

CLIP-CID: Efficient CLIP Distillation via Cluster-Instance Discrimination

Kaicheng Yang^{1*}, Tiancheng Gu^{2*}, Xiang An¹, Haiqiang Jiang¹,
Xiangzi Dai¹, Ziyong Feng¹, Weidong Cai^{2‡}, Jiankang Deng^{3‡}

¹ DeepGlint

² University of Sydney

³ Imperial College

kaichengyang@deepglint.com

Abstract

Contrastive Language-Image Pre-training (CLIP) has achieved excellent performance over a wide range of tasks. However, the effectiveness of CLIP heavily relies on a substantial corpus of pre-training data, resulting in notable consumption of computational resources. Although knowledge distillation has been widely applied in single modality models, how to efficiently expand knowledge distillation to vision-language foundation models with extensive data remains relatively unexplored. In this paper, we introduce CLIP-CID, a novel distillation mechanism that effectively transfers knowledge from a large vision-language foundation model to a smaller model. We initially propose a simple but efficient image semantic balance method to reduce transfer learning bias and improve distillation efficiency. This method filters out 43.7% of image-text pairs from the LAION400M while maintaining superior performance. After that, we leverage cluster-instance discrimination to facilitate knowledge transfer from the teacher model to the student model, thereby empowering the student model to acquire a holistic semantic comprehension of the pre-training data. Experimental results demonstrate that CLIP-CID achieves state-of-the-art performance on various downstream tasks including linear probe and zero-shot classification.

1 Introduction

With the proliferation of mobile networks and social platforms, there has been an explosion in the production of image-text pairs (Guo, Wang, and Wang 2019; Gu et al. 2024). This abundance of data has provided a strong foundation for the advancement of vision-language pre-training (Radford et al. 2021; Jia et al. 2021; Yu et al. 2020; Yang, Xu, and Gao 2020). The Contrastive Language-Image Pre-training (CLIP) (Radford et al. 2021) achieves remarkable success in multi-modal learning by aligning image-text pairs on a large-scale dataset. It learns two separate unimodal encoders for image and text using a contrastive loss, one of the most effective losses for representation learning (Tian, Krishnan, and Isola 2020; He et al. 2020; Chen et al. 2020; Chopra, Hadsell, and LeCun 2005). Nevertheless, the efficacy of CLIP heavily

* Equal contribution.

‡ Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

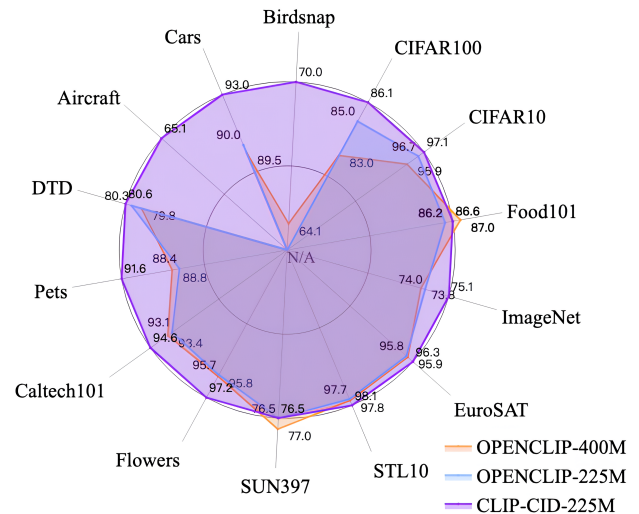


Figure 1: The linear probe performance comparison between CLIP-CID and OPENCLIP across 14 common datasets. Despite the exclusion of 43.7% of image-text pairs from the LAION400M, CLIP-CID exhibits exceptional performance.

depends on an extensive pre-training dataset. The original CLIP models are pre-trained on 400 million image-text pairs for 32 epochs, demanding thousands of GPU days. This poses a substantial challenge in scenarios with limited computational resources (Radenovic et al. 2023; Yang et al. 2023). Recently, large-scale image-text datasets crawled from websites, such as LAION400M (Schuhmann et al. 2021) and LAION5B (Schuhmann et al. 2022), have gained widespread usage in vision-language pre-training. DataComp (Gadre et al. 2024) consists of image-text pairs extracted from Common Crawl’s web data and employs various strategies such as basic filtering, CLIP score filtering, and text&image-based filtering. However, there is still a lot of semantic repetition in the training data, which not only has the potential to impact the performance of representation learning but also results in a waste of computational resources (Radenovic et al. 2023; Wang et al. 2023).

Knowledge Distillation (KD) (Hinton, Vinyals, and Dean 2015) is proposed to enhance the performance of a small stu-

dent model by transferring knowledge from a large teacher model throughout the training phase. Most existing KD approaches in the literature are primarily designed for small-scale datasets (*e.g.*, CIFAR10 and CIFAR100 (Krizhevsky, Hinton et al. 2009)), as well as small models (*e.g.*, ResNet50 and ResNet34 (He et al. 2016)). Recent studies have concentrated on distilling CLIP for particular target tasks or datasets. For example, BeamCLIP (Kim et al. 2022) introduces cross-modal similarity matching and context-based prompt augmentation to transfer knowledge from CLIP representations to a small model, achieving better performance on the ImageNet (Deng et al. 2009). ZeroSeg (Chen et al. 2023) distills the visual concepts learned by CLIP into a set of segment tokens, leading to a marked enhancement in training efficiency while preserving segmentation performance. Nevertheless, the exploration of leveraging knowledge distillation to enhance foundational models remains relatively sparse. Recent works (Wu et al. 2023; Sun et al. 2023) successfully transfer knowledge from a large foundation model to a small one. However, their attention is restricted to instance-level knowledge, thus unable to capture the semantic structure of extensive training data effectively. This limitation arises from the nature of instance-wise contrastive learning, which consistently regards samples from distinct instances as negative pairs, disregarding their semantic similarity.

