

# COTS: Collaborative Two-Stream Vision-Language Pre-Training Model for Cross-Modal Retrieval

Haoyu Lu<sup>1,2</sup> Nanyi Fei<sup>1</sup> Yuqi Huo<sup>1</sup> Yizhao Gao<sup>1</sup> Zhiwu Lu<sup>1,2,\*</sup> Ji-Rong Wen<sup>1,2</sup>

<sup>1</sup>Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China

<sup>2</sup>Beijing Key Laboratory of Big Data Management and Analysis Methods

{lhy1998, feinanyi, bnhony, gaoyizhao, luzhiwu, jrwen}@ruc.edu.cn

## Abstract

*Large-scale single-stream pre-training has shown dramatic performance in image-text retrieval. Regrettably, it faces low inference efficiency due to heavy attention layers. Recently, two-stream methods like CLIP and ALIGN with high inference efficiency have also shown promising performance, however, they only consider instance-level alignment between the two streams (thus there is still room for improvement). To overcome these limitations, we propose a novel **COL**laborative **TWO**-Stream vision-language pre-training model termed **COTS** for image-text retrieval by enhancing cross-modal interaction. In addition to instance-level alignment via momentum contrastive learning, we leverage two extra levels of cross-modal interactions in our COTS: (1) *Token-level interaction* – a masked vision-language modeling (MVLM) learning objective is devised without using a cross-stream network module, where variational autoencoder is imposed on the visual encoder to generate visual tokens for each image. (2) *Task-level interaction* – a KL-alignment learning objective is devised between text-to-image and image-to-text retrieval tasks, where the probability distribution per task is computed with the negative queues in momentum contrastive learning. Under a fair comparison setting, our COTS achieves the highest performance among all two-stream methods and comparable performance (but with  $10,800\times$  faster in inference) w.r.t. the latest single-stream methods. Importantly, our COTS is also applicable to text-to-video retrieval, yielding new state-of-the-art on the widely-used MSR-VTT dataset.*

## 1. Introduction

The pretrain-then-finetune paradigm has achieved great success in the field of natural language processing (NLP), where models are first pre-trained with large-scale data (e.g., BERT [10], RoBERTa [30], and GPT3 [5]) and then finetuned for each downstream task. Recently, this prac-

tice has also shown its effectiveness in the vision-language (VL) domain [9, 17–19, 28, 37, 52], where the performance on various VL tasks (e.g., image-text retrieval, video-text retrieval, and visual question answering) has been significantly improved by vision-language pre-training (VLP). VLP models typically take huge image-text pairs as input and aim to learn joint image-text representations with single- and cross-modal pre-training objectives, such as masked token prediction and image-text matching.

Existing VLP models can be divided into two groups: single-stream models and two-stream ones. Single-stream VLP models (see Figure 1(a)) often utilize cross-modal fusion modules (e.g., Transformer [43] layers) to model fine-grained interactions between image regions and text words. Although these models achieve promising performance, they have two limitations: (1) During inference, all possible query-candidate pairs need to be fed into the fusion modules to calculate similarity scores, resulting in huge computational cost. (2) To obtain meaningful image regions, single-stream models typically adopt object detectors, which are expensive in both computation and data annotation. For example, extracting object regions from a  $800\times 1,333$  image takes about 900ms for Faster R-CNN [39], while ViT-base [11] only needs 15ms (i.e.,  $60\times$  faster). In contrast, two-stream VLP models [22, 44] apply separate image and text encoders and match image-text pairs on the final embedding level. Although two-stream models (see Figure 1(b)–(c)) are much more efficient than single-stream ones, they only achieve sub-optimal results due to the lack of closer image-text interactions. Therefore, a few works [42, 46] (see Figure 1(b)) reconsider object detectors, and most recent ones (e.g., CLIP [37], ALIGN [18], and WenLan [17]) resort to extra large pre-training data crawled from the Internet. However, they still fail to model fine-grained interactions between the two modalities.

To address the inefficiency of single-stream VLP models and the lack of closer vision-language interactions of two-stream ones, we propose a novel **COL**laborative **TWO**-Stream vision-language pre-training model termed COTS

\*The corresponding author.

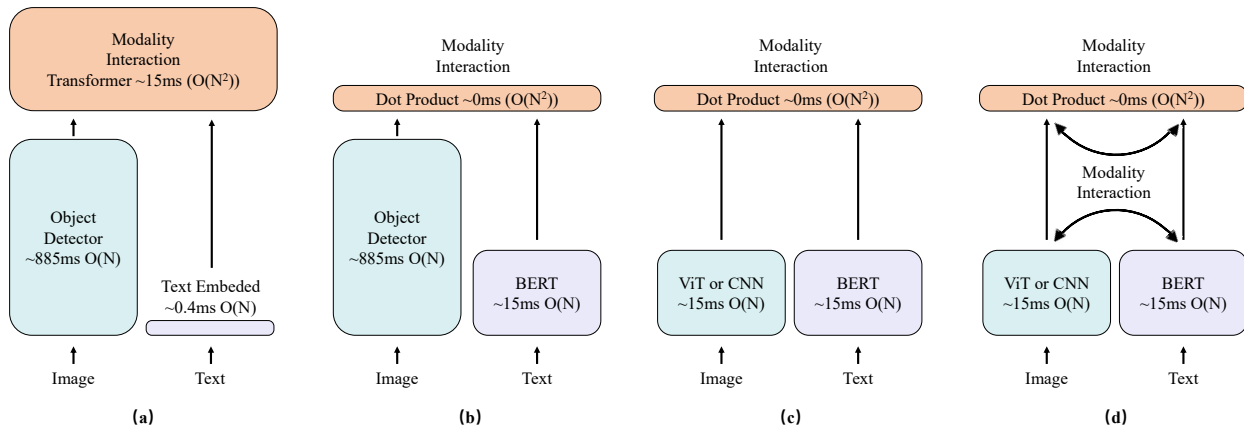


Figure 1. Four categories of vision-language pre-training (VLP) models. (a) Single-stream models (e.g., Oscar [28] and VinVL [52]). (b) Two-stream models with the object detector (e.g., LightingDot [42]). (c) Two-stream models with instance-level interaction (e.g., CLIP [37] and ALIGN [18]). (d) COTS: our two-stream model with multi-level interactions. The inference time and time complexity of each module are also reported, and more details can be found in Section 4.2.

for cross-modal retrieval, which retains the advantage of real-time inference speed and also enhances the interactions between the two modalities (see Figure 1(d)). Concretely, we consider three levels of cross-modal interactions in our COTS: (1) Instance-level interaction – an image-text matching learning objective at the final embedding level (typically adopted by two-stream VLP models) is devised via momentum contrastive learning [15], where we maintain two sample queues (one per modality) to have large size of negative samples. (2) Token-level interaction – a novel masked vision-language modeling (MVLN) learning objective is considered without using any cross-stream network module. To this end, we first tokenize both the image and the text for each input image-text pair, where variational autoencoder [21] is imposed on the visual encoder (e.g., ViT [11]) to generate visual tokens and BERT [10] is adopted for the text encoder. We then perform masked visual token prediction based on the unmasked visual tokens and the feature of each image’s paired text, and perform masked language token prediction similarly. (3) Task-level interaction – a novel KL-alignment learning objective is devised between text-to-image and image-to-text retrieval tasks by minimizing the Kullback-Leibler (KL) Divergence between probability distributions of the two retrieval tasks. For each image-text pair, the probability distribution of the text-to-image retrieval task is obtained with the similarities of the chosen text and its unpaired images in the negative image queue maintained in momentum contrastive learning, and we can obtain the other distribution similarly.

