

ConaCLIP: Exploring Distillation of Fully-Connected Knowledge Interaction Graph for Lightweight Text-Image Retrieval

Jiapeng Wang^{1*} Chengyu Wang^{2†} Xiaodan Wang³ Jun Huang² Lianwen Jin^{1†}

¹South China University of Technology, Guangzhou, China

²Alibaba Group, Hangzhou, China

³Fudan University, Shanghai, China

¹eeljpwang@mail.scut.edu.cn, eelwjin@scut.edu.cn

²{chengyu.wcy, huangjun.hj}@alibaba-inc.com

³xiaodanwang20@fudan.edu.cn

Abstract

Large-scale pre-trained text-image models with dual-encoder architectures (such as CLIP (Radford et al., 2021)) are typically adopted for various vision-language applications, including text-image retrieval. However, these models are still less practical on edge devices or for real-time situations, due to the substantial indexing and inference time and the large consumption of computational resources. Although knowledge distillation techniques have been widely utilized for uni-modal model compression, how to expand them to the situation when the numbers of modalities and teachers/students are doubled has been rarely studied. In this paper, we conduct comprehensive experiments on this topic and propose the fully-Connected knowledge interaction graph (Cona) technique for cross-modal pre-training distillation. Based on our findings, the resulting ConaCLIP achieves SOTA performances on the widely-used Flickr30K and MSCOCO benchmarks under the lightweight setting. An industry application of our method on an e-commercial platform further demonstrates the significant effectiveness of ConaCLIP.¹

1 Introduction

Text-image retrieval (TIR) aims at retrieving a list of the most relevant images from a large image collection when a specific text query is given. With the rapid development of information interaction and social intercourse, it has been regarded as a crucial component of cross-modal applications and required by various real-world scenarios, such as e-commercial platforms (sites).

Recently, inspired by the great success of pre-trained language models (Devlin et al., 2019; Liu

et al., 2019; Brown et al., 2020), research on large-scale vision-language pre-training (Tan and Bansal, 2019; Li et al., 2020; Radford et al., 2021; Li et al., 2022; Wang et al., 2022b, 2023) has achieved remarkable progress on a variety of vision-language tasks, including text-image retrieval. These existing methods can be typically classified into two categories according to the model architecture: *cross-encoder* and *dual-encoder*. *Cross-encoder* typically adds additional Transformer (Vaswani et al., 2017) layers to model the deep interaction between image and text representations. It can generally boost the retrieval performance, while resulting in an unbearably slow retrieval speed when applied to the entire image collection since the cross-modal costs are required for each image sample whenever a new text query is given. In contrast, *dual-encoder* encodes the visual and textual inputs in a wholly decoupled manner. The image representation is allowed to be pre-computed and re-used independent of the text queries. Such approaches can also utilize fast approximate nearest neighbor (ANN) search (Muja and Lowe, 2009; Jegou et al., 2010; Johnson et al., 2019) at runtime.

Although dual-encoder is usually preferred for real-world applications, the existing related models such as CLIP (Radford et al., 2021) are still less practical on edge devices with limited computing resources, or for the dynamic indexing scenario, e.g., private photos/messages collections (sites). To address this issue, we aim to start from the large-scale pre-trained dual-encoder models and focus on the pre-training distillation to present a series of much smaller, faster, and effective counterparts. Knowledge distillation (KD) (Hinton et al., 2014) is proposed to transfer knowledge with soft targets from a teacher to a student in the same modality. MoTIS (Ren and Zhu, 2022) simply repeats intra-modal InfoNCE-based (Oord et al., 2018) learning in both text and image domains for distillation. Nevertheless, when the number of modalities dou-

*Contribution during internship at Alibaba Group.

†Co-corresponding authors.

¹Related resources will be publicly available in the EasyNLP framework (Wang et al., 2022a). URL: <https://github.com/alibaba/EasyNLP>.

bles for dual-encoder structure, which means text and image teachers as well as text and image students, these methods still only involve intra-modal teacher-student knowledge interaction learning. Instead, in this paper, we comprehensively explore the fully-Connected knowledge interaction graph (Cona) between every possible teacher-student or student-student pair. As shown in Fig. 1, each two-way arrow represents the knowledge interaction learning between the two models it points to. And the aforementioned KD and MoTIS belong to a single *blue* arrow and the two *blue* arrows, respectively. Moreover, in order to better explore the potential of Cona, we implement and investigate various supervision strategies to guide the model optimization, which finally makes each type of learning contribute to the overall improvement.

