

How to Make Cross Encoder a Good Teacher for Efficient Image-Text Retrieval?

Yuxin Chen^{1,2,3*}, Zongyang Ma^{1,2,3*}, Ziqi Zhang^{1*}, Zhongang Qi², Chunfeng Yuan^{1†},
Bing Li¹, Junfu Pu², Ying Shan², Xiaojuan Qi⁵, Weiming Hu^{1,3,4}

¹State Key Laboratory of Multimodal Artificial Intelligence Systems,

Institute of Automation, Chinese Academy of Sciences; ²ARC Lab, Tencent PCG;

³School of Artificial Intelligence, University of Chinese Academy of Sciences;

⁴School of Information Science and Technology, ShanghaiTech University; ⁵The University of Hong Kong

{chenyuxin2019, mazongyang2020}@ia.ac.cn, {ziqi.zhang, cfyuan, bli, wmhu}@nlpr.ia.ac.cn

{zhongangqi, jevinpu, yingsshan}@tencent.com, xjq@eee.hku.hk

Abstract

Dominant dual-encoder models enable efficient image-text retrieval but suffer from limited accuracy, while the cross-encoder models offer higher accuracy at the expense of efficiency. Distilling cross-modality matching knowledge from cross-encoder to dual-encoder provides a natural approach to harness their strengths. Thus, we investigate the following valuable question: how to make cross-encoder a good teacher for dual-encoder? Our findings are threefold: (1) Cross-modal similarity score distribution of cross-encoder is more concentrated, while the result of dual-encoder is nearly normal, making vanilla logit distillation less effective. However, ranking distillation remains practical, as it is not affected by the score distribution. (2) Only the relative order between hard negatives conveys valid knowledge, while the order information between easy negatives has little significance. (3) Maintaining the coordination between distillation loss and dual-encoder training loss is beneficial for knowledge transfer. Based on these findings, we propose a novel Contrastive Partial Ranking Distillation (CPRD) method, which implements the objective of mimicking relative order between hard negative samples with contrastive learning. This approach coordinates with the training of the dual-encoder, effectively transferring valid knowledge from the cross-encoder to the dual-encoder. Extensive experiments on image-text retrieval and ranking tasks show that our method surpasses other distillation methods and significantly improves the accuracy of dual-encoder.

1. Introduction

Dominant pre-training works for image-text retrieval adopt dual-encoder architecture [9, 21, 28, 33] to separately extract image and text representations and employ contrastive

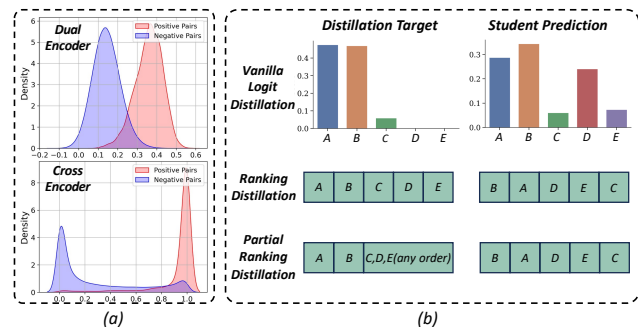


Figure 1. (a) Similarity score distribution of dual-encoder and cross-encoder. (b) Student predictions and targets for different types of distillation methods. For partial ranking distillation, the relative order between easy negatives is disregarded.

learning to facilitate global image-text alignment. Although dual-encoder has high efficiency, its retrieval accuracy is sub-optimal due to the lack of modality interactions except for the final dot product. In contrast, the cross-encoder architecture of another line of works perform deep cross-modal interaction between images and texts in a single encoder. This allows them to retrieve more accurately, but the need to calculate on all possible image-text pairs sacrifices retrieval efficiency.

To obtain an efficient and accurate retriever, a natural idea is to distill the cross-modal matching knowledge from the cross-encoder into the dual-encoder. However, the discrepancies between the two models make the existing attention and logit distillation methods less versatile or effective. On the one hand, attention distillation requires extra modules to compute cross-modal attention of dual-encoder and a cross-encoder with specific architecture to provide distillation targets. On the other hand, the differences in training objectives and architectures between two models lead to significantly distinct and non-directly alignable cross-modal similarity score distributions.

To explore a general and efficient way to achieve effective knowledge transfer from cross-encoder to dual-

* Equal contribution. † Corresponding author.

encoder, we analyze their distinct characteristics and the difficulties hindering distillation, and make the following observations: (1) As shown in Figure 1(a), the cross-modal similarity distribution of the dual-encoder constrained by contrastive learning tends to be "moderate" and normally distributed, whereas the similarity score distribution of cross-encoder optimized with image-text matching is more "aggressive" and mainly concentrates on 0/1. Vanilla logit distillation based on KL-divergence enforces the similarity score distribution of dual-encoder to approximate that of cross-encoder. It is difficult due to the substantial differences between them and may interfere with the original training objective of dual-encoder, thereby disrupting the learning of coarse image-text alignment. However, we find that mimicking the ranking of negatives derived from sorting similarity scores is an effective distillation scheme, which we refer to as ranking distillation, as it solely depends on the relative ranking of the scores and is not influenced by the sharpness of score distribution. (2) Furthermore, performance gap between dual-encoder and cross-encoder mainly stems from inaccurate ranking of top-ranked items. Supporting this claim, re-ranking top 32 items retrieved by dual-encoder with cross-encoder substantially improves retrieval accuracy. Further expanding the number of re-ranked items yields negligible gains, indicating that distillation should focus more on partial ranking of top-ranked samples rather than all samples. The illustrations of vanilla logit distillation, ranking distillation, and partial ranking distillation are shown in Figure 1(b) for clarity. (3) Different types of loss have varied effects on the image-text embedding space of dual-encoder. Therefore, it is beneficial to maintain coordination between ranking distillation loss and the original training loss of dual-encoder to avoid interfering with the learning process.