As the scale of pre-training data becomes large (e.g., tens of millions or even billions of image-text pairs crawled from the Internet), it is impossible to perform human-annotation and thus there inevitably exist noises in the large-scale data. Noisy data such as mis-matched image-text pairs and totally meaningless ones could bring negative effect for pre-

training. In this paper, we thus propose an adaptive momentum filter (AMF) module for our COTS, which can make full use of the momentum mechanism in our contrastive learning-based training algorithm. Specifically, we first calculate the similarity scores of all image-text pairs from the dynamically maintained image and text queues to obtain an extra queue. Further, we model this queue of similarity scores as a normal distribution and filter out the noisy data with the distribution mean and variance on the fly.

Our contributions are summarized as follows: (1) We propose a novel Collaborative Two-Stream (COTS) VLP model to improve the performance of two-stream models and retain their efficiency advantage at the same time. We achieve this by leveraging two extra levels of cross-modal interactions in addition to the typical instance-level alignment: a masked vision-language modeling (MVLN) learning objective for token-level interaction, and a KL-alignment learning objective for task-level interaction. (2) To alleviate the negative effect caused by the noises in large-scale pre-training data, we propose an adaptive momentum filter (AMF) module. AMF makes full use of the momentum mechanism in our instance-level alignment and adaptively filters noisy image-text pairs during pre-training. (3) Under a fair comparison setting, our COTS achieves the highest performance among all two-stream methods and performs comparably (but  $10,800\times$  faster in inference) with the latest single-stream ones. Importantly, our COTS is also applicable to text-to-video retrieval, yielding new state-of-the-art on the widely-used MSR-VTT dataset.

## 2. Related Work

**Vision-Language Pre-Training.** Recently, VLP resorts to single-stream models or two-stream ones. Single-stream models [9, 14, 16, 28, 31, 52] contain cross-modal fusion modules (e.g., Transformer [43] layers) to model closer

interactions between image regions and text words. Although single-stream models often achieve superior performance, they have several limitations in real-world scenarios: (1) When performing cross-modal retrieval during inference, all possible query-candidate pairs need to be fed into the fusion modules to calculate similarity scores, resulting in huge computational cost. (2) To obtain meaningful image regions, single-stream models often adopt object detectors, which are expensive in both computation and data annotation. In contrast, two-stream models project the two modalities into a joint embedding space and align them on the final embedding level. Early two-stream models [44, 48] only achieve sub-optimal performance because they do not consider fine-grained cross-modal interactions. More recent works (e.g., CLIP [37], ALIGN [18], and WenLan [17]) choose to improve their performance by leveraging extra large web data. However, they fail to model fine-grained interactions between the two modalities. Although the latest two-stream model LightingDot [42] considers token-level interaction, it still relies on an object detector, thus suffering from heavy computation. In this work, our COTS integrates the advantages of single-stream and two-stream models by still utilizing the two-stream architecture but enhancing the modeling of cross-modal interactions.

**Masked Vision Modeling.** Many previous works on VLP [9, 28] adopt masked vision modeling based on object tags to achieve better performance. They typically deploy a bottom-up attention mechanism [2] implemented by first extracting the object tags with Faster R-CNN [39] and then predicting the masked tags with other unmasked tags and text tokens. Although higher performance can be achieved, they commonly face two issues: (1) A heavy detector is needed to extract object tags, which is computationally expensive. For example, a Faster R-CNN detector takes 900ms to extract fine-grained region information from an image, which is nearly 60× slower than our ViT-base backbone (15ms). (2) These VLP models are not end-to-end trained, which may fail to cope with unknown objects. The latest work [19] shows that simply predicting masked raw image pixels is hard to improve the performance. Different from these works, our COTS employs a variational autoencoder [21] as an image tokenizer to tokenize a raw image into discrete image tokens for masked vision modeling, inspired by the vision Transformer BEIT [4]. The tokenizer is pre-trained in an end-to-end unsupervised training style, avoiding inducing handcrafted tags or heavy object detectors. Importantly, compared with predicting raw pixels directly, our choice of predicting masked image tokens is more meaningful as each image token contains specific high-level visual information. Overall, by combining masked vision modeling with masked language modeling, we devise a novel masked vision-language modeling (MVLM) objective for closer token-level interaction.

## 3. Methodology

### 3.1. Framework Overview

The goal of our COTS model for VLP is to learn two separate encoders that can embed image and text samples into the same semantic space for effective cross-modal retrieval. As illustrated in Figure 2, images and texts are encoded by the vision Transformer and the language Transformer, respectively. We then devise three levels of cross-modal interactions as the pre-training objectives of our COTS. Concretely, the instance-level interaction aligns the global features of paired images and texts by momentum cross-modal contrastive learning, which is inspired by the single-modal MoCo [15]. To model closer interactions than instance-level alignment, we propose to devise a masked vision-language modeling (MVLM) loss to enhance token-level interaction. MVLM has two parts: cross-modal masked vision modeling (CMVM) and cross-modal masked language modeling (CMLM). For each image, CMVM aims to predict the label of the masked image patch token based on unmasked ones together with the global feature of its paired text. CMLM does similarly on the language side. Further, we consider task-level interaction in our COTS, which aims to align the probability distributions of text-to-image and image-to-text retrieval tasks. In addition, to cope with the noises in the large-scale pre-training data, we propose an adaptive momentum filter (AMF) module, which is seamlessly integrated into the pre-training process.

Our choice of adopting the two-stream architecture in COTS has two main advantages: (1) Real-time inference speed – the separate image and text encoders allow us to compute the features of candidates beforehand for cross-modal retrieval tasks, and only a simple dot product needs to be calculated for each query-candidate pair. (2) Applicability to text-to-video retrieval – without any modification, our COTS can be directly applied to the text-to-video retrieval task, where the video representation can be obtained by averaging frame embeddings obtained by the image encoder. More details are given in Section 4.3.