We release various sizes of lightweight dual-encoder models named ConaCLIP for different real-world scenarios. Compared with the previous SOTA method (Ren and Zhu, 2022), our ConaCLIP achieves 10.6/12.9/12.8 R@1 gains on Flickr30K/MS-COCO (1K)/MS-COCO (5K) benchmarks under the same model setting. We have also verified its effectiveness on an e-commerce platform. It can achieve $1.44\times/1.92\sim 4.86\times$ inference speed-up with competitive performances given image/text queries. The main contributions of this paper can be summarized as follows:

- We propose a new pre-training distillation method with the fully-connected knowledge interaction graph (Cona) for lightweight dual-encoder models.
- We release a series of lightweight ConaCLIP models to the open-source community, which can significantly surpass previous SOTA models on the widely-used Flickr30K and MS-COCO benchmarks.
- We provide a real-world application of this method in real industrial scenarios to further demonstrate its practical values.

2 Related Work

Cross-encoder (Tan and Bansal, 2019; Li et al., 2019; Chen et al., 2020; Li et al., 2020; Chen et al., 2022) refers to multiple layers of dense cross-modal interactions, e.g., cross-attention (Vaswani et al., 2017), are typically employed to image and text representations for more fine-grained merge

and alignment. Although it often achieves superior retrieval accuracy thanks to the patch/token-level integration, the high memory cost and computation inefficiency make it impractical under time-critical real-world settings.

Oppositely, for dual-encoder (Zhang et al., 2020; Jia et al., 2021; Radford et al., 2021; Dou et al., 2022), image and text features are encoded into a joint embedding space separately, and the modality interaction is only handled by a simple cosine similarity of the final image and text feature vectors. Such approaches can be regarded as scalable and indexable: the specific choices of encoder architectures can be independent and dynamic, and the late-interaction scheme allows for efficient large-scale searching.

Pre-training distillation for lightweight dual-encoder architecture has been rarely studied. Vanilla knowledge distillation (Hinton et al., 2014) can be referred to as the knowledge transfer from a teacher to a student in the same modality based on soft targets. However, it is a general procedure without awareness and pertinence for cross-modal learning. MoTIS (Ren and Zhu, 2022) separately compresses text or image encoder with an intra-modal contrastive objective that aligns the output embeddings of the student and teacher of each modality, which can be seen as an alternative form of knowledge distillation. Nevertheless, these methods ignore or do not find an appropriate approach to leverage the cross-modal distillation process. Further than them, our method is dedicated to exploring the fully-connected knowledge interaction graph for dual-encoder distillation, which is a natural and effective extension.

3 Methodology

In this section, we first give the preliminary knowledge, then propose our pre-training distillation framework with Cona. Finally, we introduce various supervision strategies.

3.1 Preliminary

For the sake of explanation, we abbreviate *text*, *image*, *teacher* and *student* as T , I , tch and stu respectively. F represents the L2-normalized feature vector outputted by the encoder architecture E .

Before student learning, the teachers E_{tch}^T and E_{tch}^I are commonly first pre-trained using an objective that pushes the embeddings of matched text-image pairs closer while pushing those of non-

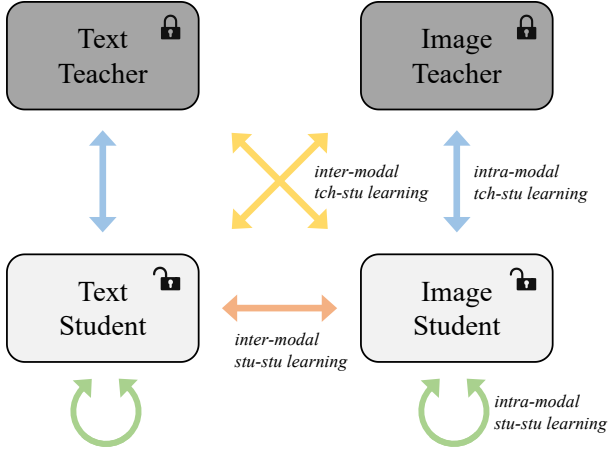


Figure 1: Our dual-encoder pre-training distillation framework with Cona. Each color of two-way arrows represents a type of knowledge interaction learning. At this stage, the teacher encoders are frozen.

matched ones apart, with large model capacity and massive data. Specifically, CLIP (Radford et al., 2021) takes the InfoNCE (Oord et al., 2018) loss as the supervision form. Without losing generality, given two outputted feature vectors F^a and $F^b \in \mathcal{R}^{N \times d}$, we define that:

$$p_{i,j}(F^a, F^b) = \frac{\exp(F_i^a F_j^{b\top} / \tau)}{\sum_k \exp(F_i^a F_k^{b\top} / \tau)}, \quad (1)$$

$$\mathcal{L}_{F^a \rightarrow F^b}^{\text{InfoNCE}} = -\frac{1}{N} \sum_{i=1}^N \log(p_{i,i}(F^a, F^b)), \quad (2)$$

where N is the mini-batch size, d is the channel size and τ is the temperature hyper-parameter. The final loss of CLIP can be formulated as:

$$\mathcal{L}^{\text{CLIP}} = \mathcal{L}_{F_{\text{tch}}^T \rightarrow F_{\text{tch}}^I}^{\text{InfoNCE}} + \mathcal{L}_{F_{\text{tch}}^I \rightarrow F_{\text{tch}}^T}^{\text{InfoNCE}}. \quad (3)$$