To address the aforementioned challenges, this paper introduces CLIP-CID, a novel distillation mechanism that effectively transfers knowledge from a large vision-language foundation model to a smaller model. In order to reduce transfer learning bias and improve distillation efficiency, we initially propose a simple but efficient method to balance the semantic concepts within the LAION400M dataset, which can filter out 43.7% of the training data while maintaining superior performance. Furthermore, we enhance the knowledge transfer from a large teacher model to a smaller student model by integrating cluster-instance discrimination, which facilitates a more comprehensive semantic understanding of the student model. As illustrated in Fig. 1, our proposed CLIP-CID demonstrates superior linear probe performance across 14 common datasets even after filtering out 43.7% of image-text pairs from the LAION400M. The main contributions of this paper are summarized as follows:

- We propose a simple but efficient image semantic balance method to reduce transfer learning bias and improve distillation efficiency, which can remove 43.7% of image-text pairs from the LAION400M while maintaining superior performance.
- We introduce the CLIP-CID, a novel distillation mechanism that integrates cluster discrimination and instance discrimination to effectively transfer knowledge from a large vision-language foundation model to a smaller model.
- We conduct extensive experiments to validate the effectiveness of our proposed approach. The experimental results prove that CLIP-CID achieves state-of-the-art performance on various downstream tasks, including zero-shot classification and linear probe.

2 Related Work

2.1 Vision-Language Pre-training

CLIP (Radford et al. 2021) has achieved remarkable attention due to its exceptional zero-shot recognition ability and successful transfer capabilities. Recent developments have improved CLIP-based methodologies (Mu et al. 2022; Geng et al. 2023). ALBEF (Li et al. 2021) introduces a contrastive loss to align the image and text representations before fusing them through cross-modal attention, which enables more grounded vision and language representation learning. ALIGN (Jia et al. 2021) leverages a dataset of over one billion noisy image alt-text pairs to scale visual and vision-language representation learning. FILIP (Yao et al. 2021) successfully exploits the finer-grained expressiveness between image patches and textual words by modifying the contrastive loss and enables offline pre-computation of image and text representations during inference. FLIP (Li et al. 2023) employs random masking and removal of a significant portion of image patches during training, enabling learning from a larger number of image-text pairs within the same wall-clock time and contrasting more samples per iteration while maintaining a comparable memory footprint. However, pre-training vision-language foundation models on extensive datasets present a substantial challenge due to the high costs involved and the significant consumption of computational resources.

2.2 Large-Scale Dataset Filtering

While open-source large-scale datasets such as LAION400M (Schuhmann et al. 2021) have filtered image-text pairs based on CLIP scores below 0.3, it is crucial to emphasize that the effectiveness of CLIP scores may be compromised in specific scenarios. For instance, images containing visible text within the image often yield high CLIP scores but could prove detrimental to representation learning (Cao et al. 2023). Recent research proposes a complexity, action, and textspotting filtering strategy (Radenovic et al. 2023) to select informative image-text pairs from noisy web-scale datasets. Additionally, the pruning method adapted from ImageNet to LAION (Abbas et al. 2024) uses a straightforward complexity measure to reduce training costs to one-quarter of the standard approach. However, the process of cleaning large-scale datasets presents a computational challenge. To tackle this issue, SemDeDup (Abbas et al. 2023) leverages embeddings from pre-trained models to identify and eliminate semantic duplicates. Different from the above methods, we propose a simple but efficient image semantic balance method to reduce transfer learning bias and improve distillation efficiency. It only entails a single traversal of the data and can filter 43.7% of image-text pairs from the LAION400M while maintaining superior performance.

2.3 Knowledge Distillation

Knowledge distillation has found wide application in the field of computer vision. Building upon the achievements of visual-language foundational models, recent research efforts (Wei et al. 2022; Dong et al. 2022) have demon-

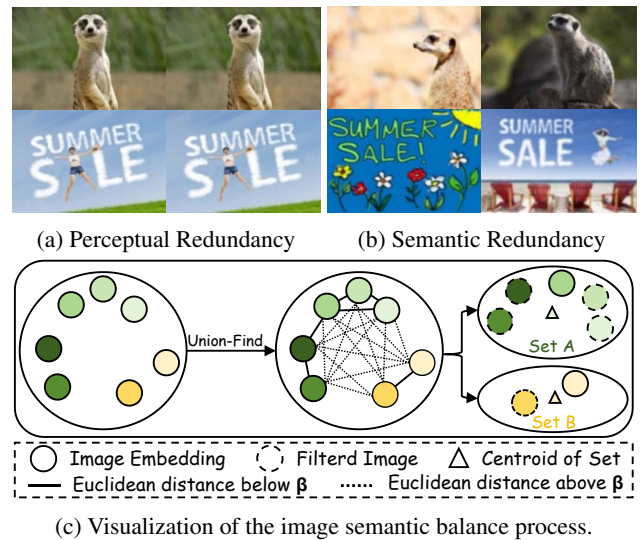
strated significant performance enhancements on specific datasets, such as ImageNet (Deng et al. 2009). Hybrid Distillation (Shi et al. 2023) achieves superior performance on COCO (Lin et al. 2014) by integrating masked autoencoders with CLIP. CLIP-TD (Wang et al. 2022) efficiently distills knowledge from CLIP into established architectures through a dynamically weighted objective concentrating on adaptively chosen tokens per instance, resulting in remarkable performance gains on visual commonsense reasoning tasks. DIME-FM (Sun et al. 2023) aims to distillate small foundation models using smaller-scale public images and unpaired sentences. However, the above methods only focus on instance-level knowledge, neglecting the semantic structure of large-scale training data. This limitation arises from instance-wise contrastive learning, which pairs samples as negatives based solely on their instance differences, without considering their semantic similarity.

3 Methodology

3.1 Image Semantic Balance

To enhance the comprehensive transfer of knowledge from the teacher model to the student model and reduce transfer learning bias, we propose a simple yet effective image semantic balance method. Our primary objective is to address both perceptual redundancy and semantic redundancy within images. As illustrated in Fig. 2a, perceptual redundancy images showcase minimal discrepancies at the pixel level. In contrast, as depicted in Fig. 2b, semantic redundancy images exhibit significant pixel-level variations while maintaining notably similar semantic information. As shown in Fig. 2d, the presence of perceptual redundancy and semantic redundancy in images causes an imbalanced distribution of semantic concepts in the dataset, leading to knowledge transfer biases in the student model during distillation.