Based on the above observations, we propose a novel Contrastive Partial Ranking Distillation (CPRD) method to enable effective ranking knowledge distillation. Specifically, we learn the ranking of hard negative samples via contrastive learning. Given an image, we first utilize the dual-encoder to identify the top-K hard negative texts and obtain the ranking of these texts. Then, we feed the image and negative texts into cross-encoder to compute the matching scores and divide the negative texts into valid and invalid negative texts. Valid negative texts have a larger matching score with the image, and their relative order contains rich cross-modal matching knowledge. Therefore, we employ contrastive learning to pull higher-ranked valid negative texts closer to image while pushing the lower-ranked texts away, ensuring the consistency between ranking of valid negative texts from dual-encoder and that from the cross-encoder. On the other hand, invalid negative texts have a smaller matching score with the image, and their relative order does not contain valid information. Thus, we

only employ contrastive learning to push all invalid negative texts away and disregard the relative order between them.

This approach does not require the similarity distributions of dual-encoders and cross-encoders to be similar, overcoming the distillation difficulties caused by the substantial differences between similarity distributions. Furthermore, we achieve the objective of hard negative ranking learning with contrastive learning, aligning seamlessly with the dual-encoder training process. The contributions of this work can be summarized as follows:

- We conduct a comprehensive investigation into the effective knowledge distillation from cross-encoders to dual-encoders and identify three key aspects.
- We propose the Contrastive Partial Ranking Distillation (CPRD) method, which attains the objective of learning the relative order between valid hard negatives via contrastive learning, enabling effective knowledge transfer from cross-encoder to dual-encoder.
- Experimental results show that our method outperforms previous distillation methods, significantly improving the retrieval and ranking accuracy of dual-encoder.

2. Related Work

Pre-training for Image-text Retrieval. Pre-training works for image-text retrieval can be categorized into dual-encoder methods and cross-encoder methods. Dual-encoder methods [9, 21, 28, 31, 33] adopt two separate encoders to extract the image and text features separately, and employ contrastive learning to align global representations in a shared embedding space. Cross-encoder methods [1, 3, 14, 15, 17, 18, 22, 34, 35] employ a single encoder for the joint encoding of image-text features. Many proxy tasks are proposed to facilitate cross-modal interaction in the cross-encoder, *e.g.*, masked language modeling (MLM), masked region modeling (MRM), and masked image modeling (MIM), *etc.* In order to improve the retrieval accuracy of dual-encoders, one line of approaches [21, 31] draw inspiration from cross-encoder and enhance the cross-modal interaction by adapting MLM, MRM, and MIM tasks to dual-encoder. However, another approach, transferring knowledge from cross-encoder to dual-encoder through distillation, has not been fully explored.

Knowledge Distillation from Cross-Encoder to Dual-Encoder. Previous works distilling knowledge from the cross-encoder to the dual-encoder can be divided into attention distillation methods[32] and logit distillation methods[24]. Attention distillation methods aim to align the cross-modal attention of two models, which requires two prerequisites: (1) Both models have to adopt attention-based backbone, *e.g.*, ViT[4], BERT[10], to produce attention map. (2) Input to dual-encoder and cross-encoder is exactly identical to guarantee their attention maps have the same shape and semantics. These make attention distilla-

tion less versatile. Inspired by the distillation works in image classification, Miech et al. [24] and Lei et al. [13] introduce logit distillation into image-text retrieval. The core idea is to constrain the consistency of the image-text similarity score distribution of the dual encoder and cross-encoder through KL-divergence based loss. But substantial similarity distribution differences between dual-encoder and cross-encoder make it difficult to transfer knowledge effectively. To explore an effective and general distillation method to improve the performance of dual-encoders, we investigate and identify three key factors and further propose the Contrastive Partial Ranking Distillation method.

Ranking Distillation for Neural Ranking. In the field of textual neural ranking, there are some studies on ranking distillation [7, 23, 29]. Sashank et al.[29] propose rank distillation losses with cross-entropy or MSE loss to constrain the consistency of positive sample scores, which is less effective due to the significant differences in similarity score distribution in image-text retrieval. Sebastian et al.[7] propose Margin-MSE, which requires the same margin between positive and negative sample scores for student and teacher models. Aditya et al.[23] further propose M3SE, requiring the student and teacher models to have the same margin between positive and the hardest negative sample. However, it only considers the hardest negative samples, limiting the knowledge that can be transferred. Moreover, the MSE loss is not coordinated with the contrastive learning for dual-encoder training, resulting in interference in the learning process. In contrast, our proposed CPRD method considers the relative order among multiple hard negative samples via contrastive learning, aligning with the original training loss of dual-encoder.

3. Method

In this section, we first introduce the architecture and training objective of vanilla dual-encoder and cross-encoder, and explain the reasons for the formation of their similarity distribution characteristics in Sec 3.1. Then, we identify the source of the performance gap between them, underscore the importance of hard negatives ranking, and elaborate on our proposed Contrastive Partial Ranking Distillation (CPRD) method in Sec 3.2. Finally, we present the overall training objective of dual-encoder in Sec 3.3.

3.1. Dual-Encoder and Cross-Encoder

Dual-Encoder. The dual-encoder takes an image and its paired text as input, and extracts global visual and textual representations \mathbf{v}_i and \mathbf{t}_i with separate encoders. The dot product of \mathbf{v}_i and \mathbf{t}_i is used to measure the image-text similarity. During pre-training, following previous works[14], we maintain two queues Q^v and Q^t to preserve the momentum features from current mini-batch $\{\hat{\mathbf{v}}_j\}_{j=1}^B$ and $\{\hat{\mathbf{t}}_j\}_{j=1}^B$ and previous iterations $\{\hat{\mathbf{v}}_j\}_{j=B+1}^{B+Nq}$

and $\{\hat{\mathbf{t}}_j\}_{j=B+1}^{B+Nq}$. For every image/text in the current mini-batch, its related text/image is deemed a positive sample, while the unmatched texts/images within the mini-batch and all samples in the previous iterations are considered as negatives. The contrastive learning (*i.e.*, InfoNCE loss) is employed to maximize the similarity between positive image-text pairs while minimizing the similarity between negative pairs, which is formulated as:

$$\begin{aligned} \mathcal{L}_{I2T} &= -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(\mathbf{v}_i^\top \hat{\mathbf{t}}_i / \tau)}{\sum_{j=1}^{B+Nq} \exp(\mathbf{v}_i^\top \hat{\mathbf{t}}_j / \tau)}, \\ \mathcal{L}_{T2I} &= -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(\mathbf{t}_i^\top \hat{\mathbf{v}}_i / \tau)}{\sum_{j=1}^{B+Nq} \exp(\mathbf{t}_i^\top \hat{\mathbf{v}}_j / \tau)}, \end{aligned} \quad (1)$$

where τ is the temperature and B is the batch size. The total loss for image-text contrastive learning is defined as:

$$\mathcal{L}_{\text{align}} = (\mathcal{L}_{I2T} + \mathcal{L}_{T2I}) / 2. \quad (2)$$

The characteristics of independent encoding for each modality in the dual-encoder and the objective of contrastive learning lead to a moderate similarity distribution.