Next, the pre-training distillation of students E_{stu}^T and E_{stu}^I begins, with parameters of teachers E_{tch}^T and E_{tch}^I frozen. MoTIS (Ren and Zhu, 2022) also adopts the InfoNCE-based loss at this stage, and implements it in both text and image domains separately:

$$\mathcal{L}^{\text{MoTIS}} = \mathcal{L}_{F_{\text{stu}}^T \rightarrow F_{\text{tch}}^T}^{\text{InfoNCE}} + \mathcal{L}_{F_{\text{stu}}^I \rightarrow F_{\text{tch}}^I}^{\text{InfoNCE}}. \quad (4)$$

According to the subscript in Eq. (4), it is easy to see that MoTIS only involves *intra-modal teacher-student* learning.

3.2 Pre-training Distillation with Cona

Unlike existing works, our method introduces the fully-connected knowledge interaction graph

(Cona) for pre-training distillation. Apart from *intra-modal teacher-student* learning, our method also includes *intra-modal student-student* learning, *inter-modal teacher-student* learning and *inter-modal student-student* learning, as shown in Fig. 1. This fully-connected learning graph established for students E_{stu}^T and E_{stu}^I serves as an integration of multi-view and multi-task learning schemes, which can strengthen the robustness and effectiveness (Caruana, 1997; Luong et al., 2016; Aghajanyan et al., 2021) required by pre-trained models.

We suggest that each type of learning process in Cona should be concretely implemented in detailed supervision strategies. Therefore, we propose and investigate various supervision strategies in the next subsection.

3.3 Supervision Strategies

Here we continue to use F^a and F^b (prediction) along with \widetilde{F}^a and \widetilde{F}^b (target) as placeholders for illustration, and present the following effective supervision strategies:

InfoNCE loss is a type of contrastive loss function. It has been formulated in Eq. (2), and successfully applied for distillation by Eq. (4).

Feature-wise distance (FD) loss directly minimizes the distance between feature vectors. We utilize squared L2-norm as the measure:

$$\mathcal{L}_{F^a \leftrightarrow F^b}^{\text{FD}} = \frac{1}{2} \frac{1}{Nd} \sum_{i=1}^N \sum_{j=1}^d (F_{i,j}^a - F_{i,j}^b)^2. \quad (5)$$

Similarity-wise distance (SD) loss minimizes the distance criterion between similarity matrices:

$$\mathcal{L}_{F^a \rightarrow F^b \leftrightarrow \widetilde{F}^a \rightarrow \widetilde{F}^b}^{\text{SD}} = \frac{1}{2} \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N (F_i^a F_j^{b\top} - \widetilde{F}_i^a \widetilde{F}_j^{b\top})^2. \quad (6)$$

Since F^a , F^b , \widetilde{F}^a and \widetilde{F}^b have been L2-normalized, the values of cosine-similarities $F_i^a F_j^{b\top}$ and $\widetilde{F}_i^a \widetilde{F}_j^{b\top}$ are in the range $[-1, 1]$. The distance between prediction $F_i^a F_j^{b\top}$ and target $\widetilde{F}_i^a \widetilde{F}_j^{b\top}$ needs to be shortened. Hence, the squared L2-norm is also adopted here.

KL-Div loss uses the Kullback–Leibler divergence to measure the difference between the predicted and the target probability distributions. Given $p_{i,j}$ acquired by softmax operation shown in Eq. (1), it

