

# TinyCLIP: CLIP Distillation via Affinity Mimicking and Weight Inheritance

Kan Wu<sup>1,\*</sup>, Houwen Peng<sup>3,\*</sup>, Zhenghong Zhou<sup>2,\*</sup>, Bin Xiao<sup>3</sup>, Mengchen Liu<sup>3</sup>, Lu Yuan<sup>3</sup>,  
Hong Xuan<sup>3</sup>, Michael Valenzuela<sup>3</sup>, Xi (Stephen) Chen<sup>3</sup>, Xinggang Wang<sup>2</sup>, Hongyang Chao<sup>1</sup>, Han Hu<sup>3</sup>  
<sup>1</sup> Sun Yat-sen University, <sup>2</sup> Huazhong University of Science & Technology, <sup>3</sup> Microsoft

## Abstract

In this paper, we propose a novel cross-modal distillation method, called TinyCLIP, for large-scale language-image pre-trained models. The method introduces two core techniques: affinity mimicking and weight inheritance. Affinity mimicking explores the interaction between modalities during distillation, enabling student models to mimic teachers' behavior of learning cross-modal feature alignment in a visual-linguistic affinity space. Weight inheritance transmits the pre-trained weights from the teacher models to their student counterparts to improve distillation efficiency. Moreover, we extend the method into a multi-stage progressive distillation to mitigate the loss of informative weights during extreme compression. Comprehensive experiments demonstrate the efficacy of TinyCLIP, showing that it can reduce the size of the pre-trained CLIP ViT-B/32 by 50%, while maintaining comparable zero-shot performance. While aiming for comparable performance, distillation with weight inheritance can speed up the training by 1.4 - 7.8 $\times$  compared to training from scratch. Moreover, our TinyCLIP ViT-8M/16, trained on YFCC-15M, achieves an impressive zero-shot top-1 accuracy of 41.1% on ImageNet, surpassing the original CLIP ViT-B/16 by 3.5% while utilizing only 8.9% parameters. Finally, we demonstrate the good transferability of TinyCLIP in various downstream tasks. Code and models will be open-sourced at [aka.ms/tinyclip](https://aka.ms/tinyclip).

## 1. Introduction

Large-scale language-image pretraining, e.g., CLIP [46], has recently gained significant attention due to its remarkable zero-shot transfer capability [46] and unprecedented performance in text-to-image generation [47]. Due to complex nature of vision and language, current approaches often resort to utilizing huge amounts of parameters to endow models with cross-modal capabilities [46, 21, 71, 72, 43, 39,




Figure 1. Comparison of the OpenCLIP [20] and TinyCLIP. TinyCLIP is pre-trained and distilled on LAION-400M [50] using OpenCLIP ViT-B/32 [20] as the teacher, whose zero-shot top-1 accuracy is 65.6%.

1]. This in turn leads to high costs in terms of storage, memory, and computation time for these models, which motivates the need for model compression to make them smaller and faster for real-world applications [64, 13].

As a core compression technique, knowledge distillation has been extensively studied and applied in single-modal settings [19, 16]. However, its potential for multi-modality remains underexplored. Unlike single-modal models, the distillation of language-image cross-modal models poses distinct challenges. First, CLIP-like language-image models commonly consist of two branches: an image encoder and a text encoder [21, 72, 52, 73]. When distilling such multi-branch models, it is crucial to consider the interaction of information across the different modality branches in both the teacher and student models. Second, the original CLIP [46] models are pre-trained on 400 million image-text pairs for 32 epochs, which requires thousands of GPU days, making distillation a significant challenge when computational resources are limited. Is there any way to reduce the cost of CLIP distillation?

To tackle these challenges, we present a novel cross-modal distillation method dubbed TinyCLIP, which introduces two key techniques: *affinity mimicking* and *weight inheritance*. In contrast to the methods that rely on either

\* Equal contribution. Kan and Zhenghong were interns of Microsoft.

† Corresponding: houwen.peng@microsoft.com

image or text features for distillation, we empirically show that distilling knowledge in an image-text affinity space is more effective. Specifically, we leverage the cosine similarity of the image and text embeddings in the teacher model to facilitate the distillation of the student model, allowing the student to mimic the teacher’s visual-linguistic feature alignment. We refer to this process as *affinity mimicking*.

To improve distillation efficiency, we introduce *weight inheritance*, a technique transferring the pre-trained weights from teacher models to their student counterparts. Since inheriting weights provides a good initialization for the student models, the distillation progress can be largely accelerated. The key challenge of weight inheritance lies in determining which weights are more advantageous. To address this issue, we introduce two solutions: manual and automatic inheritance. We surprisingly found that a simple manual selection of  $k$ -dimension or  $k$ -layer weights from the teacher model can yield satisfactory results for CLIP distillation. On the other hand, we also introduce learnable masks to automatically identify the most important weights from the teacher model. The masks are imposed independently on the vision and language branches, enabling them to capture the differences across modalities.

Moreover, we extend the proposed weight inheritance to a multi-stage progressive procedure, where each subsequent stage automatically inherits the important weights from the preceding stages. We observed that when the teacher model exhibits higher performance and shares a similar architecture with the student, weight inheritance can provide better results. This is because significant architecture differences may undermine the learned weights when transmitting from the teacher to the student. Therefore, we break the inheritance into multiple stages, allowing the student model in each stage to share a more similar structure with the predecessor teacher and inherit the weights progressively.

Our experiments show that TinyCLIP delivers competitive models at all levels of speedups and model sizes in ImageNet zero-shot evaluation and various downstream tasks. As shown in Fig. 1, pre-trained on the same LAION-400M dataset [50], TinyCLIP-ViT using 63M parameters achieves 61.4% zero-shot top-1 accuracy on ImageNet [7], being  $2.0\times$  smaller and  $1.5\times$  faster than the OpenCLIP [20] model (62.9% accuracy with 126M parameters). Meanwhile, TinyCLIP-ResNet with 38M parameters obtains 56.4% zero-shot top-1 accuracy, being  $2.1\times$  smaller and  $2.0\times$  faster than CLIP ResNet-50 (59.6% accuracy with 76M parameters). Moreover, our method can speed up the original OpenCLIP training by  $1.4\times - 7.8\times$  while achieving similar performance. Also, TinyCLIP demonstrates good transfer capacities in downstream scenarios. In summary, the contributions of this work are two-fold:

- We propose a new cross-modal distillation approach to unleash the capacity of small CLIP models, fully lever-

aging large-scale models as well as pre-training data. To our best knowledge, this is the first work exploring the pre-training distillation of language-image models.

- We present state-of-the-art language-image pre-trained models at small scale, striking the best trade-off between speed and accuracy. Extensive experiments demonstrate the superiority and good generalization ability of the small models in various downstream tasks.

## 2. Related Work

**Language-Image Pre-training** has achieved remarkable progress over the past few years [10, 4]. In particular, contrastive language-image pre-training demonstrates very impressive zero-shot transfer and generalization capacities [45, 46, 21, 72, 61, 43, 74]. One of the most representative works is CLIP [46]. A large amount of follow-up works have been proposed to optimize the pre-training framework [27, 76, 15, 68, 62, 14, 28, 67]. Meanwhile, a line of works leverage the pre-trained models for downstream tasks, such as open-vocabulary detection and segmentation [29, 75, 66, 32], video recognition [56, 23], and text-to-image generation [47, 41].