### 3.2. Training Objectives

#### 3.2.1 Token-Level Interaction

We devise a masked vision-language modeling (MVLM) loss to enhance the token-level interaction in our COTS, which can be further split into two parts: cross-modal masked vision modeling (CMVM) and cross-modal masked language modeling (CMLM). To improve the practice [9, 28] of predicting masked image region tags with heavy object detectors, we introduce CMVM based on an image tokenizer inspired by BEIT [4]. For each image, the objective of CMVM is to predict the labels of masked image tokens with the unmasked image patches and paired text.




Figure 2. A schematic illustration of the proposed COTS for cross-modal retrieval.

Formally, let  $\mathcal{D} = \{(v_i, l_i)\}_{i=1}^N$  denote the training dataset, where  $(v_i, l_i)$  is the  $i$ -th image-text pair. For each raw image  $v_i$ , we first utilize the pre-trained discrete variational auto-encoder (dVAE) [38] as the image tokenizer to obtain a sequence of  $24 \times 24$  discrete image tokens  $\mathcal{T}_i^v = \{t_{i,j}^v \in \mathcal{V}^v\}_{j=1}^{576}$ , where  $t_{i,j}^v$  is the  $j$ -th token of image  $v_i$  and  $\mathcal{V}^v$  is the vocabulary of discrete image tokens. Meanwhile, the raw image is split into  $24 \times 24$  patches, which are fed into a vision Transformer [11] to obtain their embeddings. We then predict the label of each masked token based on the summation of the masked token embedding (which is already fused with unmasked token embeddings) and the global embedding of the paired text. The CMVM loss can thus be formulated as:

$$\mathcal{L}_{\text{CMVM}} = -\mathbb{E}_{(v_i, l_i) \sim \mathcal{D}} \log P(t_{i,j}^v | t_{i,\setminus j}^v, l_i), \quad (1)$$

where  $t_{i,j}^v$  denotes the target/masked image token, and  $t_{i,\setminus j}^v = \mathcal{T}_i^v \setminus \{t_{i,j}^v\}$  denotes the unmasked image tokens.

Similar to CMVM, for each piece of text  $l_i$ , the objective of CMLM is to predict the label of each masked word token based on unmasked ones and the paired image:

$$\mathcal{L}_{\text{CMLM}} = -\mathbb{E}_{(v_i, l_i) \sim \mathcal{D}} \log P(t_{i,j}^l | t_{i,\setminus j}^l, v_i), \quad (2)$$

where  $t_{i,j}^l$  denotes the target/masked text word token, and  $t_{i,\setminus j}^l$  denotes the unmasked ones. The total loss of our token-level cross-modal interaction is then defined as:

$$\mathcal{L}_{\text{token}} = \mathcal{L}_{\text{CMVM}} + \mathcal{L}_{\text{CMLM}}. \quad (3)$$

### 3.2.2 Instance-Level Interaction

To model the instance-level interaction of two modalities (i.e., global feature alignment) in our COTS, we adopt a cross-modal momentum contrastive learning (MCL) algorithm inspired by the single-modal MoCo [15], which provides a mechanism of dynamically maintaining negative sample queues for contrastive learning. Since the two queues (one for each modality) used in our MCL successfully decouple the queue size from the mini-batch size, the size of negative samples (crucial for contrastive learning) can be much larger than the mini-batch size.

Concretely, let  $f^v$  (with parameters  $\theta^v$ ) and  $f^l$  (with parameters  $\theta^l$ ) denote the image and text encoders, respectively. We adopt two extra momentum encoders  $\hat{f}^v$  (with parameters  $\hat{\theta}^v$ ) and  $\hat{f}^l$  (with parameters  $\hat{\theta}^l$ ) for the vision and language modalities, respectively. The parameters of momentum encoders are updated by:

$$\hat{\theta}^v = m \cdot \hat{\theta}^v + (1 - m) \cdot \theta^v, \quad (4)$$

$$\hat{\theta}^l = m \cdot \hat{\theta}^l + (1 - m) \cdot \theta^l, \quad (5)$$

where  $m$  is the momentum hyper-parameter.

Further, we maintain two queues  $\mathcal{Q}^v = \{\hat{\mathbf{q}}_j^v\}_{j=1}^{N_q}$  and  $\mathcal{Q}^l = \{\hat{\mathbf{q}}_j^l\}_{j=1}^{N_q}$ , where  $\hat{\mathbf{q}}_j^v/\hat{\mathbf{q}}_j^l$  denotes the momentum feature vector, and  $N_q$  denotes the queue size. Samples in each mini-batch  $\mathcal{B} = \{(v_i, l_i)\}_{i=1}^{N_b} \subseteq \mathcal{D}$  ( $N_b = |\mathcal{B}| \ll N_q$ ) are fed into current momentum encoders to obtain their mo-

momentum feature vectors, which are then pushed into corresponding queues *after loss calculation*. Meanwhile, the earliest  $N_b$  momentum feature vectors in each queue are popped out. Given each image in a data batch, by regarding its paired text as the positive sample and all samples in  $\mathcal{Q}^l$  as negative ones, we define the image-to-text contrastive loss as ( $\tau$  is the temperature hyper-parameter):

$$\mathcal{L}_{\text{I2T}} = -\frac{1}{N_b} \sum_{(v_i, l_i) \in \mathcal{B}} \log \frac{\text{pos}(\mathbf{f}_i^v, \hat{\mathbf{f}}_i^l, \tau)}{\text{pos}(\mathbf{f}_i^v, \hat{\mathbf{f}}_i^l, \tau) + \text{neg}(\mathbf{f}_i^v, \mathcal{Q}^l, \tau)}, \quad (6)$$

where  $\mathbf{f}_i^v = f^v(v_i)$ ,  $\hat{\mathbf{f}}_i^l = \hat{f}^l(l_i)$ , and

$$\begin{aligned} \text{pos}(\mathbf{f}_i^v, \hat{\mathbf{f}}_i^l, \tau) &= \exp(\mathbf{f}_i^v \cdot \hat{\mathbf{f}}_i^l / \tau), \\ \text{neg}(\mathbf{f}_i^v, \mathcal{Q}^l, \tau) &= \sum_{\hat{\mathbf{q}}_j^l \in \mathcal{Q}^l} \exp(\mathbf{f}_i^v \cdot \hat{\mathbf{q}}_j^l / \tau). \end{aligned} \quad (7)$$

