

All You Need in Knowledge Distillation Is a Tailored Coordinate System

Junjie Zhou, Ke Zhu, Jianxin Wu*

National Key Laboratory for Novel Software Technology, Nanjing University, China
School of Artificial Intelligence, Nanjing University, China
zhoujj@lamda.nju.edu.cn, zhuk@lamda.nju.edu.cn, wujx2001@gmail.com

Abstract

Knowledge Distillation (KD) is essential in transferring dark knowledge from a large teacher to a small student network, such that the student can be much more efficient than the teacher but with comparable accuracy. Existing KD methods, however, rely on a large teacher trained specifically for the target task, which is both very inflexible and inefficient. In this paper, we argue that a SSL-pretrained model can effectively act as the teacher and its dark knowledge can be captured by the coordinate system or linear subspace where the features lie in. We then need only one forward pass of the teacher, and then tailor the coordinate system (TCS) for the student network. Our TCS method is teacher-free and applies to diverse architectures, works well for KD and practical few-shot learning, and allows cross-architecture distillation with large capacity gap. Experiments show that TCS achieves significantly higher accuracy than state-of-the-art KD methods, while only requiring roughly half of their training time and GPU memory costs.

Introduction

Deep learning has improved areas such as computer vision and natural language processing. However, large model size and heavy compute make it difficult to deploy them on devices with limited resources. Model compression methods, especially knowledge distillation (Hinton, Vinyals, and Dean 2015), effectively compress and accelerate deep networks, as evidenced in various domains (Hinton, Vinyals, and Dean 2015; Wang, Ge, and Wu 2021; Yang et al. 2020).

Knowledge Distillation (KD) has emerged as a crucial technique for enhancing the efficiency of neural networks by transferring “dark knowledge” from a large complex model (the teacher) to a smaller one (the student). In spite of its many successes, KD still encounters many drawbacks and limitations:

- **Reliant** on a teacher trained for the specific task, which is highly costly, cumbersome and inflexible. Utilizing a *self-supervised-learning (SSL) pretrained model as the teacher* is highly desired, but is rarely studied or with limited accuracy (Koochpayegani, Tejankar, and Pirsiavash 2020). Self distillation by Zhang et al. (2019) is

teacher-free, but incurs significantly higher memory and time costs.

- **Inflexible.** KD mostly leverages the logits as training signals, which is inflexible, e.g., for complex tasks like object detection (Wang, Ge, and Wu 2021). Difficulties are also evident when the teacher has huge capacity but the student is a small network (Mirzadeh et al. 2020).
- **Inefficient.** Feature-based KD (Tian, Krishnan, and Isola 2019; Chen et al. 2021a; Wang, Ge, and Wu 2021; Miles, Elezi, and Deng 2024) is more flexible than logit-based ones by mimicking final and/or intermediate features. However, both feature- and logit-based KD methods require a well-trained teacher and the teacher’s forward passes, which nearly doubles the time and memory costs.

Our goal is to achieve *teacher-free, flexible and efficient* knowledge distillation. We will extract knowledge from SSL pretrained models, which does *not* require training for any specific task, that is, being teacher-free for KD. Two key questions remain open in the literature:

1. What is the dark knowledge inside a pretrained model?
2. How to efficiently adapt it for a specific task?

Our answer to both questions is to propose a tailored coordinate system (TCS). Our key hypothesis is that *the dark knowledge within a SSL pretrained model are at least partly encoded within the coordinate system* (or linear subspace) in which their features lie in. The linear subspace can be efficiently encoded via PCA (principal component analysis).

To adapt this coordinate system to a specific task and the student network, we propose to simply compute PCA of the teacher’s features using *one single forward pass* on that task’s training data. Furthermore, the key to adapt the coordinate system is to *tailor it by choosing a subset of the coordinates specifically for the target task*.

Combining both key techniques, we argue that all you need in knowledge distillation is a tailored coordination system, i.e., the proposed TCS method.