In contrast to existing cluster-based methods that require multiple iterations (Abbas et al. 2024, 2023), our method requires only a single complete traversal of the training data to integrate images with similar semantics into the same set. The image semantic balance process is shown in Fig. 2c. Initially, we employ the OPENCLIP ViT-bigG/14 (Ilharco et al. 2021) model to extract image embeddings from LAION400M (Schuhmann et al. 2021). To address memory limitations, we divide all image embeddings $E \in \mathbb{R}^{N \times d}$ into c chunks and distribute $E_n \in \mathbb{R}^{\frac{N}{c} \times d}$ to different nodes. Then we calculate the Euclidean distance (Dokmanic et al. 2015) matrix between the current image and the images contained in various chunks, arranging the rows in ascending order according to their distances. The top- k results for each chunk are retained in the matrix $C_n \in \mathbb{R}^{\frac{N}{c} \times k}$, and these distance matrices are concatenated to form the global matrix $C \in \mathbb{R}^{N \times k}$. After that, we employ the Union-Find algorithm (Tarjan 1975) to group semantically similar images together. If the distance between different images surpasses the distance threshold β , they will be assigned to separate sets. Otherwise, they will be merged into the same set. We identify the central image (the image nearest to the centroid of the set) and eliminate the remaining images within the same set. Consequently, the total count of LAION400M im-



(a) Perceptual Redundancy (b) Semantic Redundancy
(c) Visualization of the image semantic balance process.
(d) Distribution of LAION400M and LAION225M.

Figure 2: (a) and (b) visualization of the perceptual redundancy images and semantic redundancy images. (c) visualization of the image semantic balance process. (d) distribution of LAION400M and LAION225M in 1M clusters.

ages is decreased to 225M, denoted as LAION225M. It is worth noting that the GPU is exclusively used for computing embedding distances, while the union-find algorithm is executed on the CPU. As shown in Fig. 2d, LAION225M demonstrates a smoother distribution, facilitating the student model learning a more comprehensive knowledge from the teacher model.

3.2 Cluster-level Distillation

Traditional instance-wise contrastive learning treats different instances as negative pairs, limiting its ability to capture the complete semantic information in the training data (Caron et al. 2018; Asano, Rupprecht, and Vedaldi 2020; Zhan et al. 2020; Qian et al. 2022). In this study, we introduce cluster discrimination knowledge distillation to delve into potential semantic structures inherent in the training dataset. Our method involves grouping visually similar instances into clusters, thus enabling a more comprehensive semantic representation. The cluster discrimination distillation method involves two stages: (1) Clustering, which assigns a unique class label to each image and gets cluster centers. (2) Cluster discrimination distillation, which fa-

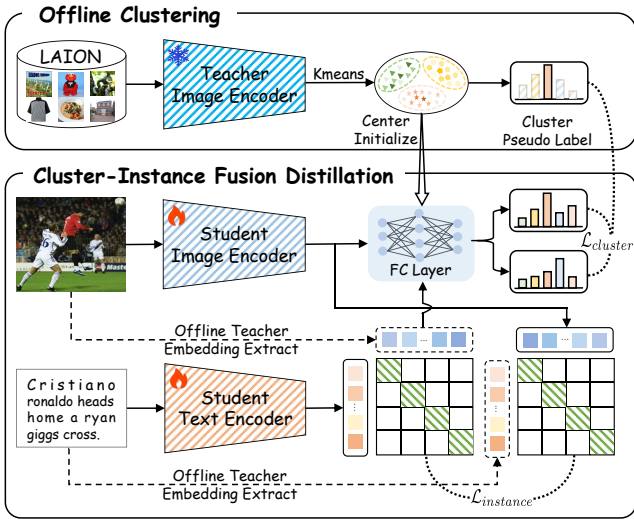


Figure 3: The architecture of our proposed cluster-instance discrimination distillation.

facilitates the transfer of global knowledge from the teacher model to the student model.

Clustering. We investigate the standard k -means algorithm, which aims to partition a given set of vectors into k distinct groups based on the nearest neighbor criterion. Given the normalized image embedding e_i , the clustering process involves jointly learning a centroid matrix $W \in \mathbb{R}^{d \times k}$ and assigning the cluster label z_i for each image by solving the following optimization problem:

$$\min_{W \in \mathbb{R}^{d \times k}} \frac{1}{N} \sum_{i=1}^N \min_{z_i \in \{0,1\}^k} \|e_i - W z_i\|_2^2 \quad \text{s.t.} \quad z_i^\top \mathbf{1}_k = 1, \quad (1)$$

where N is the number of training samples, and the centroid w_i belonging to centroid matrix $W \in \mathbb{R}^{d \times k}$ is considered the normalized prototype of i -th cluster. z_i in $\{0,1\}^k$ is a single label assignment restricted by the condition $z_i^\top \mathbf{1}_k = 1$, where $\mathbf{1}_k$ is 1-vector with size of k .

In this work, we employ the OPENCLIP ViT-bigG/14 (Harcro et al. 2021) model to extract image embeddings. The automatically clustered large-scale dataset inevitably faces challenges such as intra-class impurity and inter-class conflicts due to the presence of noise in the large uncurated web-scale datasets. The intra-class impurity can be addressed by adjusting the cluster number. Meanwhile, the inter-class conflict can be effectively mitigated by reducing the number of sampled negative instances within the minibatch. Following previous works (An et al. 2023, 2024), we leverage the benefits of efficient feature quantization (Johnson, Douze, and Jégou 2019) to cluster LAION225M into one million classes. To alleviate inter-class conflict, we employ PatialFC (An et al. 2022) and randomly sample a portion of the negative class centers during each iteration.