Most recently, there are a few works attempting to scale language-image pre-trained models. BEiT-3 [57] uses a multiway transformer to scales up multi-modal pre-trained models to 1.9B parameters. FLIP [30] scales image-text pre-training via masked modeling, enabling CLIP to enjoy faster training speed while getting better performance. On the other hand, there are also few works on scaling down CLIP models [51]. However, they all focus on specific tasks, such as generation [5, 51] and video content understanding [42]. In contrast, our work concentrates on language-image pre-training, which is the first work on CLIP compression.

**Knowledge Distillation** in a teacher-student framework [19] has been widely used to transfer knowledge from large models to small ones. Distillation in single modality, such as vision [54, 64, 37] or language [48, 22], has been extensively studied. However, the exploration of cross-modality distillation is relatively limited. A few early works exploring cross-modal distillation only focus on specific tasks, such as VQA and image caption [5, 11, 59], limiting the methodologies to be applied in general downstream tasks. In contrast, this work pays attention to general cross-modal pre-training distillation, in which the distilled models can be transferred to various downstream scenarios.

It is worth noting that our proposed weight inheritance strategy shares similar spirits to model pruning techniques [17, 65, 51]. Both approaches involve the identification of important weights while removing redundant ones. However, they have three fundamental differences. 1) Previous pruning methods predominantly focus on single-modality

models [17, 65, 3], whereas our approach shifts attention to multi-modality. Since different modalities exhibit varying redundancy characteristics, it is necessary to consider these differences when selecting important weights from each modality. 2) Unlike previous pruning methods which have emphasized the importance of the pruned architecture [12, 34, 58], we argue that the weights inherited from the original models are also highly beneficial, particularly when compressing CLIP-like language-image models. 3) Our method leverages a progressive multi-stage process, which enables the inherited weights to accelerate the convergence of small models during training.

### 3. Method

In this section, we propose TinyCLIP, a simple and effective method for distilling large-scale language-image pre-training models, such as CLIP. It consists of three components: affinity mimicking (§3.1), weight inheritance (§3.2) and multi-stage progressive distillation (§3.3).

#### 3.1. Distillation with Affinity Mimicking

A language-image pre-training model often consists of two branches, an image encoder to extract visual representations and a text encoder to extract textual representations. The visual and textual representations are linearly projected into a cross-modal common space by minimizing a contrastive loss  $\mathcal{L}_0$ . The supervision signal can be viewed as an identity matrix, in which diagonal values are set to 1 for positive image-text pairs, and all others are set to 0 for negative pairs. This signal disregards the similarity between negative pairs, which prevents the model from acquiring a nuanced understanding of the intricate relationships among negative pairs. Therefore, we introduce *affinity mimicking* to enable student models to learn similarities between negative pairs from teacher models.

More specifically, as shown in Fig. 2, we consider two types of affinity distillation losses: image-to-language loss  $\mathcal{L}_{I2T}$  and language-to-image loss  $\mathcal{L}_{T2I}$ . The former learns the alignment between the teacher and student models based on the image-to-language affinity  $A_{I2T}$ , which represents the affinity scores of an image with all the text descriptions in a batch (the yellow row highlighted in Fig. 2). The latter refers to the alignment of language-to-image affinity  $A_{T2I}$ , which compares a text description with all the images to obtain a matching score (the blue column in Fig. 2). The combination of these two losses formulates our affinity mimicking distillation, which is represented as:

$$\begin{aligned} \mathcal{L}_{distill} &= \mathcal{L}_{I2T} + \mathcal{L}_{T2I}, \\ &= CE(A_{I2T}^s, A_{I2T}^t) + CE(A_{T2I}^s, A_{T2I}^t). \end{aligned} \quad (1)$$

Here,  $CE$  represents the cross-entropy loss. The superscripts  $s$  and  $t$  denote the student and teacher models, re-




Figure 2. Affinity mimicking for language-image models. The loss includes image-to-text loss (yellow) and text-to-image loss (blue).

spectively. The elements in  $A_{I2T}$  and  $A_{T2I}$  are defined as:

$$A_{I2T}(i, j) = \frac{\exp(I_i \cdot T_j / \tau)}{\sum_{k \in \mathcal{B}} \exp(I_i \cdot T_k / \tau)}, \quad (2)$$

$$A_{T2I}(i, j) = \frac{\exp(I_i \cdot T_j / \tau)}{\sum_{k \in \mathcal{B}} \exp(I_k \cdot T_j / \tau)}, \quad (3)$$

where  $\tau$  is a temperature parameter,  $I_i$  is the feature embedding of  $i$ -th image in the batch  $\mathcal{B}$ ,  $T_j$  is the feature of  $j$ -th text description in  $\mathcal{B}$ . Affinity mimicking allows student models to imitate the behavior of the large model in learning visual-linguistic alignment. In contrast to previous methods that rely on either image or text features for distillation, we show that distilling knowledge in this image-text affinity space is more effective.

#### 3.2. Distillation with Weight Inheritance

The original CLIP models are pre-trained on 400 million image-text pairs for 32 epochs, taking thousands of GPU days. This presents a significant cost challenge for distillation. To improve the training efficiency, we introduce weight inheritance, a technique that inherits the important weights from the well-trained large teacher models to smaller student models. The key challenge of inheriting weights is identifying important weights from the large amount of weights of the teacher. We propose two approaches to select the important weights: manual weight inheritance and automatic weight inheritance.

**Manual Inheritance.** For manual weight inheritance, we first analyze the redundancy of existing CLIP pre-trained models. Fig. 4 shows that the text encoder displays more redundancy in depth (layer-wise), while the image encoder exhibits more redundancy in width (channel-wise). Based upon this discovery, we adopt the approach of uniformly selecting  $k$  layers of the text branch and directly taking the

front  $k$  channels of the image branch to select the important weights from large teacher models. These weights will serve as the initialization for small models during distillation. We surprisingly found that such a simple manual selected weights can largely accelerate CLIP distillation.