Learning Type	Supervision Strategies					
	InfoNCE	FD	SD	KL-Div	Sym-SD	Sym-KL-Div
intra-modal stu-stu learning	\	\	$\mathcal{L}_{F_{stu}^I \rightarrow F_{stu}^T \Leftrightarrow F_{tch}^I \rightarrow F_{tch}^T}^{SD}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \Leftrightarrow F_{tch}^I \rightarrow F_{tch}^I}^{SD}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^T \rightarrow F_{tch}^T}^{KL-Div}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^I \rightarrow F_{tch}^I}^{KL-Div}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \Leftrightarrow F_{tch}^I \rightarrow F_{tch}^I}^{SD}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^I \rightarrow F_{tch}^I}^{KL-Div}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^T \rightarrow F_{tch}^T}^{KL-Div}$
inter-modal stu-stu learning	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I}^{InfoNCE}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I}^{InfoNCE}$	$\mathcal{L}_{F_{stu}^I \Leftrightarrow F_{tch}^I}^{FD}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \Leftrightarrow F_{tch}^I \rightarrow F_{tch}^I}^{SD}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \Leftrightarrow F_{tch}^I \rightarrow F_{tch}^I}^{SD}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^T \rightarrow F_{tch}^T}^{KL-Div}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^I \rightarrow F_{tch}^I}^{KL-Div}$	\	\
intra-modal tch-stu learning	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I}^{InfoNCE}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I}^{InfoNCE}$	$\mathcal{L}_{F_{stu}^I \Leftrightarrow F_{tch}^I}^{FD}$ $+\mathcal{L}_{F_{stu}^I \Leftrightarrow F_{tch}^I}^{FD}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \Leftrightarrow F_{tch}^I \rightarrow F_{tch}^I}^{SD}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \Leftrightarrow F_{tch}^I \rightarrow F_{tch}^I}^{SD}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^T \rightarrow F_{tch}^T}^{KL-Div}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^I \rightarrow F_{tch}^I}^{KL-Div}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \Leftrightarrow F_{tch}^I \rightarrow F_{tch}^I}^{SD}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^I \rightarrow F_{tch}^I}^{KL-Div}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^T \rightarrow F_{tch}^T}^{KL-Div}$
inter-modal tch-stu learning	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I}^{InfoNCE}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I}^{InfoNCE}$	$\mathcal{L}_{F_{stu}^I \Leftrightarrow F_{tch}^I}^{FD}$ $+\mathcal{L}_{F_{stu}^I \Leftrightarrow F_{tch}^I}^{FD}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \Leftrightarrow F_{tch}^I \rightarrow F_{tch}^I}^{SD}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \Leftrightarrow F_{tch}^I \rightarrow F_{tch}^I}^{SD}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^T \rightarrow F_{tch}^T}^{KL-Div}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^I \rightarrow F_{tch}^I}^{KL-Div}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \Leftrightarrow F_{tch}^I \rightarrow F_{tch}^I}^{SD}$	$\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^I \rightarrow F_{tch}^I}^{KL-Div}$ $+\mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^I \parallel F_{tch}^T \rightarrow F_{tch}^T}^{KL-Div}$

Table 1: Detailed loss functions of all combinations of knowledge interaction learning and supervision strategies. "Sym-" is the symmetric version loss function. "\ " indicates the combination is meaningless.

minimizes the following optimization objective:

$$\mathcal{L}_{F^a \rightarrow F^b \parallel \widetilde{F^a} \rightarrow \widetilde{F^b}}^{KL-Div} = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^N p_{i,j}(F^a, F^b) \log \frac{p_{i,j}(F^a, F^b)}{p_{i,j}(\widetilde{F^a}, \widetilde{F^b})}. \quad (7)$$

It is worth noting that, when performing the learning process indicated by an arrow shown in Fig. 1, the common practice is to use teachers' outputs F_{tch}^T and F_{tch}^I as target in Eq. (6)(7) that students learn from. While in our case with two modalities available, we propose to use the paired arrow as the target, and we call this the **symmetric version** (for SD loss and KL-Div loss). For example, inter-modal teacher-student learning implemented with KL-Div loss can be formulated as

$$\mathcal{L}_{F_{stu}^T \rightarrow F_{tch}^I \parallel F_{tch}^T \rightarrow F_{tch}^I}^{KL-Div} + \mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^T \parallel F_{tch}^I \rightarrow F_{tch}^T}, \quad (8)$$

while its symmetric version is

$$\mathcal{L}_{F_{stu}^T \rightarrow F_{tch}^I \parallel F_{stu}^I \rightarrow F_{tch}^T}^{KL-Div} + \mathcal{L}_{F_{stu}^I \rightarrow F_{tch}^T \parallel F_{stu}^T \rightarrow F_{tch}^I}. \quad (9)$$

This modification deepens the interaction between the four encoders during optimization.

So far, any one of the learning types can be concretely implemented by any one of the supervision strategies, except for a few meaningless combinations. Detailed loss functions are listed in Tab. 1.

4 Experiments

4.1 Setup

We use Conceptual Caption (CC3M) (Sharma et al., 2018) and Conceptual 12M (CC12M) (Changpinyo et al., 2021) for pre-training distillation, which consist of 3M and 12M noisy text-image pairs respectively. During fine-tuning, we use MSCOCO (Lin

et al., 2014) and Flickr30K (Plummer et al., 2015) as benchmarks. MSCOCO has 113,287 images for training, 5K images for validation, and both 5K and 1K for testing. Flickr30K has 28,783 images for training, 1K images for validation, and 1K for testing. Following previous works, we use recall $R@k$ ($k=1,5,10$) as the main metric.

We use the open-source CLIP (Radford et al., 2021) with ViT-B/32 (Dosovitskiy et al., 2020) as the teacher model. Its image encoder is a 12-layer ViT with the hidden size to be 768 and 12 attention heads. Its text encoder is a 12-layer Transformer with hidden size to be 512 and 8 attention heads.