Cross-Encoder. Cross-encoders typically start by utilizing visual and textual encoders to extract detailed image and text representations \mathbf{V}_i and \mathbf{T}_j from the input image-text pair. Then the obtained primary representations are fed into a multi-modal encoder that employs self-attention or cross-attention to enable cross-modal interactions, thereby calculating the similarity score $\mathbf{p}_{i,j} \in [0, 1]$ for the input image-text pair. The image-text matching task serves as the training objective for the cross-encoder, which aims at judging whether the input image-text pair is matched or not and is formulated as:

$$\mathcal{L}_{itm} = -\frac{1}{|\mathcal{S}|} \sum_{(i,j) \in \mathcal{S}} y_{ij} \log \mathbf{p}_{i,j} + (1 - y_{ij}) \log(1 - \mathbf{p}_{i,j}),$$

where $y_{ij} \in \{0, 1\}$ is the ground-truth label indicating whether the image-text pair is matched or not. \mathcal{S} is the index set of image-text pairs, obtained by random sampling or hard negative sampling strategies. The objective of image-text matching pushes $\mathbf{p}_{i,j}$ to either 0 or 1, making the similarity distribution of cross-encoder more concentrated. The substantial difference between similarity distribution of dual-encoder and cross-encoder makes KL-divergence based distillation less effective.

3.2. Contrastive Partial Ranking Distillation

To achieve effective knowledge transfer from cross-encoder to dual-encoder, it is essential to first identify the source of performance gap between dual-encoder and cross-encoder. We test the performance of the pre-trained dual-encoder and cross-encoder from [15] on MSCOCO[19] dataset. Note that the retrieval results of cross-encoder are obtained by re-ranking top-K retrieved items from dual-encoder. The results are presented in Table 1, and we observe that (1) As K increases from 0 to 32, the retrieval performance significantly improves. (2) When more less-challenging negatives

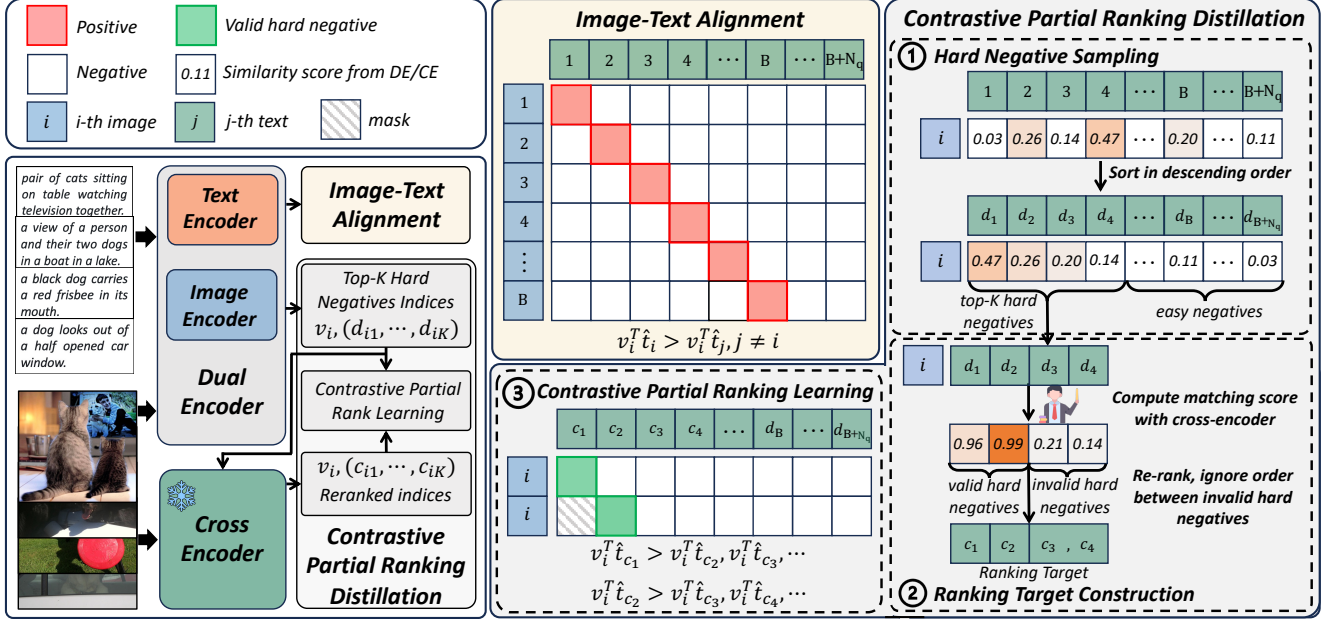


Figure 2. The illustration of the Contrastive Partial Ranking Distillation method. The left shows the overall training process and the right side elaborates on the computation process of image-text alignment and contrastive partial ranking distillation.

Table 1. Performance comparison between dual-encoder (DE) and cross-encoder (CE) with different candidate number K .

Model	image→text			text→image			R@S
	R@1	R@5	R@10	R@1	R@5	R@10	
DE	73.8	92.1	95.7	55.8	80.6	88.1	486.2
CE, K=0	73.8	92.1	95.7	55.8	80.6	88.1	486.2
CE, K=4	76.6	91.3	95.6	60.6	80.5	87.8	492.4
CE, K=16	78.1	93.0	95.8	61.4	83.5	89.3	501.1
CE, K=32	78.4	93.4	96.4	61.5	83.8	90.0	503.4
CE, K=128	78.6	93.9	96.9	61.4	83.7	89.9	504.4
CE, K=256	78.5	93.8	96.9	61.4	83.7	89.9	504.2

are introduced, *i.e.*, K continues to increase from 32 to 256, the performance saturates. It indicates that dual-encoder has filtered easy negatives and ranked most relevant items (including positives and hard negatives) into top 32, but lacks accurate ranking of them. Furthermore, given the original loss of the dual-encoder (InfoNCE loss) encompasses the objective of ranking positives ahead of hard negatives, the distillation only needs to focus on the cross-encoder’s ranking knowledge of hard negatives.