The similarity of two feature vectors is measured by dot product here. Similarly, given each text in a data batch, we define the text-to-image contrastive loss as:

$$\mathcal{L}_{\text{T2I}} = -\frac{1}{N_b} \sum_{(v_i, l_i) \in \mathcal{B}} \log \frac{\text{pos}(\mathbf{f}_i^l, \hat{\mathbf{f}}_i^v, \tau)}{\text{pos}(\mathbf{f}_i^l, \hat{\mathbf{f}}_i^v, \tau) + \text{neg}(\mathbf{f}_i^l, \mathcal{Q}^v, \tau)}, \quad (9)$$

where  $\mathbf{f}_i^l = f^l(l_i)$ , and  $\hat{\mathbf{f}}_i^v = \hat{f}^v(v_i)$ . The total loss of our instance-level cross-modal interaction is then defined as:

$$\mathcal{L}_{\text{inst}} = \mathcal{L}_{\text{I2T}} + \mathcal{L}_{\text{T2I}}. \quad (10)$$

### 3.2.3 Task-Level Interaction

As we can see from Eq. (6) that, for each image  $v_i$  in a mini-batch, the image-to-text contrastive objective is actually maximizing the probability of matching its paired text  $l_i$  against the unmatched samples in  $\mathcal{Q}^l$  (so does the text side). That is, the instance-level feature alignment only cares about maximizing one particular probability in the whole probability distribution of the image-to-text/text-to-image retrieval task, and fails to capture a higher level interaction between two modalities. To fill the void in the literature, we propose to align the probability distributions of two cross-modal retrieval tasks as our task-level interaction.

Concretely, for each image-text pair  $(v_i, l_i) \in \mathcal{B}$ , we define the probability distribution of the image-to-text task as:

$$\mathcal{D}_{\text{I2T}} = [p(\mathbf{f}_i^v, \hat{\mathbf{f}}_i^l), p(\mathbf{f}_i^v, \hat{\mathbf{q}}_1^l), \dots, p(\mathbf{f}_i^v, \hat{\mathbf{q}}_{N_q}^l)], \quad (11)$$

where

$$p(\mathbf{f}_i^v, \hat{\mathbf{f}}_i^l) = \frac{\exp(\mathbf{f}_i^v \cdot \hat{\mathbf{f}}_i^l / \tau)}{\sum_{\hat{\mathbf{f}} \in \{\hat{\mathbf{f}}_i^l\} \cup \mathcal{Q}^l} \exp(\mathbf{f}_i^v \cdot \hat{\mathbf{f}} / \tau)}, \quad (12)$$

and  $p(\mathbf{f}_i^v, \hat{\mathbf{q}}_j^l)$  ( $\hat{\mathbf{q}}_j^l \in \mathcal{Q}^l$ ,  $j = 1, 2, \dots, N_q$ ) can be calculated in the same way. Similarly, we obtain the probability distribution of the text-to-image task as:

$$\mathcal{D}_{\text{T2I}} = [p(\mathbf{f}_i^l, \hat{\mathbf{f}}_i^v), p(\mathbf{f}_i^l, \hat{\mathbf{q}}_1^v), \dots, p(\mathbf{f}_i^l, \hat{\mathbf{q}}_{N_q}^v)]. \quad (13)$$

The learning objective of our task-level cross-modal interaction is then formulated as minimizing the symmetric Kullback-Leibler (KL) Divergence between  $\mathcal{D}_{\text{I2T}}$  and  $\mathcal{D}_{\text{T2I}}$ :

$$\mathcal{L}_{\text{task}} = \frac{1}{N_b} \sum_{(v_i, l_i) \in \mathcal{B}} (\text{KL}(\mathcal{D}_{\text{I2T}} \parallel \mathcal{D}_{\text{T2I}}) + \text{KL}(\mathcal{D}_{\text{T2I}} \parallel \mathcal{D}_{\text{I2T}})). \quad (14)$$

### 3.3. Adaptive Momentum Filter

Large-scale web-crawled data inevitably contain noises, which could bring negative effect for pre-training. Therefore, based on the momentum mechanism adopted in our COTS, we propose an adaptive momentum filter (AMF) module to adaptively filter noisy image-text pairs.

As introduced in the instance-level interaction, our COTS dynamically maintains two sample queues  $\mathcal{Q}^v$  and  $\mathcal{Q}^l$  for momentum contrastive learning. Since paired images and texts are pushed into or popped out of the corresponding queue simultaneously,  $\hat{\mathbf{q}}_j^v \in \mathcal{Q}^v$  and  $\hat{\mathbf{q}}_j^l \in \mathcal{Q}^l$  ( $j = 1, 2, \dots, N_q$ ) are also paired. We can then calculate a similarity score for each pair  $(\hat{\mathbf{q}}_j^v, \hat{\mathbf{q}}_j^l)$  by dot product. In this way, we obtain an extra similarity queue  $\mathcal{Q}^s = \{\hat{\mathbf{q}}_j^v \cdot \hat{\mathbf{q}}_j^l \mid \hat{\mathbf{q}}_j^v \in \mathcal{Q}^v, \hat{\mathbf{q}}_j^l \in \mathcal{Q}^l\}_{j=1}^{N_q}$ , which is also dynamically maintained along with the two sample queues.

Note that the similarity queue  $\mathcal{Q}^s$  can be seen as a sampling of the similarity score distribution at the current training iteration. We first calculate its mean  $\mu$  and standard deviation  $\sigma$  as the estimations of those of the similarity score distribution. We then obtain the threshold value  $s_{\text{AMF}}$  based on  $\mu$  and  $\sigma$  (e.g.,  $s_{\text{AMF}} = \mu - 2\sigma$ ) for our AMF. Finally, we use this threshold to filter the current data batch  $\mathcal{B}$  before we compute the losses:

$$\mathcal{B}^* = \{(v_i, l_i) \mid \hat{\mathbf{f}}_i^v \cdot \hat{\mathbf{f}}_i^l > s_{\text{AMF}}, (v_i, l_i) \in \mathcal{B}\}. \quad (15)$$

In this work,  $s_{\text{AMF}}$  changes in different training iterations as the similarity queue is changing. Specifically, when AMF is adopted in our full COTS, we use  $\mathcal{B}^*$  instead of  $\mathcal{B}$  in each iteration for loss computation, but we still push all samples in  $\mathcal{B}$  into  $\mathcal{Q}^v$  and  $\mathcal{Q}^l$  after loss computation.

## 4. Experiments

### 4.1. Datasets and Settings

**Pre-Training Datasets.** We use two image-text datasets for pre-training our COTS: (1) **CC4M** contains 4 million images and 5.3 million captions from Conceptual Captions (CC3M) [41], SBU [34], VG [24], MSCOCO [29] and Flickr30K [36]. (2) **CC14M** consists of CC4M and CC12M [6] (about 2 million urls are now invalid), which contains 14 million images and 15.3 million captions in total. Note that CC14M is much noisier than CC4M.