**Automatic Inheritance.** Although manual inheritance can be considered an effective approach for model compression, it does present the drawback of necessitating prior knowledge in order to determine which weights to inherit. This limitation may render the technique impractical for application to diverse models. To solve this issue, we present an automatic weight inheritance scheme. Inspired by structured pruning in large language models [65], we introduce learnable masks to identify weight importance. Considering the difference across modalities, the learnable masks are imposed on vision and language branches independently, as visualized in Fig. 3. An overall sparsity constraint is introduced to guarantee the selected number of important weights to meet our compression requirements. Without loss of generality, here we use the transformer architecture as an example to introduce the detailed procedure of weight inheritance. Specifically, a standard transformer block contains a multi-head attention (MHA) layer and a feed-forward network (FFN). To capture the importance of weights in a fine-grained level, we introduce two mask variables  $\mathbf{m}_{\text{head}}, \mathbf{m}_{\text{int}} \in \{0, 1\}$  to identify the redundant attention heads in MHA and neurons in FFN respectively, while keeping the important ones. These two kinds of masks are imposed on the activation of attention heads and the intermediate layer of FFN, which is formulated as

$$\text{MHA}(\mathbf{X}) = \sum_{h=1}^{N_H} \mathbf{m}_{\text{head}}^h \cdot \text{Attn}_{\mathbf{W}_Q^h, \mathbf{W}_K^h, \mathbf{W}_V^h, \mathbf{W}_O^h}(\mathbf{X}), \quad (4)$$

$$\text{FFN}(\mathbf{X}) = \text{GeLU}(\mathbf{X}\mathbf{W}_U) \cdot \text{diag}(\mathbf{m}_{\text{int}}) \cdot \mathbf{W}_D, \quad (5)$$

where  $\mathbf{X}$  is the layer input,  $\mathbf{W}_Q^h, \mathbf{W}_K^h, \mathbf{W}_V^h, \mathbf{W}_O^h \in \mathbb{R}^{d \times d_h}$  denote the query, key, value and output matrices in MHA respectively, and  $\mathbf{W}_U \in \mathbb{R}^{d \times d_f}$  and  $\mathbf{W}_D \in \mathbb{R}^{d_f \times d}$  represent the parameters of FFN layers. Here,  $d$  denotes the hidden dimension (e.g., 768), and  $d_h = d/N_H$  denotes the output dimension of each head (e.g., 64), where  $N_H$  is the total number of heads. The trainable masks  $\mathbf{m}_{\text{head}}$  and  $\mathbf{m}_{\text{int}}$  serve as importance indicators for MHA and FFN layers. Moreover, to further learn the importance of embedding dimensions in transformer, we introduce an additional mask  $\mathbf{m}_{\text{embed}} \in \{0, 1\}$ . This mask is shared across all layers because each dimension in the hidden representation is connected to the corresponding dimension in the subsequent layer through a residual connection.

During automatic inheritance, there are two losses used for optimizing the masks and the model jointly, i.e., a sparsity loss and a distillation loss defined in Eq. (1):

$$\mathcal{L} = \mathcal{L}_{\text{distill}} + \mathcal{L}_{\text{sparsity}}, \quad (6)$$

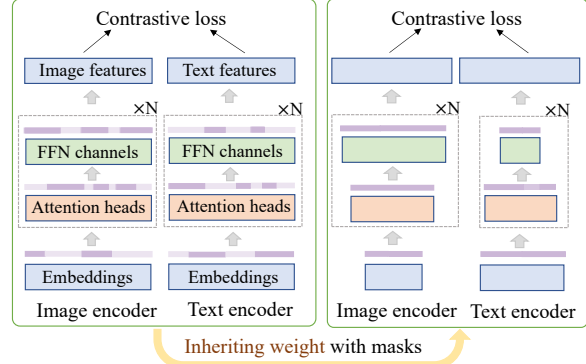


Figure 3. Weight inheritance of CLIP. Given a large pre-trained CLIP model, the inherited weights are selected by masks. The weights whose mask is 0 will be removed. For the manual way, the masks are set to 1 in the first  $k$ -dimension or uniform  $k$ -layer, the rest is set to 0. For the automatic way, the masks are learned jointly with the model under the sparsity constraint.

$$\mathcal{L}_{\text{sparsity}} = \lambda \cdot (p - q) + \beta \cdot (p - q)^2, \quad (7)$$

where  $\lambda, \beta$  are learnable multipliers that guarantee  $p = q$  [60, 65]. Here,  $q$  is the target compression rate,  $p$  is the overall compression rate of learnable masks for the model, including image encoder and text encoder:

$$p = \frac{1}{M_i + M_t} \sum_{img, txt} (4 \cdot d_h \cdot \sum_i^L \sum_j^{N_H} \sum_k^d \mathbf{m}_{\text{head}}^{i,j} \cdot \mathbf{m}_{\text{embed}}^k + 2 \cdot \sum_i^L \sum_j^{d_f} \sum_k^d \mathbf{m}_{\text{int}}^{i,j} \cdot \mathbf{m}_{\text{embed}}^k), \quad (8)$$

where  $M_i$  and  $M_t$  represent the full model size of the image encoder and text encoder, respectively,  $L$  is the total number of transformer layers.

### 3.3. Progressive Multi-Stage Distillation

When attempting to achieve a high target sparsity, i.e.,  $>70\%$ , compressing the model in a single stage can lead to a significant reduction in accuracy and even result in convergence failure. This is due to the fact that most weights of the large model are directly discarded, including those that are important for ensuring model quality and convergence. As a solution, we propose a multi-stage progressive distillation method to achieve a high compression rate without seriously sacrificing accuracy. In each stage, we use a modest degree of compression, e.g., 25%, to avoid large loss of performance and make training stable.

Specifically, each stage includes weight inheritance and affinity mimicking, as described in Algorithm 1. During the weight inheritance phase, the weights of the large model are gradually reduced until the target sparsity and key weights are retained for the small model. After that, the small model is distilled with affinity mimicking, which involves transferring the visual-linguistic affinity knowledge from

---

**Algorithm 1: TinyCLIP: Cross-Modal Distillation**

---

**Input:** a pre-trained CLIP model  $f(x; \theta_0)$  with weight  $\theta_0$ , target compression rate  $q$ , total iteration steps  $N$

```
1 Initialize the number of stages  $G$ , each stage steps  $L$ , target
  compression rate  $q_i$  of stage  $i$ , weight inheritance steps  $L_M$ ,
  input  $x$ , learnable mask  $M = \{\mathbf{m}_{\text{head}}, \mathbf{m}_{\text{int}}, \mathbf{m}_{\text{embed}}\}$ ;1
2 for  $i$  in  $[1, \dots, G]$  do
3   // Increase mask compression rate, retain important weights
4   for  $j$  in  $[1, \dots, L_M]$  do
5     Update target compression rate2  $q$  of current step based
      on  $q_i$  and step  $j$ ;
6     Calculate distillation loss between origin teacher
      model  $f(x; \theta_0)$  and masked student model
       $f(x; \theta_{i-1} \odot M_i)$ ;
7     Calculate sparsity loss between learnable mask
      compression rate3  $p$  and target compression rate  $q$ ;
8     Optimize  $\theta_{i-1}$  and  $M_i$ ;
9   end
10   $\theta_i \leftarrow \theta_{i-1} \odot M_i$  // Inherit weight with masks
11  // Cross-Modal distillation
12  for  $j$  in  $[L_M + 1, \dots, L]$  do
13    Calculate distillation loss between origin teacher
      model  $f(x; \theta_0)$  and compressed model  $f(x; \theta_i)$ ;
14    Optimize  $\theta_i$ ;
15  end
16 end
17 Return compressed model  $f(x; \theta_G)$  at compression rate  $q$ 
```

<sup>1</sup> Each stage compress 25%,  $q_i = 1 - i \cdot 25\%$ ,  $G = \lceil (1 - q)/0.25 \rceil$ ,  $L = N/G$ ,  $L_M = 3000$ ,  $M$  is initialized to 1,  $q_i$  can be adjusted according to different models.  
<sup>2</sup> Target compression rate  $q$  increases linearly from 0 to  $q_i$  with step  $j$ .  
<sup>3</sup>  $p$  is calculated by masks  $M$  in Eq. (8).