Based on the above findings, we propose Contrastive Partial Ranking Distillation (CPRD), a method that leverages contrastive learning to enforce consistency in the ranking of hard negatives between dual-encoders and cross-encoders. As shown in Figure 2, in order to transfer knowledge about hard negatives ranking effectively, our CPRD method first mining sufficient hard negative samples with dual-encoder’s similarity score, then construct the ranking target of hard negatives with cross-encoder, and finally perform partial ranking distillation via contrastive learning to transfer cross-modal matching knowledge effectively.

3.2.1 Hard Negative Mining

A straightforward approach to obtain hard negatives is selecting the most similar unmatched images/texts within the mini-batch. However, the batch size must be sufficiently large to find adequately challenging hard negatives. Therefore, we expand the mining scope to the queues \mathcal{Q}^v and \mathcal{Q}^t . Here, we demonstrate the process of image-to-text retrieval, while the text-to-image retrieval is performed symmetrically. For the i -th image, we use dual-encoder to compute its similarity scores with all negative texts in \mathcal{Q}^t , which contains samples from both current mini-batch and previous iterations. We then obtain the ranking result for the negative texts, denoted as $\mathbf{d}_i = \{d_{ij}\}_{j=1}^{B+N_q-1}$, by sorting the similarity scores in the descending order, where:

$$\mathbf{v}_i^\top \hat{\mathbf{t}}_{d_{ij}} > \mathbf{v}_i^\top \hat{\mathbf{t}}_{d_{ik}}, \forall j, k \in [1, \dots, B+N_q-1], j < k, \quad (3)$$

and indices of top- K hard negative texts is $\mathbf{h}_i = \{d_{ij}\}_{j=1}^K$.

3.2.2 Ranking Target Construction

Then we construct the ranking target of top- K hard negatives with cross-encoder. To compute the similarity scores of the top- K hard negatives with the cross-encoder, we also need to maintain two additional queues \mathcal{Q}_c^v and \mathcal{Q}_c^t with size N_c to preserve the image and text inputs from mini-batch and previous iterations for the cross-encoder. The similarity score set \mathbf{P}_i of the cross-encoder for all hard negative texts is denoted as $\mathbf{P}_i = \{p_{i,d_{ij}} | j \in [1, 2, \dots, K]\}$. The cross-encoder’s ranking of hard negative texts $\mathbf{c}_i = \{c_{ij}\}_{j=1}^K$ can then be obtained by sorting \mathbf{P}_i , where c_{ij} satisfies:

$$p_{i,c_{ij}} > p_{i,c_{ik}}, \forall j, k \in [1, \dots, K], j < k. \quad (4)$$

However, the cross-encoder’s ranking c_i can not be used as target directly because not all information in the rank-

ing c_i is beneficial for dual-encoder. There are some low-ranked negative texts in c_i that have small similarity score $p_{i,c_{ij}}$. We refer these texts as invalid hard negative texts $\{\hat{t}_{c_{ij}} | p_{i,c_{ij}} < m\}$, where m is the hyper-parameter to distinguish between valid and invalid hard negative texts. The relative order among invalid hard negative texts does not contain helpful knowledge, and enforcing the dual-encoder to learn such relative order may hurt the model performance. Therefore, we construct a partial ranking c_i^* as the target, where the relative order between valid hard negative texts is the same with c_i but the relative order between invalid ones is disregarded.

It is worth noting that using the cross-encoder to calculate similarity scores online brings additional training costs. To address this issue, we can employ an offline approach. Specifically, we first use the dual-encoder to perform a fast retrieval on the entire pre-training dataset to obtain the top- N hard negative texts for each image. Then we compute similarity scores of these negative image-text pairs with cross-encoder and store them in a similarity bank. During the training iterations, for the negative text $\hat{t}_{d_{ij}}$, if $p_{i,d_{ij}}$ is in the similarity bank, we load the pre-computed value; otherwise, we term $\hat{t}_{d_{ij}}$ as invalid negative texts. We set $N \gg K$ to ensure that most valid hard negatives are considered.

3.2.3 Contrastive Partial Ranking Learning

Given the ranking target c_i^* from cross-encoder, we expect the rankings from dual-encoder (h_i) can be *partial* consistent with c_i^* , which means only the relative order between valid hard negatives should be maintained and the relative order between invalid hard negatives is ignored. The learning objective is formulated as:

$$\min_{\theta} \text{Dist}(h_i, c_i^*), \quad (5)$$

where θ is the parameters of dual-encoder, $\text{Dist}(\cdot, \cdot)$ is the metric to measure the partial consistency between two ranking results. Optimizing this objective faces two challenges: (1) An appropriate metric to measure the partial consistency between rankings; (2) The sort operation used to compute c_i^* and h_i is non-differentiable, hindering the end-to-end training. Both are hard to tackle.

To implement effective partial ranking learning, we transform the non-differentiable ranking mimicking into the optimization of relative similarity scores. Specifically, for each valid hard negative text in c_i^* , we require its similarity with the i -th image computed by dual-encoder to be larger than the similarities of other negative texts ranked behind it. For invalid hard negative texts, we do not impose this constraint, which avoids learning relative order between them. The new objective can be formulated as:

$$v_i^\top t_{c_{ij}} > v_i^\top t_{c_{ik}}, \quad (6)$$

where J_i^* denotes the index of the first invalid hard negative text in the c_i^* .