**Downstream Datasets.** We make downstream evaluation of our COTS on three widely-used benchmark datasets:

Model	# PT Pairs	Flickr30K (1K)						MSCOCO (5K)					
		I2T Retrieval			T2I Retrieval			I2T Retrieval			T2I Retrieval		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
<b>Single-Stream:</b>													
ViBERT-Base [31]	3.1M	-	-	-	58.2	84.9	91.5	-	-	-	-	-	-
Pixel-BERT-R50 [16]	5.6M	75.7	94.7	97.1	53.4	80.4	88.5	59.8	85.5	91.6	41.1	69.7	80.5
Pixel-BERT-X152 [16]	5.6M	<b>87.0</b>	<b>98.9</b>	<b>99.5</b>	71.5	92.1	95.8	63.6	87.5	93.6	50.1	77.6	86.2
Unicoder-VL [26]	3.8M	86.2	96.3	99.0	71.5	91.2	95.2	62.3	87.1	92.8	48.4	76.7	85.9
UNITER-Base [9]	9.6M	85.9	97.1	98.8	72.5	92.4	<b>96.1</b>	64.4	87.4	93.1	50.3	78.5	87.2
ERNIE-ViL-base [50]	3.8M	86.7	97.8	99.0	74.4	92.7	95.9	-	-	-	-	-	-
VILLA-Base [14]	9.6M	86.6	97.9	99.2	<b>74.7</b>	<b>92.9</b>	95.8	-	-	-	-	-	-
OSCAR-Base [28]	6.5M	-	-	-	-	-	-	70.0	91.1	95.5	54.0	80.8	88.5
ViLT [19]	9.9M	83.5	96.7	98.6	64.4	88.7	93.8	61.5	86.3	92.7	42.7	72.9	83.1
VinVL-Base [52]	8.9M	-	-	-	-	-	-	<b>74.6</b>	<b>92.6</b>	<b>96.3</b>	<b>58.1</b>	<b>83.2</b>	<b>90.1</b>
<b>Two-Stream:</b>													
VSE $_{\infty}^{*\dagger}$ [8]	-	88.7	98.9	99.8	76.1	94.5	97.1	68.1	90.2	95.2	52.7	80.2	88.3
COOKIE $^{*\dagger}$ [45]	5.9M	89.0	98.9	99.7	75.6	94.6	97.2	71.6	90.9	95.4	54.5	81.0	88.2
Frozen in time [3]	5.5M	-	-	-	61.0	87.5	92.7	-	-	-	-	-	-
LightningDOT [42]	9.5M	83.9	97.2	98.6	69.9	91.1	95.2	60.1	85.1	91.8	45.8	74.6	83.8
COOKIE [45]	5.9M	84.7	96.9	98.3	68.3	91.1	95.2	61.7	86.7	92.3	46.6	75.2	84.1
COTS (ours)	5.3M	88.2	98.5	99.7	75.2	93.6	96.5	66.9	88.8	94.0	50.5	77.6	86.1
COTS (ours)	15.3M	90.6	98.7	99.7	76.5	93.9	96.6	69.0	90.4	94.9	52.4	79.0	86.9
COTS $^{\dagger}$ (ours)	15.3M	<b>91.7</b>	<b>99.0</b>	<b>99.9</b>	<b>78.3</b>	<b>94.9</b>	<b>97.2</b>	<b>70.6</b>	<b>91.0</b>	<b>95.3</b>	<b>53.7</b>	<b>80.2</b>	<b>87.8</b>

Table 1. Comparative results for image-text retrieval on the Flickr30K (1K) test set and MSCOCO (5K) test set. Notations: # PT Pairs – the number of image-text pairs for pre-training; I2T Retrieval – image-to-text retrieval; T2I Retrieval – text-to-image retrieval.  $\dagger$  Ensemble results of two models. \* Models that utilize 940M tagged images for visual encoder pre-training.

(1) **MSCOCO** [29] is a large image-text dataset of 123,287 images, where each image is annotated with 5 captions. As in [19], we adopt the Karpathy split of MSCOCO: 5,000 images for testing, another 5,000 for validation, and the rest 113,287 images for training. (2) **Flickr30K** [36] contains 31,000 images and 158,915 captions totally. Each image is often annotated with 5 captions. Following the split in [12], we use 1,000 images for testing, another 1,000 for validation, and the rest for training. (3) To show the general applicability of our COTS, we also conduct experiments on a video-text dataset **MSR-VTT** [47], which has 10K YouTube videos and 200K captions. As in [49], we report our results under both the 1KA and 7K splits.

**Text and Image Encoders.** In our COTS, we follow [42] and adopt a BERT-base [10] model as our text encoder, which contains a total of 12 Transformer layers with 768 hidden units and 12 heads. Further, for computation efficiency, we use ViT-B/16 [11] as our image encoder with the input image resolution of  $384 \times 384$ . Overall, *only base text and image encoders* are considered in our COTS.

**Evaluation Metrics.** The widely-used  $R@k$  ( $k = 1, 5, 10$ ) in cross-modal retrieval is reported for performance evaluation, which is the proportion of matched samples found in the top- $k$  retrieved results. Following [3], we also report the Median Rank (MR) for video-text retrieval.

**Implementation Details.** For our masked vision-language modeling (MVLM), we randomly mask 40% image patches

following [4] and mask word tokens in text with 15% probability. We adopt the Adam [20] optimizer with a weight decay of 0.02. We select hyper-parameters heuristically due to computational constraint: the momentum hyper-parameter  $m = 0.99$ , temperature  $\tau = 0.05$ , and the queue size  $N_Q$  is 12,800, 6,400, and 1,200 for pre-training, finetuning on MSCOCO, and finetuning on Flickr30K, respectively. We set the initial learning rate to  $5e-5$  for the first 5 epochs, and decay the learning rate linearly in the rest epochs. More implementation details can be found in the supp. material.

