartheta$  and  $\phi$  denote the parameters of two networks modeling the prior and the posterior distributions, respectively. And  $\text{KL}(\cdot\|\cdot)$  denotes the Kullback-Leibler divergence between two distributions. Then, we modify Equation 5 when considering  $\mathbf{z}_c$  as the confounder:

$$\begin{aligned} &\log p_{\theta, \phi}(\mathbf{y}|\text{do}(\mathbf{R})) = \\ &\mathbb{E}_{[\tilde{\mathbf{c}} \sim p(\mathbf{c}), \mathbf{z}_c \sim q_{\phi}(\mathbf{z}_c|\tilde{\mathbf{c}}), \tilde{\mathbf{m}} \sim p_{\varphi}(\mathbf{m}|\mathbf{R})]} [\log p_{\theta}(\mathbf{y}|\mathbf{R}, \tilde{\mathbf{m}}, \mathbf{z}_c, \tilde{\mathbf{c}})] \end{aligned} \quad (7)$$

where  $\tilde{\mathbf{c}} \sim p(\mathbf{c})$  means that we randomly sample concepts from the dataset and  $\mathbf{z}_c \sim q_{\phi}(\mathbf{z}_c|\tilde{\mathbf{c}})$  means that we can infer the actually mattered confounders in the latent space. It is worth noting that we are not designing a network to estimate the posterior distribution  $q_{\phi}(\mathbf{z}_c|\mathbf{y}, \mathbf{c})$  because  $\mathbf{y}$  and  $\mathbf{c}$  are not strongly coupled. Instead, as introduced in Equation 7, we model  $p(\mathbf{c})$  in the dataset rather than making a too strong assumption of  $\mathbf{c}$  paired with  $\mathbf{y}$ .

The final multi-task learning objective is maximizing:

$$\log p_{\varphi}(\mathbf{m}|\mathbf{R}) + \log p_{\theta, \phi}(\mathbf{y}|\text{do}(\mathbf{R})) + \log p_{\vartheta, \phi}(\mathbf{c}) \quad (8)$$

where  $\log p_{\theta, \phi}(\mathbf{y}|\text{do}(\mathbf{R}))$  and  $\log p_{\vartheta, \phi}(\mathbf{c})$  are computed through Equation 7 and 6 respectively.

### 3.4 Model Details

To implement the final learning objective, we introduce our model in detail, which can be divided into three main parts: the hierarchical encoder, the BOC generator, and the caption generator.

#### The Hierarchical Encoder

Given the sequence of object features  $\mathbf{R}$ , we use a hierarchical encoder for further encoding, which consists of multiple layers of multi-head attention network followed with "adding and normalization" layers:

$$\mathbf{R}^{(i+1)} = \text{Add\&Norm}(\text{Attn}^{(i)}(\mathbf{R}^{(i)}, \mathbf{R}^{(i)}, \mathbf{R}^{(i)})) \quad (9)$$

where  $i$  denotes the layer number,  $\mathbf{R}^{(0)} = \mathbf{R}$ , and  $\text{Attn}(\cdot)$  is the multi-head attention network introduced in [Vaswani *et al.*, 2017]. We choose the outputs of an intermediate layer  $\mathbf{R}^{(i_m)}$  to predict the BOC  $\mathbf{m}$ , and the outputs of the final layer  $\mathbf{R}^{(n)}$  to predict the caption  $\mathbf{y}$ .

#### The BOC Generator

The first task is a category-level classification task. For each image, we convert  $\mathbf{m}$  to the one-hot label matrix  $\mathbf{Y}^m \in \mathbb{R}^{L_m \times N_m}$  where  $L_m$  denotes the number of all categories, and  $N_m$  denotes the maximum number of detected objects belongs to each category. We use another attention-based network for prediction:

$$\mathbf{H}^c = \text{Attn}(\mathbf{Q}^c, \mathbf{R}^{(i_m)}, \mathbf{R}^{(i_m)}) \quad (10)$$