For the student model, we use ViT-S/16 with hidden size to be 384 as the image encoder, and initialize it from the pre-trained weights on ImageNet-21K (Ridnik et al., 2021). For the text encoder, we experiment with 2, 4 and 6-layer Transformer, of which the weights are initialized from the first corresponding layers of the teacher's text encoder. The details of model settings are shown in Tab. 6.

In pre-training distillation, we train the student models in 4 epochs using AdamW (Loshchilov and Hutter, 2018) with a batch size of 1024 for both images and texts, the learning rate of $3e-4$, and the weight decay of 0.1. We employ a cosine learning rate scheduler with 10,000 warm-up steps. In fine-tuning, we use the same optimization setting as in MoTIS (Ren and Zhu, 2022). Experiments are conducted on 4 NVIDIA TESLA V100 32G GPUs.

4.2 Ablation Study

Considering our complete pre-training distillation takes a relatively long time, we follow the setup of (Ren and Zhu, 2022) and train ConaCLIP on CC3M for 1 epoch with batch size 84 to conduct the ablation study. Taking Eq. (4) as the naive baseline, we aim to find out which of the proposed

Learning Type	Supervision Strategies					
	InfoNCE	FD	SD	KL-Div	Sym-SD	Sym-KL-Div
intra-modal stu-stu learning	\	\	58.8/83.7/90.1	57.1/82.7/88.8	57.1/82.0/89.2	56.8/81.6/88.6
inter-modal stu-stu learning	34.7/58.7/69.9	56.6/82.1/88.8	58.6/83.6/90.0	56.5/82.4/88.9	\	\
intra-modal tch-stu learning	57.6/82.4/89.0 [†]	57.6/82.0/88.4	58.5/83.2/89.6	55.1/80.0/87.4	58.7/83.4/89.9	56.3/81.5/88.3
inter-modal tch-stu learning	51.4/76.3/83.8	50.0/80.7/88.4	57.6/82.5/88.6	56.9/81.8/88.7	56.9/81.8/88.7	59.1/83.4/89.8

Table 2: Ablation study of text-image retrieval R@1/5/10 on Flickr30K. [†]Baseline. **Bold** denotes all R@ks have obvious improvements. All five losses **in bold** will be added to the baseline loss to finally serve as our framework.

Model	Text Encoder	Image Encoder	Flickr30K			MSCOCO (1K)			MSCOCO (5K)		
			R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
(a) Fair Comparisons											
InfoNCE-based	CLIP's[512/6]	ViT-S/16[384/12]	38.4	68.0	78.0	53.3	85.3	93.5	31.5	60.3	73.3
Cross-modal KD			41.1	70.6	80.0	54.9	86.0	93.6	33.4	61.9	74.4
MoTIS			57.0	82.1	88.8	62.7	88.2	94.5	42.6	69.6	79.4
ConaCLIP (Ours)			60.6	85.2	91.2	68.6	92.4	96.7	47.3	76.1	85.2
(b) Model Zoo and Benchmarks											
ConaCLIP-6L (Ours)	CLIP's[512/6]	ViT-S/16[384/12]	67.6	89.6	94.4	75.6	94.6	97.4	55.4	83.5	89.9
ConaCLIP-4L (Ours)	CLIP's[512/4]		67.0	89.3	94.2	75.4	94.6	97.4	55.3	83.1	89.9
ConaCLIP-2L (Ours)	CLIP's[512/2]		65.6	89.2	93.9	74.7	94.3	97.3	54.1	82.2	89.4

Table 3: (a) Fair comparisons of text-image retrieval results on Flickr30K and MSCOCO (1K and 5K). (b) Our model zoo and the corresponding benchmarks. **Bold** indicates the best performance. "[m/n]" represents n layers with the hidden size to be m.

combinations of learning types and supervision strategies can bring further improvements. The fine-tuned results on Flickr30K is shown in Tab. 2.

We can make some observations that: 1) With an appropriate choice of detailed supervision strategies, each type of learning can further bring obvious improvements on the basis of the baseline. 2) The effect of each learning type is greatly affected by the implemented loss function. It also indicates that the pre-training distillation process should be carefully explored regarding the supervision strategy. 3) Our proposed symmetric version losses (Sym-SD and Sym-KL-Div) can generally achieve superior performances to the standard ones for (intra/inter-modal) teacher-student learning.

We can also attain several findings that: 1) For (intra/inter-modal) student-student learning where students first make knowledge interaction and then learn together from teachers, SD loss performs the best. Because the actual retrieval application uses this cosine similarity to rank candidates, it can help students acquire goal-oriented knowledge more directly. It also relaxes the learning task of students from teachers' feature space to the similarity space. 2) For (intra/inter-modal) teacher-student learning, our proposed symmetric version losses are more

suitable. Compared with the standard losses, they make the knowledge interaction between teachers and students closer during optimization. In this regard, student encoders can cooperate more intimately in downstream tasks. 3) Although the naive intra-modal teacher-student learning with InfoNCE loss can serve as a competent baseline, the addition of SD and Sym-SD losses of the same learning type can complement its effectiveness. On the other hand, the other three different learning types with proper loss choices can also benefit the effect of pre-training distillation. More findings on distilling intermediate layers are shown in A.2.