As shown in Figure 1, the proposed TCS achieves the highest accuracy with lowest memory and compute costs among KD methods. In addition, it enjoys the following benefits:

- **Teacher-free**, i.e., effectively transfers knowledge from SSL pretrained models.

*J. Wu is the corresponding author.




Figure 1: Top-1 accuracy, training time and GPU memory consumption on ImageNet-1K. Our TCS achieves the fastest training, smallest GPU memory footprint and highest accuracy among KD methods. In the left figure, x-axis is logarithmic. The teacher is Swin-Tiny and the student is ResNet18.

- Easily applicable to diverse architectures such as CNNs, Transformers, and MLPs.
- Spans both traditional KD and few-shot learning.
- Allows cross-architecture and huge capacity gap between teacher and student.
- Achieves almost the same training speed as training the student from scratch.

In summary, our TCS is flexible, efficient, and not reliant on task-specific teacher. Technically, our contributions are:

- We argue that tailored coordinate system (TCS) is all you need to transfer dark knowledge from a task-agnostic teacher to a task-specific student network. Our TCS, of course, can also use a teacher trained for the target task in a supervised manner.
- Our TCS method is flexible and versatile, which accommodates diverse architectures and tasks. It is efficient in both KD training and inference in terms of speed and memory footprint, and delivers superior accuracy on both traditional KD and few-shot learning.

Related Work

Knowledge Distillation (KD) is a line of methods to transfer knowledge from large, (often) high-capacity teacher networks to smaller, more efficient student ones.

Knowledge Distillation Originally proposed by Hinton, Vinyals, and Dean (2015), KD facilitates the transfer of knowledge in two primary domains: output and representation spaces. Direct logits transfer has been extensively explored (Jin et al. 2019; Li et al. 2020; Yuan et al. 2020; Yun et al. 2020; Mirzadeh et al. 2020; Son et al. 2021; Zhou et al. 2021; Kim et al. 2021; Zhao et al. 2022; Hao et al. 2024), along with methods focusing on feature or representation distillation (Zagoruyko and Komodakis 2016; Huang and Wang 2017; Kim, Park, and Kwak 2018; Park et al. 2019; Heo et al. 2019; Tian, Krishnan, and Isola 2019; Chen et al. 2021b; Wang, Ge, and Wu 2021; Zhu, He, and Wu 2023b). The core in KD is to encourage students to emulate or replicate the predictions of teacher models. KD has been widely applied across diverse learning scenarios (Yu et al. 2017; Chen et al. 2017; Yim et al. 2017; Schmitt et al. 2018).

KD methods, however, mostly requires a teacher that is specifically trained for the target task. SSL-pretrained models are less studied in KD (Koochpayegani, Tejankar, and Pirsivash 2020). And, KD methods nearly double the compute and GPU memory costs compared to training the student from scratch, even when excluding the costs associated with training the teacher itself. A recent method (Zhang et al. 2019) introduced self-distillation by condensing knowledge internally, thus effectively reducing the network’s size. This method is teacher-free but has very large compute and GPU memory costs.

Practical Few-Shot Learning Different from classical Few-Shot Learning (FSL), practical Few-Shot Learning or pFSL is tailored to better accommodate real-world scenarios (Fu and Zhu 2024). Traditional FSL predominantly relies on knowledge derived from a base set of labeled images, employing strategies such as meta-learning or transfer learning. Meta-learning methods train across multiple episodes, enabling rapid adaptation to new tasks with minimal data (Hospedales et al. 2022), while transfer-learning typically involves pretraining on base classes followed by fine-tuning on novel classes (Chen et al. 2019; Mangla et al. 2020).

In contrast, pFSL foregoes reliance on a predefined base set, opting instead for SSL models pretrained on large-scale datasets (Fu and Zhu 2024). Our distillation methodology for few-shot learning incorporates the pFSL paradigm to enhance real-world applicability and effectiveness.