where  $\mathbf{Q}^c \in \mathbb{R}^{L_m \times d}$  denotes the category-specific query matrix in the attention layer and is trainable. Then we use a fully-connected neural network (FCNN) followed with the softmax function along the last dimension to generate the probabilities of object numbers  $\tilde{\mathbf{Y}}^m$ . Then the loss of the first task is computed through cross-entropy:

$$\mathcal{L}^m = - \sum_{j=1}^{L_m} \mathbf{Y}_j^m \log \tilde{\mathbf{y}}_j^m \quad (11)$$

where  $\mathbf{Y}_i^m$  and  $\tilde{\mathbf{y}}_i^m$  are two one-hot vectors representing the number of detected objects belong to the  $j$ -th category in  $\mathbf{Y}^m$  and  $\tilde{\mathbf{Y}}^m$ , respectively. Maximizing the log probability  $\log p_{\varphi}(\mathbf{m}|\mathbf{R})$  is implemented by minimizing  $\mathcal{L}^m$ .

#### The Caption Generator

In the second task, we mainly introduce how to implement  $p_{\theta}(\mathbf{y}|\mathbf{R}, \tilde{\mathbf{m}}, \mathbf{z}_c, \tilde{\mathbf{c}})$  and  $q_{\phi}(\mathbf{z}_c|\mathbf{c})$ .

Firstly, we need to re-encode the predicted labels  $\tilde{\mathbf{m}}$ . Taking the gold label  $\mathbf{m}$  for example, we involve another two sequences  $\mathbf{w}^m = \{w_j^m\}$  and  $\mathbf{id}^m = \{id_j^m\}$  representing the words and category indexes, respectively. For example,  $\mathbf{w}^m = [\text{person, dining table, pizza, ...}]$  and  $\mathbf{id}^m = [1, 2, 3, \dots]$ , and  $\mathbf{m}$  marks the number of corresponding categories. We encode the words of each category in  $\mathbf{w}^m$  with a LSTM layer:

$$\mathbf{h}_{j,t}^m = \text{LSTM}(\text{OH}(w_{j,t}^m \mathbf{W}_w), \mathbf{h}_{j,t-1}^m) \quad (12)$$

where  $\text{OH}(\cdot)$  denotes the one-hot operation,  $w_{j,t}^m$  denotes the  $t$ -th token in the  $j$ -th category in  $\mathbf{w}^m$ , and  $\mathbf{W}_w$  denotes the word embedding layer. Assuming  $\mathbf{H}^m = \{\mathbf{h}_{j,T_m}^m\}$  containing the hidden states at final tokens (at the  $T_m$ -th step) and  $m_j$  representing the occurrence numbers of  $j$ -th category in the BOC  $\mathbf{m}$ , we fuse three kinds of embeddings as follows:

$$\mathbf{e}_j^m = (\mathbf{h}_{j,T_m}^m + \text{OH}(id_j^m) \mathbf{W}_o) \odot \sigma(\text{OH}(m_j) \mathbf{W}_m) \quad (13)$$

where  $\odot$  denotes the element-wise multiplication and  $\sigma$  denotes the sigmoid function.  $\mathbf{W}_o, \mathbf{W}_m$  denote parameters of the category embedding and number embedding layers, respectively. In this way, the final embedding matrix  $\mathbf{E}^m = \{\mathbf{e}_j^m\} \in \mathbb{R}^{L_m \times d}$  is category-words-aware and number-aware. Then we use the max-pooling operation to obtain a vector  $\mathbf{z}_m$  representing the global information for decoding. It is worth noting that we do not further encode  $\mathbf{E}^m$  with position-aware encoders such as LSTM or attention networks, because the BOC does not contain the detailed information. For example, BOC can not indicate whether a person is hitting a ball or holding a ball. Therefore, we only apply the pooling operation and get  $\mathbf{z}_m$ .