## 4.2. Image-Text Retrieval

**Comparison to the State-of-the-Arts.** We compare our COTS with the state-of-the-art methods on two widely-used image-text datasets: Flickr30K and MSCOCO. As shown in Table 1, under a fair comparison setting (excluding VSE $_{\infty}^{*\dagger}$  [8] and COOKIE $^{*\dagger}$  [45] which utilize 940M tagged images for visual encoder pre-training), our COTS outperforms all two-stream models by large margins for all evaluation metrics. Specifically, compared with the most recent two-stream model COOKIE [45], our COTS achieves higher results by 5.2% (66.9% vs. 61.7%) for I2T R@1 on MSCOCO and 3.9% (50.5% vs. 46.6%) for T2I R@1 on MSCOCO, but with less pre-training data (5.3M vs. 5.9M). Moreover, when leveraging larger pre-training dataset and model ensemble technique, our COTS $^{\dagger}$  further improves the performance. Specifically, without using extra object detectors, our COTS $^{\dagger}$  achieves new state-of-the-art on Flickr30K




Figure 3. Comparison of inference time to different methods (VinVL [52], COOKIE [45], VSE [8]) on MSCOCO (5K) test set.

Model	I2T Retrieval			T2I Retrieval		
	R@1	R@5	R@10	R@1	R@5	R@10
ViLT [19]	56.5	82.6	89.6	40.4	70.0	81.1
CLIP [37]	58.4	81.5	88.1	37.8	62.4	72.2
ALIGN [18]	58.6	83.0	87.9	<b>45.6</b>	69.8	78.6
COTS (w/o FT)	<b>60.4</b>	<b>84.7</b>	<b>91.7</b>	43.8	<b>71.6</b>	<b>81.3</b>

Table 2. Comparative results (without finetuning) for image-text retrieval on the MSCOCO (5K) test set. FT – Finetuning.

w.r.t. both single-stream and two-stream methods. On MSCOCO, our COTS<sup>†</sup> also achieves higher performance than most single-stream methods and comparable results compared with ViLT [19] but with a 10,800× faster speed during inference (see Inference Efficiency Analysis).

**Inference Efficiency Analysis.** In real-world application scenarios, inference speed is an important evaluation metric for retrieval methods. In Figure 3, we compare our COTS with recent state-of-the-arts regarding the inference time on the MSCOCO (5K) test set. All methods are evaluated on a single Tesla V100 GPU. Compared with the single-stream ViLT [19], our COTS is 10,800× faster on the whole MSCOCO (5K) test set. This huge gap will even become dramatically larger when the size of test set  $N$  grows, as the retrieval time complexity for single-stream models is  $O(N^2)$  while it is nearly  $O(N)$  for two-stream ones. Although VSE and COOKIE are also two-stream models, our COTS is still significantly faster than them, indicating the extreme high efficiency of our COTS due to its fully tokenized Transformer-based architecture.

**Comparative Retrieval Results without Finetuning.** Following ViLT [19], we report the comparative retrieval results without finetuning on MSCOCO in Table 2. We can observe that: (1) Our COTS outperforms the latest single-stream method ViLT [19]. (2) Our COTS also beats the latest two-stream methods CLIP [37] and ALIGN [18], although it is pre-trained with much less data.

**Ablation Study Results.** In Table 3, we analyze the contributions of different pre-training objectives and the adaptive momentum filter (AMF) module in our COTS. We randomly sample 200K image-text pairs from CC12M as the pre-training dataset (termed CC200K). Zero-shot retrieval

Method	I2T Retrieval			T2I Retrieval		
	R@1	R@5	R@10	R@1	R@5	R@10
$\mathcal{L}_{inst}$	24.0	48.3	60.0	16.8	37.5	49.6
$\mathcal{L}_{inst} + \mathcal{L}_{CMLM}$	24.5	49.3	61.1	16.5	37.8	49.9
$\mathcal{L}_{inst} + \mathcal{L}_{token}$	25.6	49.9	61.9	17.1	38.3	50.4
$\mathcal{L}_{inst} + \mathcal{L}_{token} + \mathcal{L}_{task}$	26.4	50.5	<b>62.9</b>	17.5	38.5	50.6
$\mathcal{L}_{inst}$ (w/ AMF)	24.7	49.6	61.3	16.6	38.3	50.0
Our Full COTS	<b>27.1</b>	<b>51.1</b>	<b>62.9</b>	<b>17.9</b>	<b>39.2</b>	<b>51.1</b>

Table 3. Ablation study for our COTS pre-trained on the small CC200K dataset. Zero-shot image-text retrieval results are reported on the MSCOCO (5K) test set.

Model	# PT Pairs	R@1	R@5	R@10	MR↓
<b>7K split:</b>					
JSFusion [51]	-	10.2	31.2	43.2	13.0
HT MIL-NCE [33]	>100M	14.9	40.2	52.8	9.0
ActBERT [53]	>100M	16.3	42.8	56.9	10.0
HERO [27]	>100M	16.8	43.4	57.7	-
VidTranslate [23]	>100M	14.7	-	52.8	-
NoiseEstimation* [1]	>100M	17.4	41.6	53.6	8.0
UniVL* [32]	>100M	21.2	49.6	63.1	6.0
ClipBERT [25]	5.6M	22.0	46.8	59.9	6.0
TACo* [49]	>100M	24.8	52.1	64.5	5.0
COTS (ours)	5.3M	29.0	57.0	67.7	3.0
COTS (ours)	15.3M	<b>32.1</b>	<b>60.8</b>	<b>70.2</b>	<b>3.0</b>
<b>1KA split:</b>					
AVLnet* [40]	>100M	27.1	55.6	66.6	4.0
MMT* [13]	>100M	26.6	57.1	69.6	4.0
TACo* [49]	>100M	28.4	57.8	71.2	4.0
Support Set* [35]	>100M	30.1	58.5	69.3	3.0
Frozen in Time [3]	5.5M	31.0	59.5	70.5	3.0
COTS (ours)	5.3M	33.1	61.3	72.8	3.0
COTS (ours)	15.3M	<b>36.8</b>	<b>63.8</b>	<b>73.2</b>	<b>2.0</b>

Table 4. Comparison to the state-of-the-arts for text-to-video retrieval on MSR-VTT under two splits: the 7K and 1KA splits. Notations: ↓ denotes that lower results are better; \* denotes that extra modalities (e.g., motion and audio) are used.

results are reported on the MSCOCO (5K) test set. We start with our instance-level interaction loss  $\mathcal{L}_{inst}$  (without AMF) and then add other losses successively. We can observe from Table 3 that: (1) Both CMLM and CMVM bring performance improvements (see  $\mathcal{L}_{inst} + \mathcal{L}_{CMLM}$  vs.  $\mathcal{L}_{inst}$ , and  $\mathcal{L}_{inst} + \mathcal{L}_{token}$  vs.  $\mathcal{L}_{inst} + \mathcal{L}_{CMLM}$ ), indicating that token-level cross-modal interactions are beneficial to learning the aligned multi-modal representation space. (2) When task-level interaction is added (see  $\mathcal{L}_{inst} + \mathcal{L}_{token} + \mathcal{L}_{task}$  vs.  $\mathcal{L}_{inst} + \mathcal{L}_{token}$ ), the performance is further improved, which clearly validates the effectiveness of our multi-level cross-modal interactions. (3) Our AMF module works well with either instance-level or multi-level interactions (see  $\mathcal{L}_{inst}$  (w/ AMF) vs.  $\mathcal{L}_{inst}$ , and Our Full COTS vs.  $\mathcal{L}_{inst} + \mathcal{L}_{token} + \mathcal{L}_{task}$ ). (4) Combining all objectives with the AMF module (i.e., Our Full COTS) leads to the best results, indicating that each objective/module is complementary to each other.