---

the teacher to the student. This procedure is repeated until the target sparsity level is achieved.

## 4. Experiments

### 4.1. Implementation Details

*Architecture.* The original CLIP employs a transformer-based model as its text encoder, while its image encoder has two variations, *i.e.*, ViT [9] and modified ResNet [18], covering both transformer-based and CNN-based architectures. For weight inheritance, we inherited the weights from two model variants: *OpenCLIP* ViT-B/32 or ViT-B/16 [20] and *CLIP* ResNet-50 [46], which are pre-trained on LAION-2B [49] and WIT-400M [46], respectively<sup>1</sup>. For distillation, we use *OpenCLIP* ViT-B/32 [20] pre-trained on LAION-2B [49] as the teacher model, since it achieves a high zero-shot accuracy of 65.6% on ImageNet [7] with a high inference throughput. Besides, we also extend our method to DaViT [8], which is a hybrid architecture. We re-produce Florence-DaViT-5M [72, 8] on LAION-400M as a baseline. Then we distill the same model using affinity mimicking and weight inheritance. The teacher model is Florence-DaViT-D3, which is pre-trained on FLD-900M [72] and achieves a top-1 accuracy of 78.0% on ImageNet [7] with 128M parameters.

*Affinity Mimicking.* As defined in Tab. 2, we explore

<sup>1</sup>The models can be downloaded from [OpenCLIP](#) and [OpenAI](#).

different interaction schemes between teacher and student models across modalities. The contrastive loss  $\mathcal{L}_0$  and affinity mimicking  $\mathcal{L}_1$  (*i.e.*,  $\mathcal{L}_{\text{distill}}$ ) serve as the basic ones, which can be combined with other losses. The weight of each loss is set to 1 when combined. The temperature parameter  $\tau$  in Eq. (3) is set to  $\frac{1}{50}$  by default. The sensitivity analysis of  $\tau$  is presented in the *supplementary material*.

*Weight Inheritance.* We compress OpenCLIP ViT-B [20] and CLIP ResNet-50 [46] using the proposed multi-stage progressive distillation. In each stage, we apply manual or automatic weight inheritance. In manual inheritance, we compress the image encoder and the text encoder separately, by reducing the width of image encoder and the depth of text encoder. In automatic inheritance, we set the target sparsity for the whole model. In each stage, the trainable masks are initialized to 1 and updated in the first 3,000 training iterations, where the optimizer is AdamW [35, 24] with a constant learning rate of 0.01 and no weight decay. The learnable multipliers  $\lambda, \beta$  defined in Eq. (7) are both initialized to 0.01.

*Training Settings.* We train our models on two public datasets, namely LAION-400M [50] and YFCC-15M [53]. On LAION-400M, the models are compressed in 3 stages, including 100% to 75% parameters for 6 epochs, 75% to 50% parameters for 16 epochs, and 50% to 25% parameters for 16 epochs. On YFCC-15M, it contains 2 compression stages, where the training epochs are both 25 from 100% to 50% parameters, and 50% to 10%. We follow the hyper-parameter of CLIP [46] except that the learning rate is set to  $10^{-4}$  when using weight inheritance. The details are shown in the *supplementary material*. Unless otherwise specified, our ablation studies utilize the same settings, except for training the models only for 1 epoch with a learning rate of  $5 \times 10^{-4}$ . All models are trained on 32 Nvidia V100 or A100, and implemented with PyTorch [40], OpenCLIP [20], gradient cache [36], and Timm library [63].

*Evaluation Settings.* The models are evaluated on multiple benchmarks. In zero-shot transfer evaluation and robustness evaluation, we follow the same prompt engineering in CLIP [46], where 80 text templates per class are used for ImageNet [7]. In linear probe, we use Elevater toolkit [26] for evaluation, where a classification head is trained for 50 epochs with searched hyper-parameters. Note that we do not count the number of parameters in the text embedding layer. It is a look-up table whose parameter size is the same as the models with the same hidden dimension and vocabulary size. The inference throughput is measured on Nvidia V100 with CUDA 11 and PyTorch v1.12 [40], where the batch size is 1,024.

### 4.2. Comparison with State-of-the-Art Models

We compare our TinyCLIP with state-of-the-art models in Tab. 1. The zero-shot evaluation on ImageNet-1K [7],

Method	Image	Text Encoder		#Params (M)	MACs	Throughput	Training	IN-1K	Flickr30k		MSCOCO	
	Encoder	depth	width						Image+Text	(G)	(pairs/s)	datasets
<i>Training data: 15M / 20M</i>												
RILS [69]	ViT-B/16	12	512	86 + 38	20.5	818	LAION-20M	45.0	45.1	34.9	32.2	25.5
MaskCLIP [77]	ViT-B/16	12	512	86 + 38	20.5	818	LAION-20M	46.6	64.9	48.1	38.5	24.8
SLIP [38]	ViT-B/16	12	512	86 + 38	20.5	818	YFCC-15M	42.8	57.6	40.1	31.1	20.3
CLIP [46, 38]	ViT-B/16	12	512	86 + 38	20.5	818	YFCC-15M	37.6	51.6	32.2	26.5	17.1
TinyCLIP (Ours)	ViT-39M/16	6	512	39+19(2.1×)	9.5(2.2×)	1,469(1.8×)	YFCC-15M	63.5	84.4	66.7	54.9	38.9
TinyCLIP (Ours)	ViT-8M/16	3	256	8+3(11.3×)	2.0(10.3×)	4,150(5.1×)	YFCC-15M	41.1	62.3	42.3	36.2	21.5
<i>Training data: 400M</i>												
Florence [72]	DaViT-5M	6	256	5+5	1.1	2,980	LAION-400M	45.0	61.2	41.5	36.2	20.9
TinyCLIP (Ours)	DaViT-5M	6	256	5+5	1.1	2,980	LAION-400M	50.0	66.2	48.5	40.7	24.5
CLIP [46]	ResNet-101	12	512	56 + 38	12.8	1,161	WIT-400M	62.2	78.1	59.2	49.3	30.7
CLIP [46]	ResNet-50	12	512	38 + 38	9.1	1,549	WIT-400M	59.6	81.2	58.2	50.8	28.3
TinyCLIP (Ours)	ResNet-30M	9	512	30+29(1.3×)	6.9(1.3×)	1,811(1.2×)	LAION-400M	59.1	81.1	61.2	52.7	33.9
TinyCLIP (Ours)	ResNet-19M	6	512	19+19(2.0×)	4.4(2.1×)	3,024(2.0×)	LAION-400M	56.4	76.2	58.3	48.9	30.9
OpenCLIP [20]	ViT-B/32	12	512	88 + 38	7.4	2,452	LAION-2B	65.7	84.7	66.8	56.9	39.3
CLIP [46]	ViT-B/32	12	512	88 + 38	7.4	2,452	WIT-400M	63.2	80.1	59.8	51.2	30.6
OpenCLIP [20]	ViT-B/32	12	512	88 + 38	7.4	2,452	LAION-400M	62.9	79.3	62.0	53.3	35.4
TinyCLIP (Ours)	ViT-61M/32	9	512	61+29(1.4×)	5.3(1.4×)	3,191(1.3×)	LAION-400M	62.1	78.6	63.3	53.9	35.9
TinyCLIP (Ours)	ViT-40M/32	6	512	40+19(2.1×)	3.5(2.1×)	4,641(1.9×)	LAION-400M	59.7	77.3	58.9	49.8	33.1
TinyCLIP (Ours)	ViT-63M/32	auto	auto	63+31(1.3×)	5.6(1.3×)	2,905(1.2×)	LAION-400M	63.9	83.2	64.4	55.5	37.6
TinyCLIP (Ours)	ViT-45M/32	auto	auto	45+18(2.0×)	3.7(2.0×)	3,682(1.5×)	LAION-400M	61.4	80.9	62.2	52.8	34.7
TinyCLIP (Ours)	ViT-22M/32	auto	auto	22+10(3.9×)	1.9(3.9×)	5,504(2.2×)	LAION-400M	53.7	71.3	52.0	44.4	28.3
TinyCLIP (Ours)	ViT-63M/32	auto	auto	63+31(1.3×)	5.6(1.3×)	2,909(1.2×)	LAION+YFCC-400M	64.5	84.9	66.0	56.9	38.5
TinyCLIP (Ours)	ViT-45M/32	auto	auto	45+18(2.0×)	3.7(2.0×)	3,685(1.5×)	LAION+YFCC-400M	62.7	80.3	63.9	54.0	36.7