Secondly, we need to infer  $\mathbf{z}_c$  through the network  $q_{\phi}(\mathbf{z}_c|\mathbf{c})$ . We extract the high-frequency predicates and fine-grained categories from  $\mathbf{y}$  and convert them to an unordered set  $\mathbf{c}$ . We embed  $\mathbf{c}$  to  $\mathbf{E}^c \in \mathbb{R}^{L_c \times d}$  through the word embedding layer and pass it to a two-layer Transformer without the position embedding, and obtain a sequence of hidden states  $\mathbf{H}^z \in \mathbb{R}^{L_c \times d}$ . Since  $q_{\phi}(\mathbf{z}_c|\mathbf{c})$  is often approximated to a Gaussian distribution, we use the mean pooling of  $\mathbf{H}^z$  followed with two FCNNs to get the corresponding mean and

variance vectors, and then  $z_c$  is sampled from this Gaussian distribution.

Thirdly, we need to implement  $p_\theta(\mathbf{y}|\mathbf{R}, \tilde{\mathbf{m}}, z_c, \tilde{\mathbf{c}})$  where  $\tilde{\mathbf{m}} \sim p_\varphi(\mathbf{m}|\mathbf{R})$ ,  $z_c \sim q_\phi(z_c|\tilde{\mathbf{c}})$  and  $\tilde{\mathbf{c}} \sim p(\mathbf{c})$ . After randomly selecting  $\tilde{\mathbf{c}}$  from the dataset, we obtain a vector  $z_c$  by sampling from  $q_\phi(z_c|\tilde{\mathbf{c}})$ . Since  $\tilde{\mathbf{m}}$  has been already encoded to  $z_m$ , we simply use a top-down attention LSTM network [Anderson *et al.*, 2018] to decode each token:

$$\mathcal{L}^g = - \sum_{t=1}^T \log \left( p_\theta \left( y_t \mid \mathbf{y}_{<t}, \mathbf{R}^{(n)}, z_m, z_c \right) \right) \quad (14)$$

where  $y_t$  denote the  $t$ -th token of  $\mathbf{y}$ , and  $\mathbf{y}_{<t}$  denotes the sequence of tokens preceding  $y_t$ . The latent variable  $z_m$  and  $z_c$  are concatenated with the hidden states at all steps. Note that our model is also compatible with Transformer-based decoders. But recent empirical work has shown that Transformer-based decoders require meticulous designs, such as meshed attention [Cornia *et al.*, 2020] or high-order attention [Pan *et al.*, 2020], and greatly increase the time complexity when optimized with reinforcement learning (RL).

### 3.5 Multi-Agent Reinforcement Learning

RL is commonly used in the task of image captioning to reduce the train-test discrepancy [Ranzato *et al.*, 2016]. Since we formulate the image captioning as the dependent multi-task problem, the inter-task train-test discrepancy also needs to be tackled. During the MLE optimization period, Gumbel-softmax approximation is used but will improve the gradient variance because of the temperature parameter [Nie *et al.*, 2019]. Therefore, we use a *multi-agent reinforcement learning* strategy to train the model end-to-end based on the SCRL strategy on the two tasks alternatively.

**SCRL for the BOC Generator.** The Monte-Carlo sampling is used to generate  $\tilde{\mathbf{m}}$  and the arg-max operation to generate the baseline  $\bar{\mathbf{m}}$ . The gradient of the BOC generator can be expressed by:

$$\nabla_\varphi \mathcal{L}(\varphi) \approx - (r_1(\tilde{\mathbf{m}}) - r_1(\bar{\mathbf{m}})) \nabla_\varphi \log p_\varphi(\tilde{\mathbf{m}}|\mathbf{R}) \quad (15)$$

where  $r_1(\tilde{\mathbf{m}})$  is the reward provided by both two tasks. After feeding  $\tilde{\mathbf{m}}$  to the caption generator, we obtain  $\tilde{\mathbf{y}}_{\tilde{\mathbf{m}}}$  through greedy decoding and compute the reward  $r_1(\tilde{\mathbf{m}})$  by:

$$r_1(\tilde{\mathbf{m}}) = \lambda_1 s_1(\tilde{\mathbf{m}}, \mathbf{m}) + \lambda_2 s_2(\tilde{\mathbf{y}}_{\tilde{\mathbf{m}}}, \mathbf{y}) \quad (16)$$

where  $\lambda_1$  and  $\lambda_2$  are the hyperparameters to balance the two rewards, and  $s_1(\cdot)$  denotes the score function to match two BOCs, and  $s_2(\cdot)$  denotes the CIDEr-D score function [Vedantam *et al.*, 2015].