Although Equation 6 can already be used as a differentiable loss function for optimizing the dual-encoder, we find that an inappropriate loss function type may interfere with the coarse image-text alignment learned by dual-encoder. Formulating the ranking learning process in the form of contrastive learning is more coordinated with training of dual-encoder and yields better results. Therefore, we propose a new contrastive partial ranking loss based on the InfoNCE loss, which is formulated as:

$$\mathcal{L}_{ij} = -\log \frac{\exp(v_i^\top \hat{t}_{c_{ij}}/\tau)}{\sum_{k=j}^K \exp(v_i^\top \hat{t}_{c_{ik}}/\tau) + \sum_{k=K+1}^{B+N_q-1} \exp(v_i^\top \hat{t}_{d_{ik}}/\tau)},$$

$$\mathcal{L}_{\text{CPRD}}^{\text{I2T}} = \frac{1}{B} \sum_{i=1}^B \frac{1}{J_i^* - 1} \sum_{j=1}^{J_i^* - 1} \mathcal{L}_{ij}, \quad (7)$$

where \mathcal{L}_{ij} has the same effect with Equation 6. As mentioned above, $\mathcal{L}_{\text{CPRD}}^{\text{T2I}}$ can be computed in the symmetrical way and the final contrastive partial ranking distillation loss $\mathcal{L}_{\text{CPRD}}$ is defined as:

$$\mathcal{L}_{\text{CPRD}} = (\mathcal{L}_{\text{CPRD}}^{\text{I2T}} + \mathcal{L}_{\text{CPRD}}^{\text{T2I}})/2. \quad (8)$$

3.3. Training Objectives

We perform distillation training with both image-text contrastive loss and new contrastive partial ranking distillation loss to learn the coarse image-text alignment and detailed ranking between hard negatives simultaneously. The total loss is formulated as:

$$\mathcal{L} = \mathcal{L}_{\text{align}} + \mathcal{L}_{\text{CPRD}}. \quad (9)$$

4. Experiments

4.1. Datasets

Pre-training Datasets. We pre-train our model with CC4M dataset, which contains 4 million images and 5.1 million captions from four public image-text datasets, including Conceptual Captions 3M[30], SBU[25], MSCOCO[19] and Visual Genome[12].

Downstream Datasets. We conduct downstream image-text retrieval evaluation on two widely used datasets: MSCOCO[19] and Flickr30K[27]. In addition, we validate the effectiveness of our method on improving the ranking ability with the CrissCrossed Caption[26] dataset. The details of these downstream datasets and the evaluation metrics can be found in the supplemental material.

4.2. Implementation Details.

For dual-encoder, we adopt BERT_{base}[10] as text encoder and a ViT-B/16[4] pre-trained on ImageNet-1k as the visual encoder. For cross-encoder, we experiment with three models with varying performance, including ALBEF[14] pre-trained on 5M image-text pairs, ALBEF pre-trained on

Table 2. The performance comparison with different queue size N_c . “DE” and “CE” represent the baseline dual-encoder and teacher cross-encoder.

N_c	image→text			text→image			R@S
	R@1	R@5	R@10	R@1	R@5	R@10	
DE	32.0	59.4	71.5	24.4	49.5	61.0	297.8
CE	40.5	66.5	76.1	30.2	55.3	66.5	335.1
0	32.6	60.1	71.0	24.7	49.8	61.7	299.9
1024	32.6	60.3	71.9	25.9	51.0	62.7	304.4
4096	33.7	60.6	71.7	26.9	52.3	64.0	309.2
16384	<u>34.3</u>	<u>61.4</u>	<u>73.2</u>	<u>27.0</u>	<u>52.8</u>	<u>64.5</u>	<u>313.2</u>
32768	34.6	62.6	73.5	27.3	52.4	63.8	314.2

Table 3. The performance comparison with different numbers of hard negative considered for contrastive partial ranking distillation. “None” indicates the offline approach for calculating hard negative pairs’ similarity with cross-encoder, which does not need to set K.

K	image→text			text→image			R@S
	R@1	R@5	R@10	R@1	R@5	R@10	
0	32.0	59.4	71.5	24.4	49.5	61.0	297.8
2	32.4	60.7	71.8	24.6	49.6	61.4	300.5
4	33.5	61.1	<u>72.9</u>	25.9	50.9	62.2	306.5
8	34.6	<u>61.3</u>	73.2	26.7	51.6	63.5	310.9
16	34.1	61.1	72.6	27.3	53.0	<u>64.3</u>	<u>312.4</u>
32	<u>34.3</u>	<u>61.4</u>	73.2	<u>27.0</u>	<u>52.8</u>	64.5	313.2
None	34.0	60.8	72.5	27.3	52.7	64.0	311.3

15M image-text pairs, and BLIP[15] pre-trained on >200M image-text pairs. We use the AdamW[20] optimizer with a weight decay of 0.02. The learning rate is warmed up to $3e-4$ in the first 2000 iterations and decays to $1e-5$ following a cosine schedule. We pre-train the model for 20 epochs with a batch size of 512 on 8 NVIDIA V100 GPUs. We take the image resolution of 256×256 for pre-training and increase the image resolution to 384×384 for fine-tuning. The momentum coefficient for updating momentum encoders is set as 0.995. The number of hard negative K and the threshold m is set as 16 and 0.75. The queue size N_q and N_c are set as 57856 and 16384 respectively. The learnable temperature hyper-parameter for contrastive loss is initialized to 0.07. More implementation details can be found in the supplementary materials.

4.3. Teach Dual-Encoder with Cross-Encoder

In this section, we provide detailed steps to demonstrate how to use the cross-encoder as teacher for the dual-encoder effectively. All models in this section are pre-trained on CC3M using an image resolution of 224×224 with 8 NVIDIA V100 GPUs, and then tested for zero-shot image-text retrieval performance on MSCOCO. The teacher model is ALBEF pre-trained on CC3M and is kept frozen during the distillation process, while the student model is a dual-encoder trained from scratch.

The effect of queue size N_c . The queue size N_c determines the number of negative samples, which affects the number of valid hard negatives and the final distillation per-

Table 4. The performance comparison with different valid hard negative threshold m .

m	image→text			text→image			R@S
	R@1	R@5	R@10	R@1	R@5	R@10	
1.0	32.0	59.4	71.5	24.4	49.5	61.0	297.8
0.9	33.3	60.7	72.0	26.8	<u>52.6</u>	<u>64.2</u>	309.6
0.75	34.3	<u>61.4</u>	73.2	27.0	52.8	64.5	313.2
0.5	<u>33.4</u>	61.5	<u>73.0</u>	27.0	52.3	63.6	<u>310.8</u>
0.0	30.8	57.2	69.4	24.1	48.4	60.5	290.4

Table 5. The performance comparison with different distillation methods.