Our method has been established with the further integration of the highlight (in bold) combinations in Tab. 2 based on the baseline. The effect after full integration is shown in Tab. 3(a).

4.3 Performance

Fair Comparisons. In order to better verify the effectiveness of ConaCLIP, besides the previous SOTA, we also experiment with two strong baseline methods. As shown in Tab. 3(a), *InfoNCE-based* indicates the naive cross-modal contrastive learning procedure. *Cross-modal KD* represents distilling the cross-modal in-batch probability distribution of teachers into students. All these experiments

Model	Text-Image Retrieval			Image-Text Retrieval			Disk Space (MB)	QPS _t	QPS _i
	R@1	R@5	R@10	R@1	R@5	R@10			
CLIP	16.5	48.0	61.3	18.0	49.7	62.2	578	1.00×	1.00×
EC-CLIP	25.0	63.5	76.0	25.9	64.1	75.7	578	1.00×	1.00×
EC-ConaCLIP-6L	24.3	62.4	75.6	24.8	62.4	73.9	254	1.92×	1.44×
EC-ConaCLIP-4L	23.6	61.1	74.3	22.0	59.7	72.2	230	2.71×	1.44×
EC-ConaCLIP-2L	23.0	60.7	73.5	21.8	59.3	72.0	206	4.86×	1.44×

Table 4: Performance of the industry application. "EC-" is the e-commercial version of our model. QPS_t/QPS_i indicates the acceleration rate of QPS.

are conducted under the pre-training setup of (Ren and Zhu, 2022) for fair comparisons. As can be observed, 1) Cross-modal KD which introduces the knowledge distillation process obviously outperforms the standard InfoNCE-based approach. 2) MoTIS greatly surpasses InfoNCE-based and Cross-modal KD. This reveals the superiority of intra-modal teacher-student learning over inter-modal student-student learning in the case of dual-encoder distillation. 3) Our ConaCLIP shows significant improvements compared with competitors on all evaluation metrics: 3.6/3.1/2.4 R@1/5/10 gains on Flickr30K, 5.9/4.2/2.2 R@1/5/10 gains on MSCOCO (1K) and 4.7/6.5/5.8 R@1/5/10 gains on MSCOCO (5K). This fully demonstrates the effectiveness of our distillation framework with Cona.

Model Zoo and Benchmarks. In order to better promote the development of cross-modal text-image research, we release a series of lightweight dual-encoder models. Their benchmark results are shown in Tab. 3(b). In this case, the power of ConaCLIP is further unlocked and brings further improvements. Specifically, even ConaCLIP-2L can achieve 8.6/7.1/5.1 R@1/5/10 gains on Flickr30K, 12.0/6.1/2.8 R@1/5/10 gains on MSCOCO (1K) and 11.5/12.6/10.0 R@1/5/10 gains on MSCOCO (5K) compared with the previous SOTA. We have also found that the capacity of the text encoder may have limited effects on these performances. For example, ConaCLIP-4L can achieve competitive results with ConaCLIP-6L, and ConaCLIP-2L has only minor drops.

5 Industry Application

We apply the proposed technique to end-to-end cross-modal retrieval in an e-commerce platform, where we vectorize the search queries and the products and then perform product retrieval and ranking with nearest-neighbor search (Muja and Lowe,

2009; Jegou et al., 2010; Johnson et al., 2019), as shown in Fig. 2. We first collect massive data of text-image pairs from e-commerce products in our platform, where the titles of products can act as text information. We utilize most of the data to pre-train an e-commerce version of the CLIP model (denoted as EC-CLIP) with ViT-B/32 as the image encoder, which is overly large for online deployment. For the remaining data, we utilize 3M pairs for distilling the lightweight EC-ConaCLIP. To evaluate its effectiveness, we hold out a separate set of 100K pairs for fine-tuning and 5K/5K pairs used in validating/testing. In this set of experiments, we train EC-ConaCLIP for 20 epochs in pre-training distillation, and fine-tune both EC-CLIP and EC-ConaCLIP for 5 epochs. The remaining settings are the same as in Section 4.1.

In apart to the R@k metric, we also report the disk space (MB) and the acceleration rate of Query Per Second (QPS_i for image and QPS_t for text) to evaluate model’s memory footprints and inference speed. In Tab. 4, we report the averaged results where the inference speed is tested on an NVIDIA TESLA V100 (16G) GPU. As seen, the compressed EC-ConaCLIP-6L only takes 44% disk space (254MB) of EC-CLIP meanwhile being 1.44×/1.92× faster with image/text queries. It also performs on par with EC-CLIP. Our EC-ConaCLIP-2L can further achieve up to 4.86× inference speed-up with text queries, and 64% size reduction (from 578MB to 206MB). We provide some case studies in A.4.