**SCRL for the Caption Generator.** Fed with the baseline BOC  $\bar{\mathbf{m}}$  and a sampled latent confounder  $z_c \sim p(z_c)$ , the caption generator can obtain  $\tilde{\mathbf{y}}$  by the Monte-Carlo sampling and  $\bar{\mathbf{y}}$  by the greedy decoding. We use the SCRL strategy with the gradients computed as follows:

$$\nabla_\theta \mathcal{L}(\theta) \approx - (r_2(\tilde{\mathbf{y}}) - r_2(\bar{\mathbf{y}})) \nabla_\theta \log p_\theta(\tilde{\mathbf{y}}|\mathbf{R}, \bar{\mathbf{m}}, z_c) \quad (17)$$

where  $r_2(\tilde{\mathbf{y}}) = s_2(\tilde{\mathbf{y}}, \mathbf{y})$ . Furthermore, we adopt the max-max framework [Xiao *et al.*, 2020] which maximizes the rewards of the two agents alternatively.

## 4 Experiments

### 4.1 Datasets and Settings

**Datasets.** We experiment on the MSCOCO<sup>3</sup> dataset, which is the most popular dataset for image captioning. The original dataset contains about 82,783 training images and 40,504 validation images. Following most of the previous work, we first evaluate our model on the Karpathy data split [Karpathy and Li, 2015] with 5,000 images for validation, 5,000 images for testing, and the rest for training. Then we evaluate our model on the on-line COCO test server. We set the maximum decoding length to 16 which equals the maximum length of 95% of captions. We keep the words that occur no less than 5 times, resulting in a vocabulary of 10,369 words.

**Evaluation and Settings.** We apply five standard metrics for evaluation: CIDEr-D [Vedantam *et al.*, 2015], BLEU [Papineni *et al.*, 2020], METROT [Banerjee and Lavie, 2005], ROUGE [Lin, 2004], and SPICE [Anderson *et al.*, 2016]. All the metrics are computed with the publicly released code.<sup>4</sup>

**Compared Models.** We compare our models with several models including SCST [Rennie *et al.*, 2017], Up-Down [Anderson *et al.*, 2018], AoANet [Huang *et al.*, 2019], Up-Down+VC [Wang *et al.*, 2020], WeakVRD [Shi *et al.*, 2020],  $\mathcal{M}^2$  Transformer [Cornia *et al.*, 2020], and XLAN [Pan *et al.*, 2020]. Among them, XLAN is the recent SOTA with an advanced architecture, and Up-Down+VC is related to our work which uses causal intervention to reduce spurious correlations between objects. Also, we compare with the baseline model, TransLSTM, consisting of a Transformer encoder and an up-down LSTM attention decoder. It is worth noting that one may find that a recent model OSCAR [Li *et al.*, 2020] greatly outperforms the models mentioned above, but we want to clarify that OSCAR made two unfair settings: 1) using more vision-language datasets for pre-training, including Conceptual Captions, SBU captions, Flicker30k, and GQA [Li *et al.*, 2020]; 2) utilizing the ground-truth rather than generated object categories in the inference phase.

**Hyperparameters.** We use the pre-trained Faster R-CNN model [Anderson *et al.*, 2018] to represent each image as an adaptive sequence of feature vectors  $\mathbf{R}$  where the number of detected objects  $n \in [10, 100]$ . The number of encoder layers is 4 and we use the output of the second layer to predict BOC. The dimension of each feature vector  $v_i$  is 2,048 and then projected to 1,024. The dimensions of the word embedding layer, attention layers, and the LSTM-based decoder are all set to 1,024. The learning rate is initialized to 0.0001 and decreased by half when the CIDEr-D score does not increase in 2 epochs, with the minimum learning rate set to  $5 \cdot 10^{-6}$ . The batch size is set to 50. The model is firstly optimized with MLE for 30 epochs (Gumbel sampling  $\tilde{\mathbf{m}}$  after 15 epochs), and then optimized with MARL for another 35 epochs. The special setting of our model is that we can sample different  $z_c$  to generate multiple captions. Therefore, we train an image-text selector to select the best of 20 generated captions for each image (details in Appendix B in the longer version on Arxiv).