Cluster Discrimination Distillation. After clustering, we can inherit the original mechanism in the vanilla KD (Hinton, Vinyals, and Dean 2015) to implement instance-cluster

alignment. As illustrated in Fig. 3, we consider a set of training images $I = \{x_1, x_2, \dots, x_n\}$, comprising n images. Initially, we employ the student image encoder and the teacher image encoder to get normalized student image embeddings $E_i^s = \{e_1^s, e_2^s, \dots, e_n^s\}$ and normalized teacher image embeddings $E_i^t = \{e_1^t, e_2^t, \dots, e_n^t\}$ (the teacher image embeddings are extracted offline). These normalized image embeddings $e_i^s \in \mathbb{R}^d$ and $e_i^t \in \mathbb{R}^d$ are passed through a fully connected layer that is initialized using the cluster centers. It is important to note that without this initialization will lead to the collapse of model training. Subsequently, the images are partitioned into k classes, represented by prototypes $W = \{w_i\}_{i=1}^k$. With pseudo labels and cluster centers obtained from the above clustering step, the logit distillation loss can be implemented by optimizing a standard softmax classification loss as:

$$\mathcal{L}_1 = - \sum_{i=1}^n \log \frac{\exp(w_i^\top e_i^s)}{\sum_{j=1}^k \exp(w_j^\top e_i^s)}. \quad (2)$$

Furthermore, we implement distribution alignment by minimizing the Kullback-Leibler (KL) divergence between the prediction probability from the teacher and the student:

$$\mathcal{L}_d = \sum_{i=1}^n \text{KL} \left(\frac{\exp(w_i^\top e_i^s / \tau)}{\sum_{j=1}^k \exp(w_j^\top e_i^s / \tau)} \parallel \frac{\exp(w_i^\top e_i^t / \tau)}{\sum_{j=1}^k \exp(w_j^\top e_i^t / \tau)} \right), \quad (3)$$

where τ is the temperature hyper-parameter used to soften distribution representation. Finally, the cluster-level distillation loss $\mathcal{L}_{\text{cluster}}$ is defined as:

$$\mathcal{L}_{\text{cluster}} = \alpha \mathcal{L}_1 + (1 - \alpha) \mathcal{L}_d, \quad (4)$$

where α is a loss weight used to balance the influence of different losses.

3.3 Instance-level Distillation

The cluster-level distillation primarily impacts the student image encoder, facilitating the model in capturing comprehensive semantic information from the training data. However, it may unintentionally neglect the subtle nuances of fine-grained semantic details and impose limitations on image-text alignment. To address this issue, we introduce the instance-level distillation loss. Given an image-text pair, we first get the offline extracted teacher image embedding e_i^t and teacher text embedding c_i^t . After that, we use e_i^t and c_i^t to supervise the student text embedding c_i^s and student image embedding e_i^s respectively. We employ the bi-directional contrastive loss (Sohn 2016; Srivastava and Salakhutdinov 2012) to align the teacher embedding and the student embedding, which is defined as:

$$\begin{aligned} \mathcal{L}_{\text{contrast}}(e, c) &= \frac{1}{2} (\mathcal{L}_{e \rightarrow c} + \mathcal{L}_{c \rightarrow e}), \quad \text{where} \\ \mathcal{L}_{e \rightarrow c} &= - \sum_{i=1}^n \log \frac{\exp(e_i^\top c_i / \tau)}{\sum_{j=1}^n \exp(e_i^\top c_j / \tau)}, \\ \mathcal{L}_{c \rightarrow e} &= - \sum_{i=1}^n \log \frac{\exp(c_i^\top e_i / \tau)}{\sum_{j=1}^n \exp(c_j^\top e_i / \tau)}. \end{aligned} \quad (5)$$

Then the instance-level distillation loss $\mathcal{L}_{\text{instance}}$ is defined as:

Model	Dataset	Food101	CIFAR10	CIFAR100	Birdsnap	Cars	Aircraft	DTD	Pets	Caltech101	Flowers	SUN397	STL10	EuroSAT	ImageNet	Average
Model Architecture: ViT-B/32																
CLIP [†]	WIT400M	84.4	91.3	65.1	37.8	59.4	21.2	44.5	87.0	87.9	66.7	63.2	97.2	49.4	63.2	65.6
CLIP [‡]	WIT400M	82.9	88.7	63.7	35.4	57.3	18.9	43.3	84.0	89.3	66.0	61.5	96.6	43.9	61.9	63.8
OPENCLIP [‡]	LAION400M	80.5	90.6	70.7	42.6	78.1	15.9	51.3	85.8	91.2	66.0	66.7	95.2	49.7	62.1	67.6
OPENCLIP [*]	LAION225M	79.6	92.5	74.5	43.8	79.9	15.7	50.9	84.4	90.2	65.5	65.2	94.0	44.7	62.2	67.4
CLIP-CID	LAION225M	79.7	92.6	75.2	47.7	83.7	25.6	52.2	88.4	90.8	69.1	66.0	93.2	45.3	62.7	69.4
Model Architecture: ViT-B/16																
CLIP [†]	WIT400M	89.2	91.6	68.7	39.1	65.6	27.1	46.0	88.9	89.3	70.4	65.2	98.2	54.1	68.6	68.7
CLIP [‡]	WIT400M	87.8	89.6	66.4	40.9	63.5	23.1	44.8	87.3	90.4	67.6	63.0	98.0	52.9	67.2	67.3
OPENCLIP [‡]	LAION400M	85.7	91.8	71.1	46.4	82.7	16.5	50.2	88.6	91.9	66.0	68.6	96.8	51.1	65.6	69.5
OPENCLIP [*]	LAION225M	85.8	93.3	75.7	50.4	83.4	16.3	51.0	87.8	92.0	67.0	66.0	96.5	50.8	65.6	70.1
CLIP-CID	LAION225M	84.4	93.9	76.9	51.5	85.2	25.5	51.7	88.6	92.0	69.0	66.6	95.1	51.3	65.8	71.3

Table 1: Zero-shot classification comparison. We present zero-shot performance on 14 common downstream datasets. [†]: Results reported in CLIP paper. [‡]: Results we reproduced. ^{*}: Results of the OPENCLIP model trained on LAION225M.

$$\mathcal{L}_{\text{instance}} = \gamma \mathcal{L}_{\text{contrast}}(e_i^s, c_i^t) + (1 - \gamma) \mathcal{L}_{\text{contrast}}(c_i^s, e_i^t), \quad (6)$$

where γ is a loss weight. Finally, the overall loss function is defined as:

$$\mathcal{L}_{\text{overall}} = \mathcal{L}_{\text{base}} + \mathcal{L}_{\text{cluster}} + \mathcal{L}_{\text{instance}}, \quad (7)$$

where $\mathcal{L}_{\text{base}} = \mathcal{L}_{\text{contrast}}(e_i^s, c_i^s)$ represents the standard CLIP loss.