Table 1. Comparison with the state-of-the-art methods. The architecture of text encoder is Transformer [55]. “auto” denotes automatic weight inheritance, and our other models use manual weight inheritance.

Flickr30k [44] and MSCOCO retrieval [33] are reported.

When performing training on YFCC-15M [53], we use affinity mimicking and manual weight inheritance to distill OpenCLIP ViT-B/16 [20], whose zero-shot performance on ImageNet is 70.2%. By inheriting half of parameters, our TinyCLIP ViT-39M/16 can achieve a zero-shot accuracy of 63.5% on ImageNet, surpassing the original CLIP ViT-B/16 [46] by 25.9%. Furthermore, when compressing the image encoder to 10 layers with 256 dimensions and the text encoder to 3 layers with 256 dimensions, TinyCLIP ViT-8M/16 still surpasses CLIP ViT-B/16 [46] by 3.5% while using 11.3× fewer parameters and running 5.1× faster.

When conducting training on LAION-400M [50], we distill three different models, *i.e.*, Florence DaViT [72, 8], CLIP ResNet [46], and OpenCLIP ViT-B/32 [20]. The results in Tab. 1 demonstrate the efficacy of our method. More specifically, our TinyCLIP DaViT-5M achieves up to 5.0% improvements over the original model on ImageNet, and 5.0/4.5% gains in image-to-text retrieval on Flickr and MSCOCO, respectively. It is worth noting that this is a tiny language-image pre-trained model using only 10M parameters to achieve 50.0% accuracy on ImageNet.

For CLIP ResNet, TinyCLIP-30M is only trained for 4 epochs in the first stage, inheriting 75% parameters from the original model. The results indicate that there is 0.5% slight decrease in top-1 accuracy on ImageNet, however, has ~2% improvements in retrieval tasks. TinyCLIP ResNet-19M inherits the weights from the 30M model, and it is trained for 12 epochs. It reduces the parameters by 50% while getting 2× inference speedup. But it gets 2.7% performance drops on ImageNet due to the large reduction in parameters.

For OpenCLIP ViT-B/32, we distill it using our affinity

mimicking method with manual or automatic weight inheritance. We design a three-stage progressive distillation in which each stage compresses the model by 25%. The three stages are trained for 6, 16, and 16 epochs, respectively. In manual inheritance, TinyCLIP ViT-61M/32 obtains comparable performance compared to OpenCLIP ViT-B/32 [20], while reducing the parameters by 29%. On the other hand, automatic inheritance obtains better results than the manual one. For example, TinyCLIP ViT-63M/32 outperforms the manually inherited ViT-61M/32 model by more than ~1% on both classification and retrieval tasks. When auto-compressed the model by 3.9×, TinyCLIP ViT-22M/32 obtains the highest inference throughput of 5,504 and reaches 53.7% top-1 accuracy. Since YFCC-15M contains high-quality image-text pairs, if using it to replace parts of LAION data, our method can obtain additional improvements, especially for retrieval tasks.

### 4.3. Ablation Study

*Impact of affinity mimicking.* As shown in Tab. 2, we present four different interaction modes. The affinity mimicking  $\mathcal{L}_1$  (*i.e.*,  $\mathcal{L}_{distill}$ ) outperforms the contrastive loss  $\mathcal{L}_0$  by 2.1% in terms of top-1 accuracy on ImageNet [7]. The alignment in a visual-linguistic affinity space brings more similarity information. The accuracy of cross modalities interaction is lower than that of affinity mimicking by 0.2%. In this mode, the embedding features of student image and text do not have any interaction, but they are aligned to the teacher’s embedding space. Single modality interaction has lower accuracy due to missing a loss to align student’s image and text features in common embedding space.

*Impact of weight inheritance.* We verify the efficacy of

Interaction Mode	Loss Formula	Top-1 Acc
Contrastive loss [46]	$\mathcal{L}_0 = CE(\langle I_s, T_s \rangle, \mathcal{I})$	53.4 %
Affinity mimicking	$\mathcal{L}_1 = CE(\langle I_s, T_s \rangle, \langle I_t, T_t \rangle)$	55.5 %
Cross modalities	$\mathcal{L}_2 = CE(\langle I_s, T_t \rangle, \langle I_t, T_t \rangle)$	55.3 %
	$\mathcal{L}_3 = CE(\langle I_t, T_s \rangle, \langle I_t, T_t \rangle)$	
Single modality	$\mathcal{L}_4 = CE(\langle I_s, I_t \rangle, \mathcal{I})$	19.2 %
	$\mathcal{L}_5 = CE(\langle T_s, T_t \rangle, \mathcal{I})$	

Table 2. The interaction of information across image  $I$  and text  $T$  modalities in student and teacher models. The subscripts  $s$  and  $t$  denote student and teacher, respectively.  $\mathcal{I}$  is an identity matrix.  $CE$  is cross entropy function. TinyCLIP-ViT-40M/32 is trained with an interaction mode and manual weight inheritance on LAION-400M [50] for 1 epoch. The zero-shot accuracy on ImageNet [7] is reported.




Figure 4. The cosine similarity of the input and output embedding of each layer, and the zero-shot accuracy on ImageNet [7] after removing a layer for OpenCLIP ViT-B-32 [20, 46] pretrained on LAION-2B [49].