<sup>3</sup><http://cocodataset.org/>

<sup>4</sup><https://github.com/tylin/coco-caption>

	Maximum Likelihood Estimation						Reinforcement Learning					
	B1	B4	ME	RG	CD	SP	B1	B4	ME	RG	CD	SP
SCST	-	31.3	26.0	54.3	101.3	-	-	33.3	26.3	55.3	111.4	-
Up-Down	77.2	36.2	27.0	56.4	113.5	20.3	79.8	36.3	27.7	56.9	120.1	21.4
SGAE	77.6	36.9	27.7	57.2	116.7	20.9	80.8	38.4	28.4	58.6	127.8	22.1
AoANet	77.4	37.2	28.4	57.5	119.8	21.3	80.2	38.9	29.2	58.8	129.8	22.4
Up-Down + VC	-	-	-	-	-	-	-	39.5	29.0	59.0	130.5	-
WeakVRD (MT-I)	78.1	<b>38.4</b>	28.2	58.0	119.0	21.1	80.8	38.9	28.8	58.7	129.6	22.3
$\mathcal{M}^2$ Transformer	-	-	-	-	-	-	80.8	39.1	29.2	58.6	131.2	22.6
XLAN †	78.3	38.0	28.7	57.9	121.0	21.8	80.6	39.4	<b>29.5</b>	59.2	131.1	<b>23.1</b>
TransLSTM (Baseline)	77.1	37.1	28.2	57.3	117.3	21.2	80.2	38.6	28.5	58.3	128.7	22.2
DMTCI (Avg.)	78.5	38.0	28.5	58.1	120.2	22.0	81.3	39.7	29.1	59.1	131.7	22.9
DMTCI + Trained Selector (20)	<b>79.1</b>	38.2	<b>28.8</b>	<b>58.2</b>	<b>123.7</b>	<b>22.1</b>	<b>81.3</b>	<b>39.7</b>	29.2	<b>59.2</b>	<b>132.0</b>	23.0
DMTCI + Gold Categories (Avg.)	79.3	38.2	29.0	58.4	122.9	22.2	81.9	40.3	29.5	59.6	135.0	23.1
DMTCI + Optimum Selector (20)	82.6	45.0	31.1	62.2	141.9	23.9	82.7	42.5	30.0	61.0	139.0	23.3
XLAN ‡	78.0	38.2	28.8	58.0	122.0	21.9	80.8	39.5	29.5	59.2	132.0	23.4

Table 1: Single-model performances on the MSCOCO dataset (Karpathy split) in both MLE and RL period. B1, B4, ME, RG, CD, and SP denote BLEU1, BLEU4, METEOR, ROUGE, CIDEr-D, and SPICE, respectively. Most results of SOTA models are cited except that models tagged by † are reproduced with the officially released codes. We also show the officially reported results of XLAN marked with ‡.

Model	B4		ME		RG		CD	
	c5	c40	c5	c40	c5	c40	c5	c40
SCST	35.2	64.5	27.0	35.5	56.3	70.7	114.7	116.0
Up-Down	36.9	68.5	27.6	36.7	57.1	72.4	117.9	120.5
SGAE	38.5	69.7	28.2	37.2	58.6	73.6	123.8	126.5
AoANet	37.3	68.1	28.3	37.2	57.9	72.8	124.0	126.2
Up-Down + VC	37.8	69.1	28.5	37.6	58.2	73.3	124.1	126.2
WeakVRD (MT-I)	<b>38.6</b>	<b>70.1</b>	28.6	37.8	<b>58.8</b>	<b>74.5</b>	125.1	126.7
XLAN †	38.3	69.2	28.5	38.2	58.4	74.0	125.0	127.1
DMTCI	38.5	70.0	<b>29.0</b>	<b>38.4</b>	58.7	74.2	<b>125.1</b>	<b>127.6</b>

Table 2: The single-model performance on COCO online test server.