4 Experiments and Results

4.1 Experimental Settings

Implementation Details. In this paper, we utilize the OPENCLIP bigG/14 as the teacher model. The student model adopts the same architecture as CLIP (Radford et al. 2021). We employ AdamW (Loshchilov and Hutter 2017) as the optimizer, initializing it with a learning rate of $1e - 3$ and a weight decay of 0.2. To prevent collapsing, the initial learning rate of the fully connected layer is set to $1e - 6$. Based on empirical observations, we set the loss weight α and γ to 0.999 and 0.5 respectively. We set β_1 to 0.9 and β_2 to 0.98 for improved training stability. The input image size is 224×224 , and the input text sequence length is truncated or padded to 77. The temperature parameter τ is initialized to 0.07. We conduct distillation training on ViT-B/32 and ViT-B/16 models for 32 epochs, using a batch size of 32, 768 and 24, 576 on 64 NVIDIA H800 GPUs.

Downstream Datasets. To prove the effectiveness of our method, we present linear probe and zero-shot classification performance across 14 datasets, including Food101 (Bossard, Guillaumin, and Van Gool 2014), CIFAR10 & CIFAR100 (Krizhevsky, Hinton et al. 2009), Birdsnap (Berg et al. 2014), Stanford Cars (Krause et al. 2013),

Aircraft (Maji et al. 2013), DTD (Cimpoi et al. 2014), Oxford Pets (Parkhi et al. 2012), Caltech101 (Fei-Fei, Fergus, and Perona 2004), Flowers102 (Nilsback and Zisserman 2008), SUN397 (Xiao et al. 2010), STL10 (Coates, Ng, and Lee 2011), EuroSAT (Helber et al. 2019), and ImageNet (Deng et al. 2009). Additionally, to evaluate the robustness of our model, we conduct zero-shot robustness comparison on ImageNet-V2 (Recht et al. 2019), ImageNet-A (Recht et al. 2019), ImageNet-R (Hendrycks et al. 2021), ObjectNet (Barbu et al. 2019), and ImageNet-Sketch (Wang et al. 2019).

4.2 Experimental Results

Zero-shot Classification. We present our performance on 14 zero-shot classification datasets. The prompt templates and class names are consistent with CLIP (Radford et al. 2021). As shown in Tab. 1, the OPENCLIP trained on LAION225M achieves comparable performance with trained on LAION400M. This is primarily attributed to the removal of semantically redundant images, which enhances both the semantic balance and diversity of the pre-training dataset. Furthermore, through the integration of cluster-instance discrimination distillation, CLIP-CID ViT-B/32 and CLIP-CID ViT-B/16 achieve an average performance of 69.4% and 71.3% respectively across the 14 datasets, surpassing OPENCLIP trained on LAION400M by 1.8% and 1.8%. This performance improvement demonstrates the effectiveness of the cluster-instance discrimination distillation in enhancing student representations.

Linear Probe. Following the same evaluation setting as CLIP (Radford et al. 2021), we freeze our model and only train a logistic regression classifier. In Tab. 2, we present linear probe performance on 14 downstream datasets. Similarity with zero-shot classification, we observe that after removing 43.7% of the training data, the OPENCLIP

Model	Dataset	Food101	CIFAR10	CIFAR100	Birdsnap	Cars	Aircraft	DTD	Pets	Caltech101	Flowers	SUN397	STL10	EuroSAT	ImageNet	Average
Model Architecture: ViT-B/32																
CLIP [†]	WIT400M	88.8	95.1	80.5	58.5	81.8	52.0	76.5	90.0	93.0	96.9	76.6	98.3	97.0	76.1	82.9
CLIP [‡]	WIT400M	88.6	95.0	80.2	61.8	81.3	50.6	76.4	89.3	92.7	94.6	76.9	98.2	95.0	75.0	82.5
OPENCLIP [‡]	LAION400M	87.0	95.9	83.0	64.1	89.5	54.3	79.8	88.8	93.4	95.8	77.0	97.8	95.9	73.8	84.0
OPENCLIP*	LAION225M	86.2	96.7	85.0	62.8	90.0	55.4	80.3	88.4	93.1	95.7	76.5	97.7	95.8	74.0	84.1
CLIP-CID	LAION225M	86.6	97.1	86.1	70.0	93.0	65.1	80.6	91.6	94.6	97.2	76.5	98.1	96.3	75.1	86.3
Model Architecture: ViT-B/16																
CLIP [†]	WIT400M	92.8	96.2	83.1	67.8	86.7	59.5	79.2	93.1	94.7	98.1	78.4	99.0	97.1	80.2	86.1
CLIP [‡]	WIT400M	92.7	96.0	82.5	72.4	86.4	59.7	78.9	93.1	94.0	96.5	78.6	99.1	95.3	79.3	86.0
OPENCLIP [‡]	LAION400M	90.8	96.4	84.0	71.7	92.0	59.5	81.5	91.7	94.7	97.3	79.4	98.7	96.1	77.7	86.5
OPENCLIP*	LAION225M	90.1	97.3	86.4	73.2	92.8	59.0	81.4	91.7	94.6	96.9	78.5	98.5	96.8	77.6	86.8
CLIP-CID	LAION225M	90.5	97.4	87.2	77.2	93.8	70.9	81.8	92.9	95.0	98.4	78.3	98.6	96.8	78.0	88.3

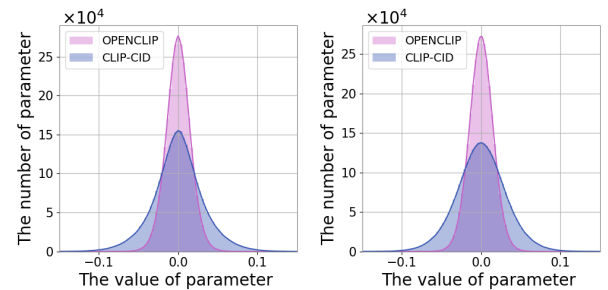
Table 2: Linear probe comparison. We present linear probe performance on 14 common downstream datasets. [†]: Results reported in CLIP paper. [‡]: Results we reproduced. *: Results of the OPENCLIP model trained on LAION225M.