Method	image→text			text→image			R@S
	R@1	R@5	R@10	R@1	R@5	R@10	
None	32.0	59.4	71.5	24.4	49.5	61.0	297.8
KL	32.5	59.6	71.4	26.0	50.9	62.9	303.3
M3SE	30.8	57.2	69.4	24.1	48.4	60.5	290.4
R-M3SE	32.8	60.4	71.0	25.6	50.5	62.0	302.3
CPRD $_{m^*}$	<u>33.4</u>	<u>60.8</u>	<u>72.4</u>	<u>26.3</u>	<u>51.5</u>	<u>63.4</u>	<u>307.7</u>
CPRD	34.3	61.4	73.2	27.0	52.8	64.5	313.2

formance. Therefore, choosing a sufficiently large queue size is crucial. As shown in Table 2, when we set the queue size to 0, only the negative samples in the current batch are utilized for distillation, resulting in a limited number of valid hard negatives and less valuable knowledge. Consequently, the performance improvement compared to the baseline is relatively weak. When we increase the queue size to 1024, the number of valid hard negatives increases, and the model performance achieve significant improvement. As the queue size further increases, the model performance gradually improves and tends to saturate. This is because, in the training dataset, the number of valid hard negatives for a specific sample is limited. Further increasing the size of the queue leads to a saturation in the number of valid hard negatives. We set the queue size to 16384 by default.

The effect of hard negatives number K. Similarly, the choice of K also affects the number of valid hard negatives and the effectiveness of knowledge distillation. As shown in Table 3, when set K to 0, we do not consider the ranking of any negative samples, causing the model to degrade to the baseline model. Gradually increasing K leads to more valid hard negative samples being considered, making the distillation more effective and the model performance improves. When K is increased to 32, the model performance tends to saturate, as most of the valid hard negatives are considered. Moreover, increasing K results in heavier computation, thus we set K to 16 by default. It is worth noting that we can calculate the similarity of hard negative pairs in an offline manner to avoid extra training costs. For each image/text in the training dataset, we use the cross-encoder to calculate the scores of the top-1000 hard negative sample pairs as mentioned in Section 3.2.2. As shown in Table 3, the offline calculation approach achieves comparable performance with the online approach. It should be noted that when using the offline approach, we directly search for valid

Table 6. Comparative results for fine-tuned image-text retrieval results on the Flickr30K (1K) test set and MSCOCO (5K) test set. We make comparisons with dual-encoder methods and cross-encoder methods. “DE” and “CE” represent the baseline dual-encoder and the teacher cross-encoder. Our method improves the performance of baseline model significantly, surpasses previous state-of-the-art dual-encoder methods by a large margin, and achieves comparable performance with some cross-encoder methods while keeping the high retrieval efficiency. **Higher R@K indicates better performance. PT Pairs:** the number of image-text pairs for pre-training. † is ensemble result of two models. * models use 940M tagged images for visual encoder pre-training.

Model	PT Pairs	Flickr30K (1K test set)						MSCOCO (5K test set)							
		image→text			text→image			R@S	image→text			text→image			R@S
R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5		R@10	R@1	R@5	R@10	R@1	R@5	
Cross-Encoder															
Pixel-BERT-X152 [8]	5.6M	87.0	98.9	99.5	71.5	92.1	95.8	544.8	63.6	87.5	93.6	50.1	77.6	86.2	458.6
VILLA-base [6]	9.6M	86.6	97.9	99.2	74.7	92.9	95.8	547.1	-	-	-	-	-	-	-
Oscar-base [18]	6.5M	-	-	-	-	-	-	-	70.0	91.1	95.5	54.0	80.8	88.5	479.9
ViLT [11]	9.9M	83.5	96.7	98.6	64.4	88.7	93.8	525.7	61.5	86.3	92.7	42.7	72.9	83.1	439.2
VinVL-base [35]	8.9M	-	-	-	-	-	-	-	74.6	92.6	96.3	58.1	83.2	90.1	494.9
ALBEF [14]	5.1M	94.3	99.4	99.8	82.8	96.7	98.4	571.4	73.1	91.4	96.0	56.8	81.5	89.2	488.0
Dual-Encoder															
VSE ∞ *† [2]	-	88.7	98.9	99.8	76.1	94.5	97.1	555.1	68.1	90.2	95.2	52.7	80.2	88.3	474.7
COOKIE*† [33]	5.9M	89.0	98.9	99.7	75.6	94.6	97.2	555.0	71.6	90.9	95.4	54.5	81.0	88.2	481.6
LightingDOT [31]	9.5M	83.9	97.2	98.6	69.9	91.1	95.2	535.9	60.1	85.1	91.8	45.8	74.6	83.8	441.2
COOKIE [33]	5.9M	84.7	96.9	98.3	68.3	91.1	95.2	534.5	61.7	86.7	92.3	46.6	75.2	84.1	446.6
COTS [21]	5.3M	88.2	98.5	99.7	75.2	93.6	96.5	551.7	66.9	88.8	94.0	50.5	77.6	86.1	463.9
COTS [21]	15.3M	90.6	98.7	99.7	76.5	93.9	96.6	556.0	69.0	90.4	94.9	52.4	79.0	86.9	472.6
DE	5.1M	89.1	98.8	99.7	74.1	92.8	96.2	550.7	66.7	88.9	94.4	48.5	76.7	85.1	460.3
+CPRD	5.1M	90.8	98.7	99.4	76.4	93.9	96.8	556.0	69.7	91.1	95.8	52.1	78.9	86.8	474.4
CE	5.1M	94.5	99.0	99.7	78.6	94.3	97.1	563.2	74.4	92.3	96.5	58.0	82.8	89.7	493.7
+CPRD	5.1M	90.7	99.0	99.7	78.6	94.9	97.4	560.3	70.8	91.7	96.2	53.4	80.6	88.5	481.2
CE	15.2M	95.9	99.8	100.	85.6	97.5	98.9	577.7	77.6	94.3	97.2	60.7	84.3	90.5	504.6

hard negatives from the current batch and queue without the need of setting K.