Period	Ablated Models	B1	B4	ME	RG	CD	SP
MLE	TransLSTM	77.1	37.1	28.2	57.3	117.3	21.2
	Pipeline	77.1	36.7	27.9	56.9	116.1	20.8
	DMTCI-M	77.3	37.8	28.3	57.8	119.5	21.8
	DMTCI-LDC	77.4	37.2	28.1	57.5	118.8	21.6
	DMTCI-Full (Avg.)	78.5	38.0	28.5	58.1	120.2	22.0
	DMTCI-Full (Best)	<b>79.1</b>	<b>38.2</b>	<b>28.8</b>	<b>58.2</b>	<b>123.7</b>	<b>22.1</b>
RL	TransLSTM	80.2	38.6	28.5	58.3	128.7	22.2
	Pipeline	79.9	38.1	28.3	58.0	126.6	22.1
	DMTCI-M	81.1	39.3	28.9	59.0	131.3	22.7
	DMTCI-LDC	80.7	38.7	28.7	58.7	130.1	22.3
	DMTCI-Full (Avg.)	81.3	39.7	29.1	59.1	131.7	22.9
	DMTCI-Full (Best)	<b>81.3</b>	<b>39.7</b>	<b>29.2</b>	<b>59.2</b>	<b>132.0</b>	<b>23.0</b>

Table 3: The performances of ablated models on the Karpathy split in both MLE and RL periods.

## 4.2 Main Results

Table 1 presents the performance on the Karpathy test split of our model as well as the compared models. As shown, our model DMTCI achieves average CIDEr-D scores of 120.2 and 131.7 and outperforms the baseline TransLSTM by 2.9 and 3.0 points in MLE and RL periods, respectively. When compared with the related model with causal intervention, "Up-Down+VC", our model outperforms it by 1.2 points on the CIDEr-D score. Furthermore, when combining with the trained selector, our model outperforms TransLSTM by 6.4

and 3.3 points in MLE and RL periods, respectively. And our model achieves competitive results with the recent SOTA model XLAN, without using meticulously designed attention mechanisms. It is worth noting that we also reproduce XLAN with the officially released codes, because in recent work, a slightly tricky method is to obtain an ensemble model by training the model multiple times. Thus, the single-model performance is reported under the best trial.

If we have the gold BOC, the averaged CIDEr-D score can be further improved to 122.9 and 135.2 in MLE and RL periods respectively. It suggests that future work can focus on improving the performance of the BOC generator. Moreover, if we have an optimal selector, we can further improve the CIDEr-D score to 141.9 in the MLE period even without RL, which can be considered as the upper bound of our model. And RL slightly decreases the upper bound, which means that optimizing with RL drops some useful patterns.

Table 2 presents the single-model results of the models on the online COCO test server. Compared with recent SOTA models which have reported single-model performance, DMTCI achieves competitive performance and achieves the best CIDEr-D score with 40 reference captions.

## 4.3 Ablation Study

To analyze which mechanisms are driving the improvements, we present an ablation study in Table 3. We obtain ablated models including: 1) Pipeline, learning the two tasks sequentially; 2) DMTCI-M, only considering the mediator; 3) DMTCI-LDC, only considering the confounder. To avoid misleading, we denote our full model by DMTCI-Full. Table 3 shows the performance of these ablated models. We can see that Pipeline suffers from the train-test discrepancy and falls behind the baseline model TransLSTM. DMTCI-M increases the CIDEr-D score by 2.2 and 2.6 points over TransLSTM in MLE and RL periods, respectively. DMTCI-LDC outperforms TransLSTM while performing slightly worse than the DMTCI-M. When compared with the two ablated mod-