Method

As depicted in Figure 2, our TCS starts by extracting the teacher’s features through the training dataset. The extracted features are used to compute the teacher’s coordinate system by Principal Component Analysis (PCA). This procedure does *not* involve data augmentation, which means *only one* forward pass is needed and the cost involving the teacher is in fact negligible.

We then align the student’s features to fit into this coordinate system. This alignment process is supported by our iterative feature selection method, which tailors the coordinate system to suit the target task. eLSH is another module we proposed by a simple modification to LSH feature mimicking of Wang, Ge, and Wu (2021), which further improves our TCS efficiently.

We present our TCS method and its modules one by one.

Preparations

In contrast to methods that utilize intermediate features from various layers within the backbone as supervisory signals, our TCS is inspired by the “feature mimicking” approach suggested by Wang, Ge, and Wu (2021), which only mimics features from the final (penultimate) layer in the backbone.

For an input image x , the teacher generates a feature vector $f^t \in \mathbb{R}^{D_t}$, while the student produces $f^s \in \mathbb{R}^{D_s}$. To reconcile the differing dimensionalities between them, a linear layer is inserted after the student backbone. It changes f^s from D_s dimensions to D_t by a fully connected (FC) layer, and the transformed feature vector is $\hat{f}^s \in \mathbb{R}^{D_t}$.

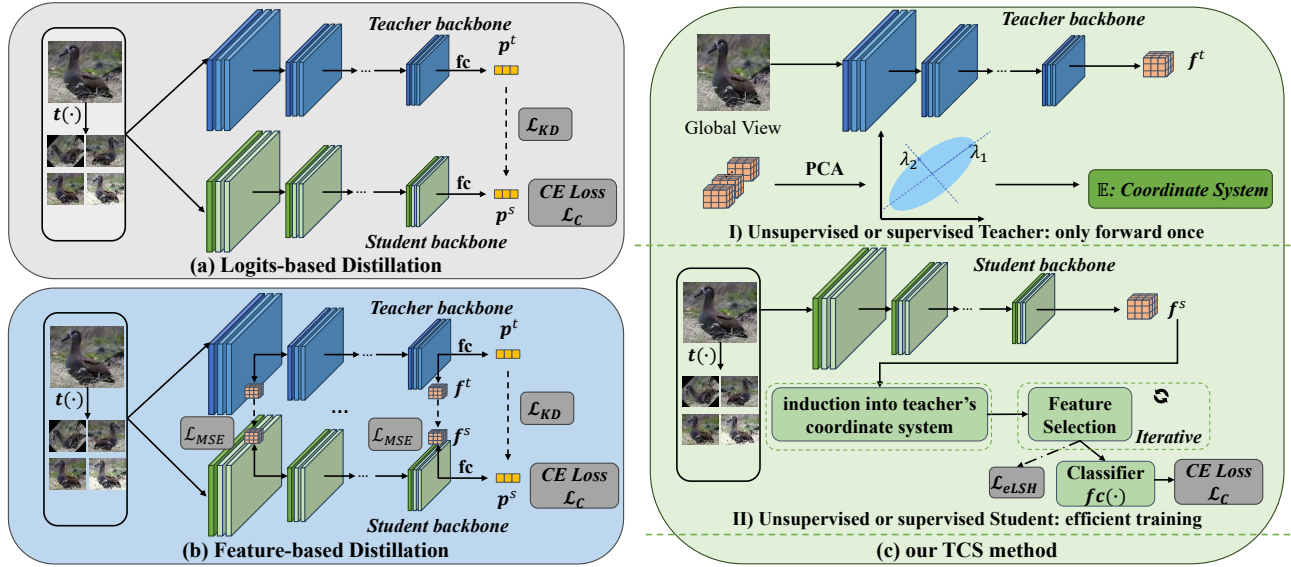


Figure 2: Illustration of existing KD methods and our TCS method. (a) logits-based distillation, where the student learns only from the final predictions of the teacher (p^t); (b) feature-based distillation, in which the student mimics intermediate features of the teacher (f^t), too; and (c) our TCS method, which aligns student features into the teacher’s coordinate system and *tailor* it to the target task by feature selection. TCS is optionally augmented by an eLSH module—an efficient unsupervised feature mimicking method. This figure is best viewed in color.