The effect of valid hard negative threshold m . The threshold m controls the difficulty of negative samples for which dual-encoder needs to learn their relative order. As shown in Table 4, when setting $m = 0$, all relative order among top-K hard negatives are learned by dual-encoder, and the performance declines significantly because dual-encoder is also enforced to learn the relative order between easy negatives, which contains no valid knowledge and introduces interference for the dual-encoder. When setting $m = 1$, none of negative samples are considered in the partial ranking learning, and thus the model degenerates into the baseline model. We also empirically test with $m = 0.5, 0.75, 0.9$. We find that the performance of all three models is better than the baseline model, showing the robustness of our method toward the threshold m . The best result is achieved when $m = 0.75$. Therefore, we set the default value of m to 0.75. These experimental results validate our findings that only relative order between valid hard negatives conveys valuable knowledge.

Comparison with Different Distillation Methods. Based on the optimal setting identified through the above experiments, we explore different distillation methods to transfer knowledge from to dual-encoder. KL is the KL-divergence loss which constrains the consistency between similarity score distribution. M3SE[23] requires that (1) the similarity margin between positive and hardest negative should be same for dual-encoder and cross-encoder, (2) similarity of

Table 7. Comparison for image-text retrieval results (without fine-tuning) on the MSCOCO (5K) test set.

Model	image→text			text→image			R@S
	R@1	R@5	R@10	R@1	R@5	R@10	
Pre-train from Scratch							
CLIP [28]	58.4	81.5	88.1	37.8	62.4	72.2	400.4
ALIGN [9]	58.6	83.0	87.9	45.6	69.8	78.6	423.5
COTS [21]	60.4	84.7	91.7	43.8	<u>71.6</u>	<u>81.3</u>	433.5
DE	<u>61.7</u>	<u>86.0</u>	<u>92.3</u>	41.5	70.4	80.6	432.5
+CPRD	62.4	86.5	93.0	<u>45.1</u>	74.1	83.6	444.7
CE	73.2	91.5	95.7	54.8	79.7	87.0	481.9
Post Pre-train							
DE	65.2	88.4	93.5	50.0	76.2	84.8	458.1
+CPRD	65.5	88.5	93.6	51.7	77.8	86.0	463.1
CE	71.9	91.1	95.2	58.7	81.7	88.5	487.1

other negative samples should be smaller than that of hardest negative.

As shown in Table 5, KL method achieves limited performance improvement due to the difference between similarity score distribution of dual-encoder and cross-encoder. M3SE even results in performance decline because the similarity range is different for dual-encoder and cross-encoder. To tackle this problem, we propose a modification R-M3SE to first re-scale similarity score with min-max normalization, and we observe the retrieval performance is improved but still worse than our method. To validate whether the form of loss function affects the distillation, we propose a variant of our method CPRD $_{m^*}$, which adjusts the threshold m adaptively to ensure there always has only one valid hard negative. CPRD $_{m^*}$ has the highly-similar objective with R-M3SE but is implemented with infoNCE loss. It can be seen

that CPRD_{m*} surpasses R-M3SE on all metrics, validating that it is beneficial to make the distillation loss coordinated with the training of dual-encoder.

4.4. Comparison with SOTA

Image-Text Retrieval. We compare with state-of-the-art image-text retrieval methods on Flickr30K and MSCOCO datasets. The experimental results under fine-tuning setting are shown in Table 6. Compared with our dual-encoder baseline, our proposed distillation method achieves significant improvement. On the Flickr30K dataset, we achieve higher performance by 1.7% and 2.3% on the R@1 of image-to-text and text-to-image retrieval. On the MSCOCO dataset, we also surpass the baseline by 3.0% and 3.6% on the R@1 of two retrieval tasks. We also experiment with a stronger cross-encoder trained with 15.2M data as teacher, the retrieval performance is further improved, showing that a stronger teacher can transfer more valid knowledge to student even with the same training data.

Under a fair comparison experimental setting (excluding VSE_∞*[†] and COOKIE*[†] as they use 940M tagged images for visual-encoder pre-training), our method also outperforms other dual-encoder methods by a large margin under all evaluation metrics. Specifically, compared with the current state-of-the-art dual-encoder method COTS [21] with 5.3M pre-training data, our method with similar data size achieves higher performance by 1.8% and 1.6% on the R@1 of MSCOCO dataset. Moreover, our method with 5.1M pre-training data outperforms COTS pre-trained on 15.3M image-text pairs. Furthermore, our method also achieves comparable performance with the cross-encoder methods VinVL-base and ALBEF while much more efficient.

The experimental results without fine-tuning are shown in Table 7. Our distillation method improves the performance of dual-encoder on all evaluation metrics. With a similar pre-training data size, our method outperforms the COTS [21] by 2.0% and 1.3% on the R@1 of two retrieval tasks. Our method also outperforms CLIP [28] and achieves comparable performance with ALIGN [9], which utilize 78× and 356× pre-training data than our method respectively. Moreover, we also evaluate our method under post pre-train setting, where we initialize our dual-encoder with the pre-trained weights from BLIP-129M and continue to pre-train the model with image-text contrastive loss and our distillation loss. Our method can further improve the performance of the pre-trained model with only a small additional computation cost (10k iterations).

Image-Text Ranking. To further validate the effectiveness of our CPRD method on improving the ranking ability of dual-encoder, we perform evaluation on the Criss-Crossed caption dataset for image-text ranking task. We report the Spearman’s R bootstrapped correlation for image-text ranking task, *i.e.*, “SITS” in Table 8, which reflects

Table 8. Spearman’s R Bootstrap Correlation (×100) on Criss-crossed Captions dataset. “FT” indicates the model is fine-tuned on the MSCOCO dataset. **STS**, **SIS** and **SITS** represent the task of semantic text similarity, semantic image similarity and semantic image-text similarity.

Model	STS	SIS	SITS	Mean Avg
	avg ± std	avg ± std	avg ± std	
VSE++[5]	74.4±0.4	73.3±0.9	55.2±1.5	67.6
VSRN [16]	73.0±0.4	70.1±1.0	60.4±1.3	67.8
DE _{I2T} [26]	50.9±0.6	81.3±0.7	61.6±1.4	64.6
DE _{T2T+I2T} [26]	74.2±0.4	74.5±0.9	61.9±1.3	70.2
ALIGN[14]	72.9±0.4	77.2±0.8	67.6±1.2	72.6
DE	74.1±0.4	75.1±0.8	61.8±1.4	70.3
+CPRD	74.9±0.3	75.6±0.8	64.3±1.3	71.6
CE	–	–	67.3±1.2	–
DE-FT	73.9±0.4	75.7±0.8	65.1±1.3	71.5
+CPRD-FT	74.4±0.4	76.1±0.9	<u>66.9±1.2</u>	<u>72.5</u>
CE-FT	–	–	69.4±1.1	–

the consistency between human ranking and model ranking. Compared with vanilla dual-encoder, our CPRD method achieves 2.5 improvement. The performance of our method is further improved after fine-tuning on the MSCOCO dataset, and achieves comparable results with ALIGN. In addition, we also report the results on text-text similarity ranking (**STS**) and image-image similarity ranking (**SIS**). It is worth noting that although our CPRD method only distills the knowledge of cross-modal matching, it can also improve the accuracy of measuring intra-modal similarity.