	TransLSTM	<u>Two women</u> sitting on a couch in a living room.
	DMTCI-M	A couple of people sitting on a couch talking on a laptop.
	DMTCI-Full	<b>A man and a woman</b> sitting on a couch in a living room.
	GT1	A woman sitting on the back of a couch next to a man wearing glasses.
	GT2	A man and a woman sitting on a couch.
	TransLSTM	A woman is <u>blow drying</u> her hair with a <u>baseball bat</u> .
	DMTCI-M	A woman brushing her hair with a brush.
	DMTCI-Full	<b>A young girl</b> brushing her hair with a brush.
	GT1	A kid standing in a room with a brush.
	GT2	A young girl brushing her hair in front of a curtain.
	TransLSTM	A man standing in a living room playing a video game.
	DMTCI-M	A man standing in a living room holding a <u>Wii controller</u> .
	DMTCI-Full	<b>A man</b> is playing a video game with a <u>Wii controller</u> .
	GT1	A man getting into playing the game of Wii.
	GT2	A man that is playing a Wii game.

Figure 3: Case study. The **underlined red** words lead to inconsistent information, and the words with **gray** background are not informative enough. The **blue** words represent consistent and informative expressions. GT1 and GT2 represent two reference sentences.

els, our full model achieves the best CIDEr-D score in both MLE and RL periods. And the selector further brings a large improvement in the MLE period and does not help much in the RL period.

#### 4.4 Case Study

To help understand how the mediators and the confounders affect the learning, we show some cases in Figure 3. The inconsistent descriptions are marked by the red underlines and the not informative enough descriptions are marked by the gray background. The bold blue words represent applaudable expressions that are content consistent and informative. Compared with TransLSTM, DMTCI-M can alleviate the inconsistency problem by involving the mediator. When further combining with the latent de-confounding approach, DMTCI-Full generates consistent and informative captions.

Furthermore, we are trying to discuss why involving the mediators and the confounders work. At present, deep learning models have strong abilities to capture correlations. However, since it is too strong, models can often utilize the co-occurrence of concepts to minimize the training loss, ignoring the exact appearance of the detected objects. Then in the inference stage, it is possible to infer non-exist objects relying on the co-occurrence (e.g., "brush" and "woman"). Involving mediators and confounders during training could alleviate this problem by focusing on the exactly appeared objects.

### 5 Conclusion

In this paper, we propose a dependent multi-task framework with causal intervention for image captioning. We first involve BOC as the mediator to help the model understand the pre-extracted features, then we propose a latent de-confounding approach to alleviate the spurious correlation problem. Furthermore, we apply a multi-agent reinforcement learning strategy to better utilize the task dependencies and enable end-to-end training. The experiments show that our

model yields competitive results with recent SOTA models and the proposed framework could improve content consistency and informativeness.

### Acknowledgments

The authors would like to thank the anonymous reviewers for their constructive comments. This work was supported by the National Key Research and Development Program of China under Grant 2018YFC0830400, and Shanghai Municipal Science and Technology Major Project under Grant 2021SHZDZX0102.

### References

- [Anderson *et al.*, 2016] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. SPICE: semantic propositional image caption evaluation. In *ECCV 2016*, 2016.
- [Anderson *et al.*, 2018] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018.
- [Banerjee and Lavie, 2005] Satanjeev Banerjee and Alon Lavie. METEOR: an automatic metric for MT evaluation with improved correlation with human judgments. In *ACL*, 2005.
- [Chen *et al.*, 2015] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO captions: Data collection and evaluation server. *CoRR*, abs/1504.00325, 2015.
- [Chen *et al.*, 2020] Wenqing Chen, Jidong Tian, Liqiang Xiao, Hao He, and Yaohui Jin. Exploring logically dependent multi-task learning with causal inference. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu, editors, *EMNLP*, 2020.
- [Cornia *et al.*, 2020] Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. Meshed-memory transformer for image captioning. In *CVPR*, 2020.
- [Donahue *et al.*, 2015] Jeff Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Trevor Darrell, and Kate Saenko. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR 2015*, 2015.
- [Fang *et al.*, 2015] Hao Fang, Saurabh Gupta, Forrest N. Iandola, Rupesh Kumar Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C. Platt, C. Lawrence Zitnick, and Geoffrey Zweig. From captions to visual concepts and back. In *CVPR*, 2015.
- [Guo *et al.*, 2020] Longteng Guo, Jing Liu, Xinxin Zhu, Peng Yao, Shichen Lu, and Hanqing Lu. Normalized and geometry-aware self-attention network for image captioning. In *CVPR*, 2020.

- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, 1997.
- [Huang *et al.*, 2019] Lun Huang, Wenmin Wang, Jie Chen, and Xiaoyong Wei. Attention on attention for image captioning. In *ICCV*, 2019.
- [Jang *et al.*, 2017] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *ICLR*, 2017.
- [Karpathy and Li, 2015] Andrej Karpathy and Fei-Fei Li. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015.
- [Kim *et al.*, 2019] Dong-Jin Kim, Jinsoo Choi, Tae-Hyun Oh, and In So Kweon. Dense relational captioning: Triple-stream networks for relationship-based captioning. In *CVPR*, 2019.
- [Kingma and Welling, 2014] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014.
- [Li *et al.*, 2020] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, 2020.
- [Lin, 2004] Chinyew Lin. Rouge: A package for automatic evaluation of summaries. In *ACL*, 2004.
- [Louizos *et al.*, 2017] Christos Louizos, Uri Shalit, Joris M. Mooij, David A. Sontag, Richard S. Zemel, and Max Welling. Causal effect inference with deep latent-variable models. In *NeurIPS*, 2017.
- [Nie *et al.*, 2019] Weili Nie, Nina Narodytska, and Ankit Patel. Relgan: Relational generative adversarial networks for text generation. In *ICLR*, 2019.
- [Pan *et al.*, 2020] Yingwei Pan, Ting Yao, Yehao Li, and Tao Mei. X-linear attention networks for image captioning. In *CVPR*, 2020.
- [Papineni *et al.*, 2020] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2020.
- [Pearl, 2010] Judea Pearl. On measurement bias in causal inference. In Peter Grünwald and Peter Spirtes, editors, *UAI 2010, Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence, Catalina Island, CA, USA, July 8-11, 2010*, pages 425–432. AUAI Press, 2010.
- [Ranzato *et al.*, 2016] Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. Sequence level training with recurrent neural networks. In *ICLR*, 2016.
- [Rennie *et al.*, 2017] Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. In *CVPR*, 2017.
- [Schuster *et al.*, 2015] Sebastian Schuster, Ranjay Krishna, Angel X. Chang, Li Fei-Fei, and Christopher D. Manning. Generating semantically precise scene graphs from textual descriptions for improved image retrieval. In *VL@EMNLP*, 2015.
- [Shi *et al.*, 2020] Zhan Shi, Xu Zhou, Xipeng Qiu, and Xiaodan Zhu. Improving image captioning with better use of captions. In *ACL*, 2020.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [Vedantam *et al.*, 2015] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *CVPR*, 2015.
- [Vinyals *et al.*, 2015] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *CVPR*, 2015.
- [Wang *et al.*, 2020] Tan Wang, Jianqiang Huang, Hanwang Zhang, and Qianru Sun. Visual commonsense R-CNN. In *CVPR*, 2020.
- [Xiao *et al.*, 2020] Liqiang Xiao, Lu Wang, Hao He, and Yaohui Jin. Copy or rewrite: Hybrid summarization with hierarchical reinforcement learning. In *AAAI*, 2020.
- [Xu *et al.*, 2015] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C. Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.
- [Yang *et al.*, 2019] Xu Yang, Kaihua Tang, Hanwang Zhang, and Jianfei Cai. Auto-encoding scene graphs for image captioning. In *CVPR*, 2019.
- [Yao *et al.*, 2018] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Exploring visual relationship for image captioning. In *ECCV*, 2018.