6 Conclusion

In this paper, we propose Cona for pre-training distillation with dual-encoder architecture. It gathers every type of knowledge interaction learning with appropriate supervision choice to benefit the cross-modal distillation. The resulting ConaCLIP

achieves superior performances on both general benchmarks and industry applications.

For future work, we will explore more variants of visual encoders, and continue to tap the potential of dual-encoder distillation.

Acknowledgements

This research is supported in part by Alibaba Innovative Research Foundation (No. D8200510), NSFC (Grant No. 61936003), Zhuhai Industry Core and Key Technology Research Project (No. 2220004002350). This work is also supported by Alibaba Cloud Group, through Research Talent Program with South China University of Technology.

References

- Armen Aghajanyan, Ankit Gupta, Akshat Shrivastava, Xilun Chen, Luke Zettlemoyer, and Sonal Gupta. 2021. Muppet: Massive multi-task representations with pre-finetuning. In *EMNLP*, pages 5799–5811.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *NeurIPS*, 33:1877–1901.
- Rich Caruana. 1997. Multitask learning. *Machine Learning*, 28(1):41–75.
- Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. 2021. Conceptual 12M: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *CVPR*.
- Feilong Chen, Xiuyi Chen, Shuang Xu, and Bo Xu. 2022. Improving cross-modal understanding in visual dialog via contrastive learning. In *ICASSP*, pages 7937–7941.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. UNITER: Universal image-text representation learning. In *ECCV*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional Transformers for language understanding. In *NAACL-HLT*, pages 4171–4186.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*.
- Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Pengchuan Zhang, Lu Yuan, Nanyun Peng, et al. 2022. An empirical study of training end-to-end vision-and-language transformers. In *CVPR*, pages 18166–18176.
- Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2014. Distilling the knowledge in a neural network. *NeurIPS Deep Learning Workshop*.
- Herve Jegou, Matthijs Douze, and Cordelia Schmid. 2010. Product quantization for nearest neighbor search. *IEEE TPAMI*, 33(1):117–128.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, pages 4904–4916.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, pages 12888–12900.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. VisualBERT: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. 2020. Oscar: Object-semantic aligned pre-training for vision-language tasks. In *ECCV*, pages 121–137.
- Xuwei Li, Songyuan Li, Bourahla Omar, Fei Wu, and Xi Li. 2021. ResKD: Residual-guided knowledge distillation. *IEEE TIP*, 30:4735–4746.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *ECCV*, pages 740–755.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2018. Decoupled weight decay regularization. In *ICLR*.
- Minh-Thang Luong, Quoc V Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2016. Multi-task sequence to sequence learning. *ICLR*.
- Marius Muja and David G Lowe. 2009. Fast approximate nearest neighbors with automatic algorithm configuration. *VISAPP*, 2(331-340):2.

- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, pages 2641–2649.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763.
- Siyu Ren and Kenny Zhu. 2022. Leaner and faster: Two-stage model compression for lightweight text-image retrieval. In *NAACL-HLT*, pages 4085–4090.
- Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. 2021. ImageNet-21K pretraining for the masses. In *NeurIPS Datasets and Benchmarks Track*.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, pages 2556–2565.
- Hao Tan and Mohit Bansal. 2019. LXMERT: Learning cross-modality encoder representations from Transformers. In *EMNLP*, pages 5100–5111.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *NeurIPS*, 30.
- Chengyu Wang, Minghui Qiu, Taolin Zhang, Tingting Liu, Lei Li, Jianing Wang, Ming Wang, Jun Huang, and Wei Lin. 2022a. Easynlp: A comprehensive and easy-to-use toolkit for natural language processing. In *EMNLP*, pages 22–29.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022b. OFA: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *ICML*, pages 23318–23340.
- Xiaodan Wang, Lei Li, Zhixu Li, Xuwu Wang, Xiangru Zhu, Chengyu Wang, Jun Huang, and Yanghua Xiao. 2023. AGREE: aligning cross-modal entities for image-text retrieval upon vision-language pre-trained models. In *WSDM*, pages 456–464.
- Quanzheng Xu, Liyu Liu, and Bing Ji. 2022. Knowledge distillation guided by multiple homogeneous teachers. *Information Sciences*, 607:230–243.
- Shengyu Zhang, Tan Jiang, Tan Wang, Kun Kuang, Zhou Zhao, Jianke Zhu, Jin Yu, Hongxia Yang, and Fei Wu. 2020. DeVLBERT: Learning deconfounded visio-linguistic representations. In *ACMMM*, pages 4373–4382.
- Yichen Zhu and Yi Wang. 2021. Student customized knowledge distillation: Bridging the gap between student and teacher. In *ICCV*, pages 5057–5066.