5. Conclusion

In this work, we investigate how to effectively distill cross-modal knowledge from cross-encoder to dual-encoder for image-text retrieval. We identify three key factors and propose a novel Contrastive Partial Ranking Distillation(CPRD) method. Our method focuses on learning relative order among valid hard negatives while disregarding relative order among invalid hard negatives and easy negatives. We implement our method with contrastive learning, which aligns with training of dual-encoder and transfers knowledge effectively without disrupting the learning of image-text alignment. Comprehensive experiments on image-text retrieval and ranking show the superiority of our method compared to other distillation methods. Moreover, our method significantly improves retrieval and ranking accuracy of dual-encoder under various experiment settings.

Acknowledgments This work is supported by the National Science and Technology Major Project (Grant No.2022ZD0118501), Beijing Natural Science Foundation (Grant No. JQ21017, L223003), the Natural Science Foundation of China (Grant No. 62222206, 62036011, 62192782, 62302501, 62202470, U2033210), The Project of Beijing Science and technology Committee(Project No.Z231100005923046), Research Fund of ARC Lab, Tencent PCG.

References

- [1] Hangbo Bao, Wenhui Wang, Li Dong, Qiang Liu, Owais Khan Mohammed, Kriti Aggarwal, Subhojit Som, Songhao Piao, and Furu Wei. Vlmo: Unified vision-language pre-training with mixture-of-modality-experts. *Advances in Neural Information Processing Systems*, 35:32897–32912, 2022. 2
- [2] Jiacheng Chen, Hexiang Hu, Hao Wu, Yuning Jiang, and Changhu Wang. Learning the best pooling strategy for visual semantic embedding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15789–15798, 2021. 7
- [3] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Learning universal image-text representations. 2019. 2
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 2, 5
- [5] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improving visual-semantic embeddings with hard negatives. 8
- [6] Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. Large-scale adversarial training for vision-and-language representation learning. *Advances in Neural Information Processing Systems*, 33:6616–6628, 2020. 7
- [7] Sebastian Hofstätter, Sophia Althammer, Michael Schröder, Mete Sertkan, and Allan Hanbury. Improving efficient neural ranking models with cross-architecture knowledge distillation. *arXiv preprint arXiv:2010.02666*, 2020. 3
- [8] 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. 7
- [9] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. 1, 2, 7, 8
- [10] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019. 2, 5
- [11] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning*, pages 5583–5594. PMLR, 2021. 7
- [12] 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*, 123:32–73, 2017. 5
- [13] Jie Lei, Xinlei Chen, Ning Zhang, Mengjiao Wang, Mohit Bansal, Tamara L Berg, and Licheng Yu. Loopitr: Combining dual and cross encoder architectures for image-text retrieval. *arXiv preprint arXiv:2203.05465*, 2022. 3
- [14] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021. 2, 3, 5, 7, 8
- [15] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 2, 3, 6
- [16] Kunpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. Visual semantic reasoning for image-text matching. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4654–4662, 2019. 8
- [17] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. 2
- [18] Xiujuan 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 *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX 16*, pages 121–137. Springer, 2020. 2, 7
- [19] 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 *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 3, 5
- [20] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. 6
- [21] Haoyu Lu, Nanyi Fei, Yuqi Huo, Yizhao Gao, Zhiwu Lu, and Ji-Rong Wen. Cots: Collaborative two-stream vision-language pre-training model for cross-modal retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15692–15701, 2022. 1, 2, 7, 8
- [22] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32, 2019. 2
- [23] Aditya Menon, Sadeep Jayasumana, Ankit Singh Rawat, Seungyeon Kim, Sashank Reddi, and Sanjiv Kumar. In defense of dual-encoders for neural ranking. In *International Conference on Machine Learning*, pages 15376–15400. PMLR, 2022. 3, 7
- [24] Antoine Miech, Jean-Baptiste Alayrac, Ivan Laptev, Josef Sivic, and Andrew Zisserman. Thinking fast and slow: Efficient text-to-visual retrieval with transformers. In *Proceed-*

- ings of the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9826–9836, 2021. [2](#), [3](#)
- [25] Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2text: Describing images using 1 million captioned photographs. *Advances in neural information processing systems*, 24, 2011. [5](#)
- [26] Zarana Parekh, Jason Baldridge, Daniel Cer, Austin Waters, and Yinfei Yang. Crisscrossed captions: Extended intramodal and intermodal semantic similarity judgments for ms-coco. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2855–2870, 2021. [5](#), [8](#)
- [27] 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 *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649, 2015. [5](#)
- [28] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [1](#), [2](#), [7](#), [8](#)
- [29] Sashank Reddi, Rama Kumar Pasumarthi, Aditya Menon, Ankit Singh Rawat, Felix Yu, Seungyeon Kim, Andreas Veit, and Sanjiv Kumar. Rankdistil: Knowledge distillation for ranking. In *International Conference on Artificial Intelligence and Statistics*, pages 2368–2376. PMLR, 2021. [3](#)
- [30] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018. [5](#)
- [31] 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 *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 982–997, 2021. [2](#), [7](#)
- [32] Zekun Wang, Wenhui Wang, Haichao Zhu, Ming Liu, Bing Qin, and Furu Wei. Distilled dual-encoder model for vision-language understanding. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8901–8913, 2022. [2](#)
- [33] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2208–2217, 2021. [1](#), [2](#), [7](#)
- [34] Yan Zeng, Xinsong Zhang, and Hang Li. Multi-grained vision language pre-training: Aligning texts with visual concepts. *arXiv preprint arXiv:2111.08276*, 2021. [2](#)
- [35] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5579–5588, 2021. [2](#), [7](#)