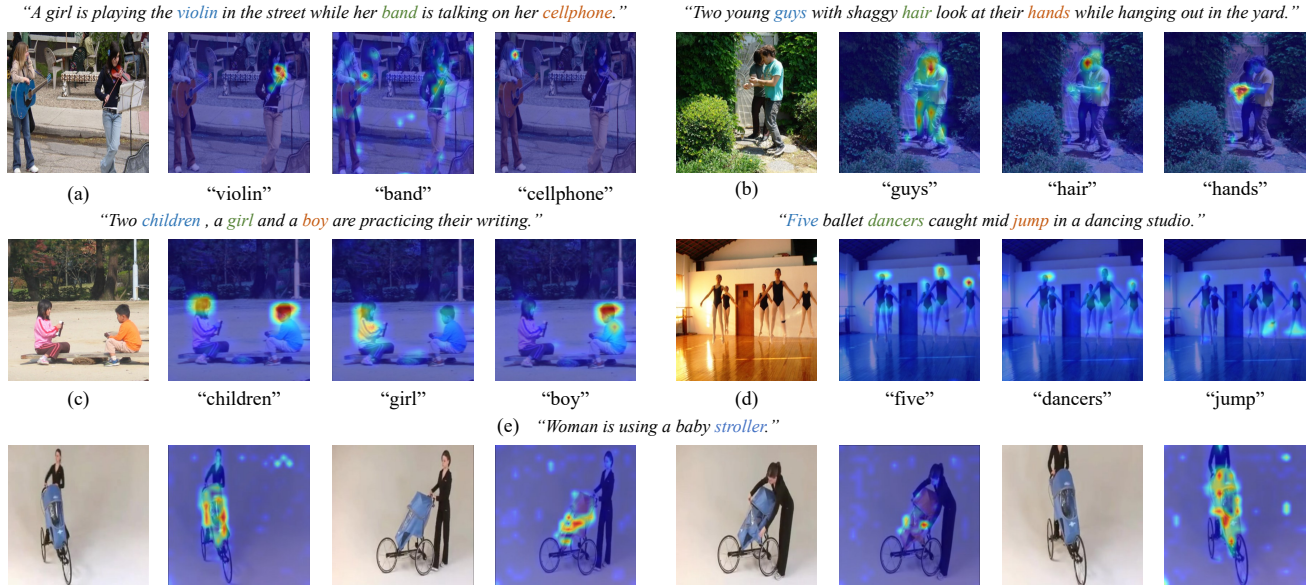


Figure 4. Visualizations of attention maps of our COTS using GAE [7] on images/video frames responding to individual words. (a) – (d) Image attention maps w.r.t. different words. (e) Video frame attention maps w.r.t. the word “stroller”.

### 4.3. Video-Text Retrieval

We further compare our COTS with the state-of-the-art methods on the video-text retrieval task. To directly deploy our COTS, we do not consider utilizing complex methods or additional modules to model the temporal information of videos. Instead, we simply use the mean frame embeddings as video representations and then calculate similarity scores by dot product with text embeddings. We report the text-to-video retrieval results on the MSR-VTT dataset in Table 4. Note that only text-to-video retrieval is considered as in the latest work [3]. It can be seen that: (1) Our COTS significantly outperforms the state-of-the-arts even without modeling the temporal information of videos, which demonstrates the general applicability and the great potentiality of our COTS. (2) Our COTS leads to better results than methods utilizing extra modalities (e.g., motion and audio) or those pre-trained on extra large video data (e.g., the HowTo100M dataset [33] with more than 100 million video-text pairs), indicating that a well pre-trained vision-language model may be the key to video-text retrieval.

### 4.4. Visualization Results

Figure 4 shows the visualized attention maps of our COTS on images/video frames responding to individual words. We can see from Figures 4(a)–(b) that our COTS can well locate different objects (even fine-grained ones like “violin” and “cellphone” in Figure 4(a), “hair” and “hands” in Figure 4(b)) in the same image. Figure 4(c) shows how our COTS determines gender information. Given the word “children”, COTS focuses on the faces. When recognizing “girl”, COTS pays attention to the girl’s long hair and pink

clothes (and the same for the word “boy”). Interestingly, our COTS can also capture abstract concepts (“five”) and actions (“jump”) as shown in Figure 4(d). COTS focuses on five dancers for both “five” and “dancers”, but pays more attention for the number “five”. And it focuses on feet when it comes to “jump”. Figure 4(e) presents attention maps w.r.t. “stroller” on four frames from the same video, showing that our COTS can also work well for the video modality.

## 5. Conclusion

In this paper, we have investigated how to improve the performance of the two-stream vision-language pre-training (VLP) while still maintaining its advantage of high efficiency for image-text retrieval. Specifically, we propose a novel COLlaborative Two-Stream VLP model termed COTS by leveraging three levels of cross-modal interactions in image-text retrieval. That is, we consider token-level interaction by masked vision-language modeling with both tokenized images and texts, instance-level interaction by cross-modal momentum contrastive learning, and task-level interaction by aligning two task distributions. Extensive experiments validate the effectiveness and high efficiency of our COTS in image-text retrieval. It is also shown to have general applicability as it achieves new state-of-the-art on video-text retrieval without any modification.

**Acknowledgements** This work was supported in part by National Natural Science Foundation of China (61976220 and 61832017), Beijing Outstanding Young Scientist Program (BJJWZYJH012019100020098), and Large-Scale Pre-Training Program 468 of BAAI.