A Appendix

A.1 Model Settings

We give detailed parameters on the settings of our ConaCLIP models in Tab. 6, such as the number of parameters, layers, heads, etc.

Model Setting	ConaCLIP-6L	ConaCLIP-4L	ConaCLIP-2L
Number of Parameters	66M	60M	53M
Text Encoder Layers	6	4	2
Text Encoder Heads	8	8	8
Text Encoder Hidden Size	512	512	512
Vocabulary Size	49408	49408	49408
Text Length	77	77	77
Image Encoder Layers	12	12	12
Image Encoder Heads	6	6	6
Image Encoder Hidden Size	384	384	384
Image Patch Size	16	16	16
Image Size	224	224	224

Table 6: Detailed parameters on the settings of our ConaCLIP models.

A.2 Negative Results on Distilling Intermediate Layers

We also present an exploratory study on distilling the knowledge of intermediate layers from teacher encoders. We first evenly divide the encoder of each student/teacher into six parts along the number of layers, and then perform our distillation technique on the feature representations of each part. The experiment results are shown in Tab. 5.

We can observe that additional distillation with features of the intermediate layers does not bring about positive improvement. This inspires us that we should mainly focus on the representation matching ability of the output of the last layer for the cross-modal retrieval task. Due to the difference of capabilities between models of different sizes, they can choose different paths to learn the goal-oriented features in the same task during distillation (Li et al., 2021; Zhu and Wang, 2021; Xu

Applied Parts	R@1	R@5	R@10
6th (Baseline)	60.6	85.2	91.2
5-6	59.5	84.2	90.7
4-6	59.5	84.3	90.9
3-6	57.9	83.5	90.2
2-6	58.6	84.1	90.9
1-6	59.2	84.5	90.8

Table 5: An exploratory study on distilling intermediate layers. The R@1/5/10 results on Flickr30K are listed. Each student/teacher encoder is evenly divided into six parts along the number of layers, and distillation is performed on the feature representations of each part.

et al., 2022). In our application, we suggest that it can be inappropriate to force small models to learn the same path as the large ones.

A.3 Application in E-Commerce Product Retrieval

We apply the proposed distillation technique to end-to-end cross-modal retrieval in an e-commerce platform, where we vectorize the search queries and the products and then perform product retrieval and ranking with nearest-neighbor search. The whole framework is shown in Fig. 2.

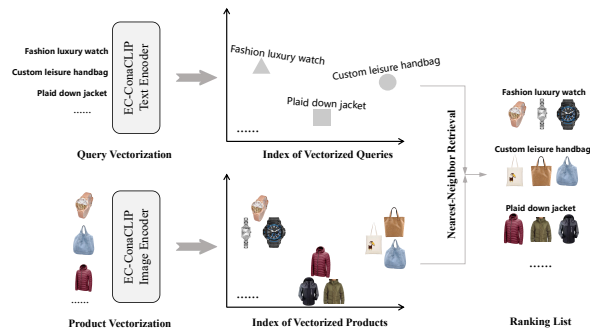


Figure 2: The application of our ConaCLIP in e-commerce retrieval.

A.4 Case Study

Case	Query	CLIP	EC-ConaCLIP (Ours)
1	Waterproof large capacity lightweight fashion unicorn cartoon kids girl middle school backpack.		
2	Stainless steel induction steamers pot, 2 layers double handle food cooking pots with lid.		
3	Children's sand hammer wooden bell multi-color children's development toy.		
4	Hot red large size sports tights high waist yoga pants.		
5	Tempered glass waterproof platform 5kg digital food electric kitchen scale.		

Table 7: Case studies in e-commerce retrieval. Given the same text query, we show the image retrieval results of the open-source CLIP and our EC-ConaCLIP.

Tab. 7 shows the case studies in our e-commerce retrieval scenario. For the same text query, we

show the top-1 image retrieval results of the open-source CLIP model and our EC-ConaCLIP model respectively.

From these cases, we can find that our model can better capture conceptual and fine-grained fashion information during cross-modal text-image retrieval, and maintain the cross-modal alignment effect of text-image samples after the lightweight distillation. For example, in Case 1, our model more accurately captures the cartoon subject in the target commodity as "unicorn". In Case 2, our model pays more attention to fine-grained information "2 layers double handle", while maintaining the correct perception of other information such as "Stainless steel", "steamers pot" and "with lid". In Case 3, our EC-ConaCLIP better captures the color clue of "Hot red". Although the retrieval result of CLIP also conforms to the information of "sports tights high waist yoga pants", its color is more like "dark red".

Based on our distillation technique, the resulting model can sufficiently learn the perception ability of the teacher model about commodity fashion concepts and reduce matching errors.