When the student is not trained from scratch, the linear transformation matrix W in this FC layer is initialized by using the feature sets $F^t = \{f_i^t\}_{i=1}^N$ ($N \times D_t$) and $F^s = \{f_i^s\}_{i=1}^N$ ($N \times D_s$), where N denotes the number of training images. W is initialized as the optimal least-square solution that regresses F^t from F^s :

$$W \leftarrow (F_s^T F_s)^{-1} F_s^T F_t. \quad (1)$$

Dark knowledge from a SSL-pretrained teacher

There is no doubt that a SSL-pretrained model (Zhu et al. 2024; Zhu, He, and Wu 2023a) contains useful dark knowledge for various downstream tasks, as even finetuning such models with few-shot samples lead to excellent recognition and detection results (Fu and Zhu 2024). However, the difficulty is *how to transfer such dark knowledge*?

Logit- or feature-based KD are both implicit—they mimic either the logit or feature of the teacher. But, because the SSL-pretrained model is *not* trained for the target task, its logit or feature is at best suboptimal (if not harmful).

Our theoretical hypothesis and empirical finding is that the *coordinate system, or the linear subspace where the teacher’s features reside in*, contains enough dark knowledge for various downstream target tasks.

We are inspired by the finding of Yu and Wu (2023), which states that the features of various models are low-rank and the PCA reduced lower-dimensional versions of these features contain almost the same amount of useful information as the original features.

In other words, the few orthonormal eigenvectors that define the projection directions in PCA form a new and (semantically almost equivalent) coordinate system for the

teacher’s features. And, because *we compute PCA using the target task’s training data*, this coordinate system already incorporates some characteristics of the target task in addition to the dark knowledge from the SSL-pretrained model.

Of course, there must be significant domain gap between the general-purpose SSL-pretrained model and the target task, or, we must *not* assume all PCA dimensions are useful for the target task. We further hypothesize that we can *select a subset of the PCA dimensions (or coordinates)* that suit the target task. Hence, our overall hypothesis is that we can *tailor* (i.e., feature selection) the teacher’s coordinate system defined via PCA. Hence, our proposed method is called TCS (tailored coordinate system).

We want to emphasize that even though our hypotheses assume a SSL-pretrained teacher (Zhu, Liu, and Huang 2023), the logic also clearly applies when the teacher is trained using the target task’s data. That is, our hypotheses are also suitable for traditional KD tasks.

Induction into teachers’ coordinate system

The teacher’s feature matrix, $F^t = \{f_i^t\}_{i=1}^N$, is $N \times D^t$ in size. The first step in PCA is to center the features by subtracting the mean of each column from its respective values:

$$\mu = \frac{1}{N} \sum_{i=1}^N f_i^t, \quad (2)$$

$$X = F^t - \mathbf{1}_N \mu^T, \quad (3)$$

where $\mathbf{1}_N$ is an $N \times 1$ vector of all ones.

PCA is then efficiently computed via the Singular Value Decomposition (SVD) of the centered matrix $X = U \Sigma V^T$,

where U and V are orthogonal matrices containing the left and right singular vectors of X , respectively, and Σ is a diagonal matrix containing the singular values. The columns of V (right singular vectors) represent the principal components of X . These principal components define the new coordinate system along which the variations of the data are maximized.