## References

- [1] Elad Amrani, Rami Ben Ari, Daniel Rotman, and Alex Bronstein. Noise estimation using density estimation for self-supervised multimodal learning. In *AAAI*, pages 6644–6652, 2021. 7
- [2] 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*, pages 6077–6086, 2018. 3
- [3] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, pages 1728–1738, 2021. 6, 7, 8
- [4] Hangbo Bao, Li Dong, and Furu Wei. BEiT: BERT pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 3, 6
- [5] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *NeurIPS*, pages 1877–1901, 2020. 1
- [6] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12M: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *CVPR*, pages 3558–3568, 2021. 5
- [7] Hila Chefer, Shir Gur, and Lior Wolf. Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *ICCV*, pages 397–406, 2021. 8
- [8] Jiacheng Chen, Hexiang Hu, Hao Wu, Yuning Jiang, and Changhu Wang. Learning the best pooling strategy for visual semantic embedding. In *CVPR*, pages 15789–15798, 2021. 6, 7
- [9] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. UNITER: Universal image-text representation learning. In *ECCV*, pages 104–120, 2020. 1, 2, 3, 6
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186, 2018. 1, 2, 6
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 1, 2, 4, 6
- [12] Andrea Frome, Greg S Corrado, Jonathon Shlens, Samy Bengio, Jeffrey Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. DeViSE: a deep visual-semantic embedding model. In *NeurIPS*, pages 2121–2129, 2013. 6
- [13] Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *ECCV*, pages 214–229, 2020. 7
- [14] Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. Large-scale adversarial training for vision-and-language representation learning. In *NeurIPS*, pages 6616–6628, 2020. 2, 6
- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. 2, 3, 4
- [16] Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. Pixel-BERT: Aligning image pixels with text by deep multi-modal transformers. *arXiv preprint arXiv:2004.00849*, 2020. 2, 6
- [17] Yuqi Huo, Manli Zhang, Guangzhen Liu, Haoyu Lu, Yizhao Gao, Guoxing Yang, Jingyuan Wen, Heng Zhang, Baogui Xu, Weihao Zheng, et al. WenLan: Bridging vision and language by large-scale multi-modal pre-training. *arXiv preprint arXiv:2103.06561*, 2021. 1, 3
- [18] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V Le, Yunhsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, pages 4904–4916, 2021. 1, 2, 3, 7
- [19] Wonjae Kim, Bokyoung Son, and Ildoo Kim. ViLT: Vision-and-language transformer without convolution or region supervision. In *ICML*, pages 5583–5594, 2021. 1, 3, 6, 7
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [21] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In *ICLR*, 2013. 2, 3
- [22] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*, 2014. 1
- [23] Bruno Korbar, Fabio Petroni, Rohit Girdhar, and Lorenzo Torresani. Video understanding as machine translation. *arXiv preprint arXiv:2006.07203*, 2020. 7
- [24] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, 2017. 5
- [25] Jie Lei, Linjie Li, Luowei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. Less is more: ClipBERT for video-and-language learning via sparse sampling. *CVPR*, pages 7331–7341, 2021. 7
- [26] Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. Unicoder-VL: A universal encoder for vision and language by cross-modal pre-training. In *AAAI*, pages 11336–11344, 2020. 6
- [27] Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. HERO: Hierarchical encoder for video+ language omni-representation pre-training. *EMNLP*, pages 2046–2065, 2020. 7
- [28] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, pages 121–137, 2020. 1, 2, 3, 6

- [29] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, pages 740–755, 2014. 5, 6
- [30] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 1
- [31] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, pages 13–23, 2019. 2, 6
- [32] Huaishao Luo, Lei Ji, Botian Shi, Haoyang Huang, Nan Duan, Tianrui Li, Xilin Chen, and Ming Zhou. UniVL: A unified video and language pre-training model for multimodal understanding and generation. *arXiv preprint arXiv:2002.06353*, 2020. 7
- [33] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. HowTo100M: Learning a text-video embedding by watching hundred million narrated video clips. In *ICCV*, pages 2630–2640, 2019. 7, 8
- [34] Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2text: Describing images using 1 million captioned photographs. In *NeurIPS*, pages 1143–1151, 2011. 5
- [35] Mandela Patrick, Po-Yao Huang, Yuki Asano, Florian Metzger, Alexander Hauptmann, João Henriques, and Andrea Vedaldi. Support-set bottlenecks for video-text representation learning. In *ICLR*, 2021. 7
- [36] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, pages 2641–2649, 2015. 5, 6
- [37] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 1, 2, 3, 7
- [38] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092*, 2021. 4
- [39] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NeurIPS*, pages 91–99, 2015. 1, 3
- [40] Andrew Rouditchenko, Angie Boggust, David Harwath, Dhiraj Joshi, Samuel Thomas, Kartik Audhkhasi, Rogério Feris, Brian Kingsbury, Michael Picheny, Antonio Torralba, et al. AVLnet: Learning audio-visual language representations from instructional videos. *arXiv preprint arXiv:2006.09199*, 2020. 7
- [41] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, pages 2556–2565, 2018. 5
- [42] Siqi Sun, Yen-Chun Chen, Linjie Li, Shuohang Wang, Yuwei Fang, and Jingjing Liu. LightningDOT: Pre-training visual-semantic embeddings for real-time image-text retrieval. In *NAACL-HLT*, pages 982–997, 2021. 1, 2, 3, 6
- [43] 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*, pages 5998–6008, 2017. 1, 2
- [44] Liwei Wang, Yin Li, and Svetlana Lazebnik. Learning deep structure-preserving image-text embeddings. In *CVPR*, pages 5005–5013, 2016. 1, 3
- [45] Keyu Wen, Jin Xia, Yuanyuan Huang, Linyang Li, Jiayan Xu, and Jie Shao. COOKIE: Contrastive cross-modal knowledge sharing pre-training for vision-language representation. In *ICCV*, pages 2208–2217, 2021. 6, 7
- [46] Yiling Wu, Shuhui Wang, Guoli Song, and Qingming Huang. Learning fragment self-attention embeddings for image-text matching. In *ACM-MM*, pages 2088–2096, 2019. 1
- [47] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. MSR-VTT: A large video description dataset for bridging video and language. In *CVPR*, pages 5288–5296, 2016. 6
- [48] Fei Yan and Krystian Mikolajczyk. Deep correlation for matching images and text. In *CVPR*, pages 3441–3450, 2015. 3
- [49] Jianwei Yang, Yonatan Bisk, and Jianfeng Gao. TACO: Token-aware cascade contrastive learning for video-text alignment. In *ICCV*, pages 11562–11572, 2021. 6, 7
- [50] Fei Yu, Jiji Tang, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. Ernie-vil: Knowledge enhanced vision-language representations through scene graphs. In *AAAI*, pages 3208–3216. AAAI Press, 2021. 6
- [51] Youngjae Yu, Jongseok Kim, and Gunhee Kim. A joint sequence fusion model for video question answering and retrieval. In *ECCV*, pages 487–503, 2018. 7
- [52] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. VinVL: Revisiting visual representations in vision-language models. In *CVPR*, pages 5579–5588, 2021. 1, 2, 6, 7
- [53] Linchao Zhu and Yi Yang. ActBERT: Learning global-local video-text representations. In *CVPR*, pages 8743–8752, 2020. 7