Weight Inheritance	IN-1K top1 acc (%)	Flickr30k		MSCOCO	
		I→T@1	T→I@1	I→T@1	T→I@1
None	37.8	50.2	33.4	29.7	16.6
Manual	54.0	70.4	52.9	46.5	29.3
Automatic	54.9	73.5	55.9	47.9	30.8

Table 3. Manual inheritance vs. Automatic inheritance. The model OpenCLIP ViT-B/32 [20] is compressed by 50% with distillation on LAION-400M [50] for 1 epoch.

the proposed two kinds of weight inheritance in Tab. 3. Compared to the model without using weight inheritance, manual inheritance brings 16.2% accuracy improvements on ImageNet [7]. Automatic weight inheritance can further improve the accuracy by 0.9%. On image-text retrievals, the improvements of automatic inheritance is relatively higher, *i.e.*, 1.4 – 3.1% over the manual method.

*The redundancy of language-image models.* We analyze the redundancy of image encoder and text encoder by similar strategy [6, 2, 25] and zero-shot accuracy by removing a layer. As shown in Fig. 4, we observe a high correlation between cosine similarity and accuracy. Comparing with image encoder, text encoder has higher cosine similarity. It shows the redundancy of the text encoder is higher than that of the image encoder. Consequently, compressing the text encoder along the depth dimension can be achieved without significant performance degradation.

We further explore the redundancy by the proposed automatic weight inheritance. As shown in Fig. 5, for width

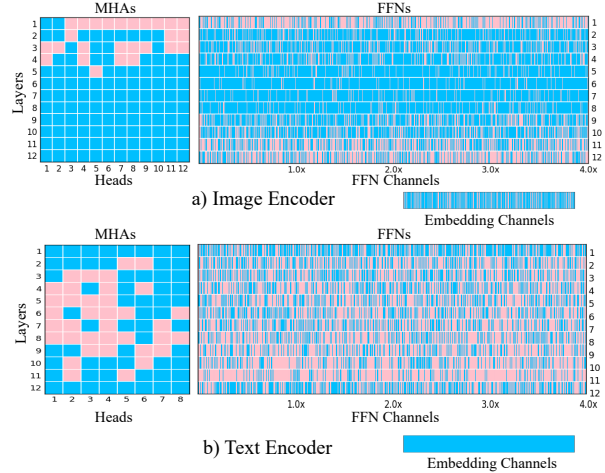


Figure 5. Visualization of learnable masks at 50% sparsity on hidden dimension, attention heads and MLP intermediate dimension. The blocks in red color are removed. The original model is OpenCLIP ViT-B/32 [20].

Pretrained Teacher Model	Teacher Acc.(%)	Inherited Ratio	Student Acc.(%)
w/o weight inheritance	-	0	36.2
CLIP ViT-B/32 [46]	63.2	59 / 126	52.4(+16.2)
OpenCLIP ViT-B/32 [20]	62.9	59 / 126	53.5(+17.3)
OpenCLIP ViT-B/16 [20]	67.1	59 / 124	52.8(+16.6)
OpenCLIP ViT-L/14 [20]	75.3	59 / 390	45.1(+8.9)
OpenCLIP ViT-H/14 [20]	78.0	59 / 935	41.1(+4.9)

Table 4. Ablation study on inherited teacher model. The student model TinyCLIP ViT-40M/32 with 59M parameters is inherited from a teacher model, then trained without distillation on LAION-400M [50] for 1 epoch. The top-1 zero-shot accuracy on ImageNet [7] is reported.

analysis, it is observed that the embedding channels of the text encoder experience only a marginal reduction (*e.g.*, 512 → 511), whereas the embedding channels of the image encoder are reduced to a greater extent (*e.g.*, 768 → 526). Regarding depth analysis, it was observed that a significant proportion of MHA layers and FFN channels could be removed from the text encoder, indicating redundancy within it. This finding suggests that the image encoder can be compressed along the width dimension, while the text encoder can be compressed along the depth dimension.

*Impact of teacher models.* We investigate which model is better for weight inheritance. As analyzed in Tab. 4, a strong teacher does not necessarily guarantee better results for weight inheritance. Instead, teacher models with similar architectures and higher performance offer a better choice. For instance, although ViT-H/14 ranks as the highest-performing teacher model, it lags behind other models in terms of weight inheritance.

*Impact of multi-stage progressive distillation.* Multi-stage progressive distillation shrinks the gap between the compressed model and the inherited model in each stage.

Method	Image Encoder	ImageNet																						
		Food101	CIFAR10	CIFAR100	SUN397	Stanford Cars	FGVC Aircraft	VOC2007	DTD	Oxford Pets	Caltech101	Flowers102	MINIST	FER2013	STL10	EuroSAT	RESISC45	GTSRB	KITTI	Country211	PCam	HatefulMemes	Rendered SST2	ImageNet
<i>Zero-shot performance</i>																								
DeCLIP [31]	ViT-B/32	76.3	90.0	67.9	63.5	50.6	8.9	80.5	45.3	84.4	89.3	83.2	12.8	16.3	97.5	40.0	48.4	10.3	35.3	11.5	54.7	52.9	50.1	66.1
FILIP [70]	ViT-B/32	82.8	86.9	65.5	69.1	55.5	57.2	-	49.3	88.1	91.9	85.3	-	-	-	-	49.9	-	-	-	-	-	-	68.8
CLIP [46]	ViT-B/32	84.4	91.3	65.1	63.2	59.4	21.2	83.1	44.5	87.0	87.9	66.7	51.9	47.3	97.2	49.4	60.3	32.2	39.4	17.8	58.4	57.6	59.6	63.2
OpenCLIP [20]	ViT-B/32	80.9	90.7	70.6	66.8	79.3	16.6	82.2	54.4	86.5	90.1	66.0	37.4	42.3	95.6	51.6	57.6	42.0	31.6	14.8	50.1	52.9	52.3	62.9
<b>TinyCLIP (Ours)</b>	ViT-45M/32	78.3	92.2	71.8	66.0	79.9	23.0	83.2	53.8	85.3	88.7	60.7	62.5	45.2	96.2	52.2	61.8	41.1	16.9	16.5	51.5	53.5	52.6	61.2
<i>Linear probe performance</i>																								
CLIP [46]	ViT-B/32	88.8	95.1	80.5	76.6	81.8	52.0	87.7	76.5	90.0	93.0	96.9	99.0	69.2	98.3	97.0	90.5	85.3	66.2	27.8	83.9	66.7	70.8	76.1
<b>TinyCLIP (Ours)</b>	ViT-45M/32	84.0	97.4	86.0	68.9	85.4	39.0	85.4	71.2	89.9	93.4	82.1	99.5	66.9	97.6	97.5	88.1	98.5	79.7	16.1	90.4	58.4	62.5	73.2

Table 5. Zero-shot and linear-probe classification top-1 accuracy on 23 datasets.