Model	Dataset	IN-V2	IN-A	IN-R	Object	IN-S	Average
Model Architecture: ViT-B/32							
CLIP [‡]	WIT400M	54.9	31.9	67.1	52.6	38.8	49.1
OPENCLIP [‡]	LAION400M	54.6	22.2	72.3	52.7	46.5	49.7
OPENCLIP*	LAION225M	54.0	21.2	71.4	51.3	46.5	48.9
CLIP-CID	LAION225M	54.7	21.9	72.5	52.9	47.1	49.8
Model Architecture: ViT-B/16							
CLIP [‡]	WIT400M	61.1	49.6	76.0	59.0	45.1	58.1
OPENCLIP [‡]	LAION400M	59.1	33.4	76.8	57.0	49.8	55.2
OPENCLIP*	LAION225M	58.2	31.4	76.5	55.7	50.3	54.4
CLIP-CID	LAION225M	59.3	32.9	77.1	57.5	50.5	55.5

Table 3: Zero-shot robustness comparison. [‡]: Results we reproduced. *: Results of the OPENCLIP model trained on LAION225M.

trained on the filtered LAION225M achieves similar performance with trained on the entire LAION400M. By employing our proposed cluster-instance discrimination distillation method, CLIP-CID demonstrates an average performance improvement of 2.3% and 1.8% across 14 datasets. Notably, our method exhibits superior performance on CIFAR10&CIFAR100, Oxford Pets, Birdsnap, Stanford Car, and Aircraft, further substantiating the effectiveness of our approach in significantly enhancing the representation power for instance discrimination.

Zero-shot Robustness Evaluation. In Tab. 3, we present a robustness evaluation across different model sizes. We observe that excluding 43.7% of the training data leads to a marginal decrease in robustness. Subsequent integration of the cluster-instance discrimination distillation mechanism enables our model to acquire knowledge from the teacher model efficiently and effectively. Consequently, our model consistently exhibits superior robustness compared to OPENCLIP. In Fig. 4, we visually depict the weight distribution of the last fully connected layer in both the middle and last transformer layers.



(a) Middle Transformer Layer (b) Last Transformer Layer

Figure 4: Weight distribution of the last fully connected layer in the middle and last transformer layers.

tribution of the last fully connected layer in both the middle and last transformer layers. After distillation, we observe an expanded range of values within the weight distribution of our model, accompanied by a reduction in the number of elements proximal to zero. This phenomenon reflects the improvement of capacity since the weights can assume a more diverse array of potential values or states (Shen and Savvides 2020).

4.3 Ablation Study

Ablation on Threshold β . To explore the optimal image filtering ratio of the LAION400M, we perform an ablation study on threshold β . The value of β is associated with the number of sets, which directly influences the proportion of removed images. We train standard OPENCLIP ViT-B/32 on the filtered dataset with various values of the threshold β . As shown in Tab. 4, setting β to 0.07 results in the removal of 43.7% of image-text pairs, and we observe the optimal performance. However, increasing β to 0.08 raised the filtration

β	Dataset	Filtration Ratio	Linear probe	Zero-shot
0.06	LAION237M	40.8%	83.4	66.2
0.07	LAION225M	43.7%	84.1	67.4
0.08	LAION210M	47.5%	83.7	67.1

Table 4: Ablation on different threshold β .

$\mathcal{L}_{\text{base}}$	$\mathcal{L}_{\text{cluster}}$	$\mathcal{L}_{\text{instance}}$	Dataset	Linear probe	Zero-shot
✓	✗	✗	LAION225M	84.1	67.4
✓	✓	✗	LAION225M	85.7	68.3
✓	✗	✓	LAION225M	85.3	68.8
✓	✓	✓	LAION225M	86.3	69.4

Table 5: Ablation on different loss combinations.

rate to 47.5%, leading to a significant performance decline. **Ablation on Distillation Loss.** We perform ablation experiments on distillation losses to validate our proposed cluster-instance distillation mechanism. As shown in Tab. 5, the integration of cluster discrimination distillation loss improves linear probe performance from 84.1% to 85.7%, while the marginal gain in zero-shot performance is only 0.9%. This is because the cluster-level distillation loss mainly affects the image encoder, enhancing its ability to learn comprehensive semantic information while imposing constraints on image alignment. After introducing the instance discrimination distillation loss, it improves cross-modal alignment and captures more fine-grained semantics, which significantly boosts the zero-shot result from 68.3% to 69.4%.

Ablation on Cluster Centers. The number of cluster centers is a vital factor in managing inter-class and intra-class conflicts. As demonstrated in Tab. 6, we present the average linear probe and zero-shot classification performance of CLIP-CID ViT-B/32. The increase in the number of cluster centers from 0.1 million to 1 million resulted in a corresponding enhancement in model performance. This enhancement can be attributed to the increased intra-cluster purity, signifying improved discrimination and representation capabilities of the model. However, due to increased inter-cluster conflicts, the performance deteriorates as the number of cluster centers rises from 1M to 5M.

Cluster Centers	Dataset	Linear probe	Zero-shot
0.1M	LAION225M	83.8	67.8
0.5M	LAION225M	84.6	68.3
1M	LAION225M	86.3	69.4
2M	LAION225M	85.8	69.0
5M	LAION225M	85.2	68.8

Table 6: Ablation on different numbers of cluster centers.

Ablation on Teacher Models. To explore the impact of the difference in parameter quantity between teacher and student models, we compare the performance of CLIP-CID ViT-B/32 distilled from different scale teacher models. The experiment results are shown in Tab. 7, we find a larger teacher model such as OPENCLIP ViT-bigG/14 produces better student performance for both linear probe and zero-shot classification.

Visualization of PCA Components. We present the results

Teacher Model	Dataset	Linear probe	Zero-shot
OPENCLIP ViT-L/14	LAION225M	85.5	68.7
OEPNCLIP ViT-bigG/14	LAION225M	86.3	69.4

Table 7: Ablation on different scales of teacher models.




Figure 5: Visualization of PCA components. We extract three principal components from the collected patch features of each image. The principal components are then visualized using separate color channels. Similar colors within patches indicate semantic similarities. We use \circ to accentuate the primary distinction.

of Principal Component Analysis (PCA) applied to patch features extracted by OPENCLIP ViT-B/32 and our CLIP-CID ViT-B/32. We retain patches with positive values after applying a threshold to the first component, effectively separating the main object from the background. Subsequently, a second PCA is computed on the remaining patches. To visualize the results, we assign three distinct colors to the first three components of each model. In Fig. 5, our model demonstrates superior semantic understanding by consistently preserving the consistent color representation of object parts across diverse images within the same category. For example, the visualization of our model consistently maintains the head of the goose in a consistent red color. However, the OPENCLIP displays three different colors. We also provide visualization of embeddings, clusters, and class activation maps, please refer to the supplementary material.