Number of Stages	Model Size	Epoch	ImageNet	MSCOCO	
			top1 acc (%)	I→T R@1	T→I R@1
1	100% → 25%	2	47.0	39.4	23.1
			54.0	46.5	29.3
2	100% → 50%	1	49.0	40.2	25.3
			54.0	46.5	29.3

Table 6. Ablation study on multi-stage progressive distillation. The model OpenCLIP ViT-B/32 [20] is compressed to 25% model size on LAION-400M [50] for 2 epochs.




Figure 6. Training efficiency. OpenCLIP ViT-B/32 [20] is trained from scratch on LAION-400M [50] for 32 epochs. TinyCLIP ViT/32 is trained for 6+16 epochs using the proposed 2-stage progressive distillation, where the teacher is OpenCLIP ViT-B/32.

We compare it with single-stage in Tab. 6. Using the same training epochs, the model trained with 2-stage surpasses that with 1-stage by 2.0% in terms of zero-shot accuracy on ImageNet. It demonstrates that multi-stage progressive compression is superior when transferring the knowledge of the pre-trained model to the compressed small model.

*Training cost.* We also study the distillation efficiency. As shown in Fig. 6, the original OpenCLIP ViT-B/32 [20] takes around 287 GPU days when trained on LAION-400M for 32 epochs. Our models with 75% and 50% parameters reach the similar accuracy as OpenCLIP counterparts, getting a speed up of 7.8 $\times$  and 1.4 $\times$ , respectively. The underlying reason is that our training benefits more from the affinity mimicking supervision and the good initialization using weight inheritance. This also proves that weight inheritance is capable of accelerating cross-modal distillation.

#### 4.4. Transfer Learning Results

*Classification on 23 datasets.* We also evaluate the performance of zero-shot and linear-probe classification on 23

Model	IN-V2	IN-A	IN-R	ObjectNet	IN-Sketch
CLIP ViT-B/32 [46]	56.0	31.6	69.4	29.9	42.3
OpenCLIP ViT-B/32 [20]	55.1	21.7	73.4	28.9	49.4
<b>TinyCLIP ViT-63M/32(Ours)</b>	55.7	22.8	74.1	31.2	50.8
<b>TinyCLIP ViT-45M/32(Ours)</b>	52.6	19.8	71.5	29.1	48.8

Table 7. Zero-shot robustness evaluation. CLIP [46] is trained on WIT-400M, OpenCLIP [20] and TinyCLIP is trained on LAION-400M [50]. Top-1 accuracy is reported.

datasets by the toolkit Elevater [26]. As shown in Tab. 5, our model wins on 7 datasets in zero-shot benchmark, and 9 datasets on linear-probe benchmark. The performance of our model is close to OpenCLIP ViT-B/32 [20], where the two models are both trained on Laion-400M [50].

*Zero-shot robustness evaluation.* We verify the robustness of our models on 5 datasets, which are out of ImageNet distribution. As shown in Tab. 7, our TinyCLIP ViT-63M/32 outperforms OpenCLIP ViT-B/32 [20], and wins 3 datasets when compared with CLIP ViT-B/32 [46]. Our TinyCLIP ViT-45M/32 uses 50% fewer parameters yet still wins on IN-R and IN-Sketch, compared to CLIP ViT-B/32. This demonstrates the robustness of the distilled TinyCLIP models and the efficacy of our proposed distillation method.

## 5. Conclusions

In this paper, we propose TinyCLIP, a method for distilling large-scale vision-language pre-trained models. It mainly introduced two core techniques named affinity mimicking and weight inheritance. Extensive experiments and ablation studies have demonstrated the effectiveness of TinyCLIP, showing that it can largely reduce model size while maintaining competitive performance. In the future, we will explore ways to further improve cross-modal distillation efficiency on extremely small models.

## Acknowledgement

Prof. Hongyang Chao was partially supported by NSFC (U1611461, 61672548) and NSFC (U22A2095). Prof. Xinggong Wang was partly supported by NSFC (No. 62276108).



## References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. In *NeurIPS*, 2022. 1
- [2] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. Deep equilibrium models. *NeurIPS*, 32, 2019. 7
- [3] Arnav Chavan, Zhiqiang Shen, Zhuang Liu, Zechun Liu, Kwang-Ting Cheng, and Eric P Xing. Vision transformer slimming: Multi-dimension searching in continuous optimization space. In *CVPR*, 2022. 3
- [4] Fei-Long Chen, Du-Zhen Zhang, Ming-Lun Han, Xiu-Yi Chen, Jing Shi, Shuang Xu, and Bo Xu. Vlp: A survey on vision-language pre-training. *Machine Intelligence Research*, 2023. 2
- [5] Wenliang Dai, Lu Hou, Lifeng Shang, Xin Jiang, Qun Liu, and Pascale Fung. Enabling multimodal generation on clip via vision-language knowledge distillation. In *ACL*, 2022. 2
- [6] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Lukasz Kaiser. Universal transformers. In *ICLR*, 2019. 7
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 2, 5, 6, 7
- [8] Mingyu Ding, Bin Xiao, Noel Codella, Ping Luo, Jingdong Wang, and Lu Yuan. Davit: Dual attention vision transformers. In *ECCV*, 2022. 5, 6
- [9] 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. *ICLR*, 2021. 5
- [10] Yifan Du, Zikang Liu, Junyi Li, and Wayne Xin Zhao. A survey of vision-language pre-trained models. In *IJCAI*, 2022. 2
- [11] Zhiyuan Fang, Jianfeng Wang, Xiaowei Hu, Lijuan Wang, Yezhou Yang, and Zicheng Liu. Compressing visual-linguistic model via knowledge distillation. In *ICCV*, 2021. 2
- [12] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *ICLR*, 2019. 3
- [13] Prakhar Ganesh, Yao Chen, Xin Lou, Mohammad Ali Khan, Yin Yang, Hassan Sajjad, Preslav Nakov, Deming Chen, and Marianne Winslett. Compressing large-scale transformer-based models: A case study on bert. *ACL*, 9, 2021. 1
- [14] Yuting Gao, Jinfeng Liu, Zihan Xu, Jun Zhang, Ke Li, and Chunhua Shen. Pyramidclip: Hierarchical feature alignment for vision-language model pretraining. In *NeurIPS*, 2022. 2
- [15] Yunhao Ge, Jie Ren, Ming-Hsuan Yang, Yuxiao Wang, Andrew Gallagher, Hartwig Adam, Laurent Itti, Balaji Lakshminarayanan, and Jiaping Zhao. Improving zero-shot generalization and robustness of multi-modal models. In *NeurIPS ML Safety Workshop*, 2022. 2
- [16] Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. Knowledge distillation: A survey. *IJCV*, 129:1789–1819, 2021. 1
- [17] Manish Gupta and Puneet Agrawal. Compression of deep learning models for text: A survey. *TKDD*, 2022. 2, 3
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [19] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv*, 2015. 1, 2
- [20] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hananeh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip, July 2021. 1, 2, 5, 6, 7, 8
- [21] 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 *ICML*, 2021. 1, 2
- [22] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. Tinybert: Distilling bert for natural language understanding. In *EMNLP*, 2020. 2
- [23] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *ECCV*, 2022. 2
- [24] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2015. 5
- [25] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *ICLR*, 2020. 7
- [26] Chunyuan Li, Haotian Liu, Liunian Harold Li, Pengchuan Zhang, Jyoti Aneja, Jianwei Yang, Ping Jin, Yong Jae Lee, Houdong Hu, Zicheng Liu, and Jianfeng Gao. Elevator: A benchmark and toolkit for evaluating language-augmented visual models. *NeurIPS*, 2022. 5, 8
- [27] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. 2
- [28] Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Deepak Gotmare, Shafiq Joty, Caiming Xiong, and Steven Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In *NeurIPS*, 2021. 2
- [29] Liunian Harold Li\*, Pengchuan Zhang\*, Haotian Zhang\*, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and Jianfeng Gao. Grounded language-image pre-training. In *CVPR*, 2022. 2
- [30] Yanghao Li, Haoqi Fan, Ronghang Hu, Christoph Feichtenhofer, and Kaiming He. Scaling language-image pre-training via masking. *arXiv preprint arXiv:2212.00794*, 2022. 2
- [31] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. In *ICLR*, 2022. 8
- [32] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yanan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana

- Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. *arXiv preprint arXiv:2210.04150*, 2022. [2](#)
- [33] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. [6](#)
- [34] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. In *ICLR*, 2019. [3](#)
- [35] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018. [5](#)
- [36] Jiawei Han, Jamie Callan, Luyu Gao, Yunyi Zhang. Scaling deep contrastive learning batch size under memory limited setup. In *ReplANLP*, 2021. [5](#)
- [37] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. Improved knowledge distillation via teacher assistant. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 5191–5198, 2020. [2](#)
- [38] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. In *ECCV*, 2022. [6](#)
- [39] Basil Mustafa, Carlos Riquelme Ruiz, Joan Puigcerver, Rodolphe Jenatton, and Neil Houlsby. Multimodal contrastive learning with limoe: the language-image mixture of experts. In *NeurIPS*, 2022. [1](#)
- [40] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, 2019. [5](#)
- [41] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *ICCV*, 2021. [2](#)
- [42] Renjing Pei, Jianzhuang Liu, Weimian Li, Bin Shao, Songcen Xu, Peng Dai, Juwei Lu, and Youliang Yan. Clipping: Distilling clip-based models for video-language understanding. *OpenReview*, 2023. [2](#)
- [43] Hieu Pham, Zihang Dai, Golnaz Ghiasi, Kenji Kawaguchi, Hanxiao Liu, Adams Wei Yu, Jiahui Yu, Yi-Ting Chen, Minh-Thang Luong, Yonghui Wu, et al. Combined scaling for open-vocabulary image classification. *arXiv*, 2021. [1](#), [2](#)
- [44] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, 2015. [6](#)
- [45] Di Qi, Lin Su, Jia Song, Edward Cui, Taroon Bharti, and Arun Sacheti. Imagebert: Cross-modal pre-training with large-scale weak-supervised image-text data. *arXiv preprint arXiv:2001.07966*, 2020. [2](#)
- [46] 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 *ICML*, 2021. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [47] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. [1](#), [2](#)
- [48] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019. [2](#)
- [49] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5b: An open large-scale dataset for training next generation image-text models. In *NeurIPS*, 2022. [5](#), [7](#)
- [50] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [51] Dachuan Shi, Chaofan Tao, Ying Jin, Zhendong Yang, Chun Yuan, and Jiaqi Wang. Upop: Unified and progressive pruning for compressing vision-language transformers. *arXiv preprint arXiv:2301.13741*, 2023. [2](#)
- [52] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *CVPR*, 2022. [1](#)
- [53] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 2016. [5](#), [6](#)
- [54] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, 2021. [2](#)
- [55] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. [6](#)
- [56] Mengmeng Wang, Jiazheng Xing, and Yong Liu. Actionclip: A new paradigm for video action recognition. *arXiv preprint arXiv:2109.08472*, 2021. [2](#)
- [57] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022. [2](#)
- [58] Yulong Wang, Xiaolu Zhang, Lingxi Xie, Jun Zhou, Hang Su, Bo Zhang, and Xiaolin Hu. Pruning from scratch. In *AAAI*, 2020. [3](#)
- [59] Zhecan Wang, Noel Codella, Yen-Chun Chen, Luowei Zhou, Jianwei Yang, Xiyang Dai, Bin Xiao, Haoxuan You, Shih-Fu Chang, and Lu Yuan. Clip-td: Clip targeted distillation for vision-language tasks. *arXiv preprint arXiv:2201.05729*, 2022. [2](#)
- [60] Ziheng Wang, Jeremy Wohlwend, and Tao Lei. Structured pruning of large language models. In *EMNLP*, 2020. [4](#)

- [61] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. In *ICLR*, 2021. 2
- [62] Yixuan Wei, Yue Cao, Zheng Zhang, Zhuliang Yao, Zhenda Xie, Han Hu, and Baining Guo. icar: Bridging image classification and image-text alignment for visual recognition. In *NeurIPS*, 2022. 2
- [63] Ross Wightman. Pytorch image models, 2019. 5
- [64] Kan Wu, Jinnian Zhang, Houwen Peng, Mengchen Liu, Bin Xiao, Jianlong Fu, and Lu Yuan. Tinyvit: Fast pretraining distillation for small vision transformers. In *ECCV*, 2022. 1, 2
- [65] Mengzhou Xia, Zexuan Zhong, and Danqi Chen. Structured pruning learns compact and accurate models. In *ACL*, 2022. 2, 3, 4
- [66] Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In *ECCV*, 2022. 2
- [67] Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Liqun Chen, Belinda Zeng, Trishul Chilimbi, and Junzhou Huang. Vision-language pre-training with triple contrastive learning, 2022. 2
- [68] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Bin Xiao, Ce Liu, Lu Yuan, and Jianfeng Gao. Unified contrastive learning in image-text-label space. In *CVPR*, 2022. 2
- [69] Shusheng Yang, Yixiao Ge, Kun Yi, Dian Li, Ying Shan, Xiaohu Qie, and Xinggang Wang. Masked visual reconstruction in language semantic space. *arXiv preprint arXiv:2301.06958*, 2023. 6
- [70] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. In *ICLR*, 2022. 8
- [71] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. In *NeurIPS*, 2022. 1
- [72] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. 1, 2, 5, 6
- [73] Yan Zeng, Xinsong Zhang, and Hang Li. Multi-grained vision language pre-training: Aligning texts with visual concepts. *arXiv preprint arXiv:2111.08276*, 2021. 1
- [74] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *CVPR*, pages 18123–18133, 2022. 2
- [75] Haotian Zhang, Pengchuan Zhang, Xiaowei Hu, Yen-Chun Chen, Liunian Harold Li, Xiyang Dai, Lijuan Wang, Lu Yuan, Jenq-Neng Hwang, and Jianfeng Gao. Glipv2: Unifying localization and vision-language understanding. In *NeurIPS*, 2022. 2
- [76] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luowei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *CVPR*, 2022. 2
- [77] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *ECCV*, 2022. 6