4.4 Conclusion

In this paper, we introduce CLIP-CID, a novel distillation mechanism that effectively transfers knowledge from a large vision-language foundation model to a smaller model. To mitigate transfer learning bias and enhance distillation efficiency, we propose an efficient image semantic balance method, which can filter 43.7% of image-text pairs from the LAION400M while maintaining superior performance. After that, we employ cluster-instance discrimination to facilitate knowledge transfer from the teacher to the student model, enabling the student model to achieve a comprehensive semantic understanding of the pre-training data. Experimental results demonstrate that CLIP-CID surpasses existing methods in various downstream tasks, including linear probe and zero-shot classification.

References

- Abbas, A.; Rusak, E.; Tirumala, K.; Brendel, W.; Chaudhuri, K.; and Morcos, A. S. 2024. Effective pruning of web-scale datasets based on complexity of concept clusters. *arXiv preprint arXiv:2401.04578*.
- Abbas, A.; Tirumala, K.; Simig, D.; Ganguli, S.; and Morcos, A. S. 2023. SemDeDup: Data-efficient learning at web-scale through semantic deduplication. *arXiv preprint arXiv:2303.09540*.
- An, X.; Deng, J.; Guo, J.; Feng, Z.; Zhu, X.; Yang, J.; and Liu, T. 2022. Killing two birds with one stone: Efficient and robust training of face recognition cnns by partial fc. In *CVPR*.
- An, X.; Deng, J.; Yang, K.; Li, J.; Feng, Z.; Guo, J.; Yang, J.; and Liu, T. 2023. Unicom: Universal and Compact Representation Learning for Image Retrieval. In *ICLR*.
- An, X.; Yang, K.; Dai, X.; Feng, Z.; and Deng, J. 2024. Multi-label Cluster Discrimination for Visual Representation Learning. In *ECCV*.
- Asano, Y. M.; Rupprecht, C.; and Vedaldi, A. 2020. Self-labelling via simultaneous clustering and representation learning. In *ICLR*.
- Barbu, A.; Mayo, D.; Alverio, J.; Luo, W.; Wang, C.; Gutfreund, D.; Tenenbaum, J.; and Katz, B. 2019. Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. *NeurIPS*.
- Berg, T.; Liu, J.; Woo Lee, S.; Alexander, M. L.; Jacobs, D. W.; and Belhumeur, P. N. 2014. Birdsnap: Large-scale fine-grained visual categorization of birds. In *CVPR*.
- Bossard, L.; Guillaumin, M.; and Van Gool, L. 2014. Food-101—mining discriminative components with random forests. In *ECCV*.
- Cao, L.; Zhang, B.; Chen, C.; Yang, Y.; Du, X.; Zhang, W.; Lu, Z.; and Zheng, Y. 2023. Less is More: Removing Text-regions Improves CLIP Training Efficiency and Robustness. *arXiv preprint arXiv:2305.05095*.
- Caron, M.; Bojanowski, P.; Joulin, A.; and Douze, M. 2018. Deep clustering for unsupervised learning of visual features. In *ECCV*.
- Chen, J.; Zhu, D.; Qian, G.; Ghanem, B.; Yan, Z.; Zhu, C.; Xiao, F.; Culatana, S. C.; and Elhoseiny, M. 2023. Exploring open-vocabulary semantic segmentation from clip vision encoder distillation only. In *CVPR*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *ICML*.
- Chopra, S.; Hadsell, R.; and LeCun, Y. 2005. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing textures in the wild. In *CVPR*.
- Coates, A.; Ng, A.; and Lee, H. 2011. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*.
- Dokmanic, I.; Parhizkar, R.; Ranieri, J.; and Vetterli, M. 2015. Euclidean distance matrices: essential theory, algorithms, and applications. *IEEE Signal Processing Magazine*.
- Dong, X.; Bao, J.; Zhang, T.; Chen, D.; Zhang, W.; Yuan, L.; Chen, D.; Wen, F.; and Yu, N. 2022. Bootstrapped masked autoencoders for vision BERT pretraining. In *ECCV*.
- Fei-Fei, L.; Fergus, R.; and Perona, P. 2004. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *CVPR*.
- Gadre, S. Y.; Ilharco, G.; Fang, A.; Hayase, J.; Smyrnis, G.; Nguyen, T.; Marten, R.; Wortsman, M.; Ghosh, D.; Zhang, J.; et al. 2024. Datacomp: In search of the next generation of multimodal datasets. *NeurIPS*.
- Geng, S.; Yuan, J.; Tian, Y.; Chen, Y.; and Zhang, Y. 2023. HiCLIP: Contrastive Language-Image Pretraining with Hierarchy-aware Attention. In *ICLR*.
- Gu, T.; Yang, K.; An, X.; Feng, Z.; Liu, D.; Cai, W.; and Deng, J. 2024. RWKV-CLIP: A Robust Vision-Language Representation Learner. *arXiv:2406.06973*.
- Guo, W.; Wang, J.; and Wang, S. 2019. Deep multimodal representation learning: A survey. *IEEE Access*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *CVPR*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Helber, P.; Bischke, B.; Dengel, A.; and Borth, D. 2019. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*.
- Hendrycks, D.; Basart, S.; Mu, N.; Kadavath, S.; Wang, F.; Dorundo, E.; Desai, R.; Zhu, T.; Parajuli, S.; Guo, M.; et al. 2021. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *CVPR*.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Ilharco, G.; Wortsman, M.; Wightman, R.; Gordon, C.; Carlini, N.; Taori, R.; Dave, A.; Shankar, V.; Namkoong, H.; Miller, J.; Hajishirzi, H.; Farhadi, A.; and Schmidt, L. 2021. OpenCLIP.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*.
- Johnson, J.; Douze, M.; and Jégou, H. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*.
- Kim, B.; Choi, S.; Hwang, D.; Lee, M.; and Lee, H. 2022. Transferring pre-trained multimodal representations with cross-modal similarity matching. *NeurIPS*.

- Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3d object representations for fine-grained categorization. In *ICCV*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Li, J.; Selvaraju, R.; Gotmare, A.; Joty, S.; Xiong, C.; and Hoi, S. C. H. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *NeurIPS*.
- Li, Y.; Fan, H.; Hu, R.; Feichtenhofer, C.; and He, K. 2023. Scaling language-image pre-training via masking. In *CVPR*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv:1711.05101*.
- Maji, S.; Rahtu, E.; Kannala, J.; Blaschko, M.; and Vedaldi, A. 2013. Fine-grained visual classification of aircraft. *arXiv:1306.5151*.
- Mu, N.; Kirillov, A.; Wagner, D.; and Xie, S. 2022. Slip: Self-supervision meets language-image pre-training. In *ECCV*.
- Nilsback, M.-E.; and Zisserman, A. 2008. Automated flower classification over a large number of classes. In *Sixth Indian Conference on Computer Vision, Graphics & Image Processing*.
- Parkhi, O. M.; Vedaldi, A.; Zisserman, A.; and Jawahar, C. 2012. Cats and dogs. In *ICCV*.
- Qian, Q.; Xu, Y.; Hu, J.; Li, H.; and Jin, R. 2022. Unsupervised visual representation learning by online constrained k-means. In *CVPR*.
- Radenovic, F.; Dubey, A.; Kadian, A.; Mihaylov, T.; Vandenhende, S.; Patel, Y.; Wen, Y.; Ramanathan, V.; and Mahajan, D. 2023. Filtering, distillation, and hard negatives for vision-language pre-training. In *CVPR*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Recht, B.; Roelofs, R.; Schmidt, L.; and Shankar, V. 2019. Do imagenet classifiers generalize to imagenet? In *ICML*.
- Schuhmann, C.; Beaumont, R.; Vencu, R.; Gordon, C.; Wightman, R.; Cherti, M.; Coombes, T.; Katta, A.; Mullis, C.; Wortsman, M.; et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *NeurIPS*.
- Schuhmann, C.; Vencu, R.; Beaumont, R.; Kaczmarczyk, R.; Mullis, C.; Katta, A.; Coombes, T.; Jitsev, J.; and Komatsuzaki, A. 2021. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*.
- Shen, Z.; and Savvides, M. 2020. Meal v2: Boosting vanilla resnet-50 to 80%+ top-1 accuracy on imagenet without tricks. *arXiv preprint arXiv:2009.08453*.
- Shi, B.; Zhang, X.; Wang, Y.; Li, J.; Dai, W.; Zou, J.; Xiong, H.; and Tian, Q. 2023. Hybrid Distillation: Connecting Masked Autoencoders with Contrastive Learners. *arXiv preprint arXiv:2306.15876*.
- Sohn, K. 2016. Improved deep metric learning with multi-class n-pair loss objective. *NeurIPS*.
- Srivastava, N.; and Salakhutdinov, R. R. 2012. Multimodal learning with deep boltzmann machines. *NeurIPS*.
- Sun, X.; Zhang, P.; Zhang, P.; Shah, H.; Saenko, K.; and Xia, X. 2023. DIME-FM: DIstilling Multimodal and Efficient Foundation Models. *arXiv preprint arXiv:2303.18232*.
- Tarjan, R. E. 1975. Efficiency of a good but not linear set union algorithm. *Journal of the ACM (JACM)*.
- Tian, Y.; Krishnan, D.; and Isola, P. 2020. Contrastive multiview coding. In *ECCV*.
- Wang, A. J.; Lin, K. Q.; Zhang, D. J.; Lei, S. W.; and Shou, M. Z. 2023. Too Large; Data Reduction for Vision-Language Pre-Training. *arXiv preprint arXiv:2305.20087*.
- Wang, H.; Ge, S.; Lipton, Z.; and Xing, E. P. 2019. Learning Robust Global Representations by Penalizing Local Predictive Power. In *NeurIPS*.
- Wang, Z.; Codella, N.; Chen, Y.-C.; Zhou, L.; Yang, J.; Dai, X.; Xiao, B.; You, H.; Chang, S.-F.; and Yuan, L. 2022. Clip-td: Clip targeted distillation for vision-language tasks. *arXiv preprint arXiv:2201.05729*.
- Wei, Y.; Hu, H.; Xie, Z.; Zhang, Z.; Cao, Y.; Bao, J.; Chen, D.; and Guo, B. 2022. Contrastive learning rivals masked image modeling in fine-tuning via feature distillation. *arXiv preprint arXiv:2205.14141*.
- Wu, K.; Peng, H.; Zhou, Z.; Xiao, B.; Liu, M.; Yuan, L.; Xuan, H.; Valenzuela, M.; Chen, X. S.; Wang, X.; et al. 2023. Tinyclip: Clip distillation via affinity mimicking and weight inheritance. In *CVPR*.
- Xiao, J.; Hays, J.; Ehinger, K. A.; Oliva, A.; and Torralba, A. 2010. Sun database: Large-scale scene recognition from abbey to zoo. In *ICCV*.
- Yang, K.; Deng, J.; An, X.; Li, J.; Feng, Z.; Guo, J.; Yang, J.; and Liu, T. 2023. Alip: Adaptive language-image pre-training with synthetic caption. In *ICCV*.
- Yang, K.; Xu, H.; and Gao, K. 2020. Cm-bert: Cross-modal bert for text-audio sentiment analysis. In *Proceedings of the 28th ACM international conference on multimedia*, 521–528.
- Yao, L.; Huang, R.; Hou, L.; Lu, G.; Niu, M.; Xu, H.; Liang, X.; Li, Z.; Jiang, X.; and Xu, C. 2021. Filip: Fine-grained interactive language-image pre-training. *arXiv preprint arXiv:2111.07783*.
- Yu, W.; Xu, H.; Meng, F.; Zhu, Y.; Ma, Y.; Wu, J.; Zou, J.; and Yang, K. 2020. Ch-sims: A chinese multimodal sentiment analysis dataset with fine-grained annotation of modality. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, 3718–3727.
- Zhan, X.; Xie, J.; Liu, Z.; Ong, Y.-S.; and Loy, C. C. 2020. Online deep clustering for unsupervised representation learning. In *CVPR*.