Induction of a student feature \hat{f}_s into this coordinate system is pretty simple: by projecting \hat{f}_s onto this coordinate system to obtain a new representation \tilde{f}^s :

$$\tilde{f}^s = (\hat{f}^s - \mu)V. \quad (4)$$

The features in \tilde{f}^s encode knowledge from the pretrained model. As showed by the distributed representation nature of deep features, one single dimension of f^t often does not correspond to semantic concepts well, but the combination of them can (Bengio, Courville, and Vincent 2013). For example, the linear combination formed by the first column in V (i.e., the first dimension in \tilde{f}^s) may have high correlation with the concept ‘‘animal’’. By encouraging the first dimension of $\tilde{f}^s = (\hat{f}^s - \mu)V$ and that of $(f^t - \mu)V$ to be similar, the student will benefit from this coordinate system when the target task is related to animals. Instead, classic KD method ask \hat{f}^s and f^t to be similar, which cannot enjoy such benefits due to the distributed representation nature.

Iterative feature selection

On one hand, it is also obvious that not all target tasks are related to all concepts as in the general-purpose SSL-pretrained teacher. For example, the target task may be flower recognition, which is irrelevant to animals. Many dimensions in \tilde{f}^s may well be irrelevant or even noise in the target task. We counter this difficulty by feature selection, i.e., choosing a subset of dimensions from \tilde{f}^s that are useful for the target task. The subset of features that are picked up forms the Tailored Coordinate System (TCS) for the student.

Since $\tilde{f}^s \in \mathbb{R}^{D_t}$, we introduce a trainable mask $\mathbf{m} \in \mathbb{R}^{D_t}$, which is initialized as $\mathbf{1}_{D_t}$ and gradually changes some dimensions in \mathbf{m} to 0. This simple feature selection strategy prunes irrelevant dimensions using the target task’s training data. Subsequently, the student features become

$$\bar{f}^s = \tilde{f}^s \odot \mathbf{m}. \quad (5)$$

For feature selection, we resort to the ℓ_1 -norm:

$$\mathcal{L} = \mathcal{L}_{CE} + \lambda \|\mathbf{m}\|_{\ell_1}, \quad (6)$$

where \mathcal{L}_{CE} is the loss for training the student model (with cross entropy being the most frequently used one), and λ is the hyperparameter to balance \mathcal{L}_{CE} and $\|\mathbf{m}\|_{\ell_1}$.

Our iterative feature selection module retains relevant dimensions by iteratively removing the influence of irrelevant ones. We accumulate the gradient with respect to \mathbf{m} :

$$\Delta^{(\tau)} \leftarrow \frac{1}{n} \sum_{i=1}^n \nabla_{\mathbf{m}} \mathcal{L}(\theta^{(\tau)}, \mathbf{m}^{(\tau)}), \quad (7)$$

where n is the batch size, τ indexes training iterations, and θ represents the parameters of the student model. We collect these gradient contributions over iterations, $\sum_{i=0}^{\tau} \Delta^{(i)}$,

to continuously assess each dimension’s relevance. The selection strategy at iteration τ is:

$$S^\tau \leftarrow \text{TopDims}\left(\sum_{i=0}^{\tau} \Delta^{(i)}, r_\tau \cdot D_t\right), \quad (8)$$

where r_τ represents the target selection ratio at that iteration and $\text{TopDims}(\mathbf{a}, b)$ picks the largest b dimensions from \mathbf{a} . The update rule for \mathbf{m} in each iteration is then:

$$\mathbf{m}^{(\tau+1)} \leftarrow \mathbf{m}^{(\tau)} \odot (1 - S^\tau) + \left(1 - \frac{r_\tau}{r}\right) \odot S^\tau. \quad (9)$$

This formula ensures a smooth transition towards the final target selection ratio r , which is set to 0.5 in all experiments.

It is worth noting that operations after the student’s backbone (linear layer to obtain \hat{f}^s , PCA to obtain \tilde{f}^s and iterative feature selection to obtain \bar{f}^s) are all *linear*. During inference, TCS can seamlessly integrate all these modules into the student’s classifier. Hence, TCS does *not incur extra parameters or computational overhead* during the student’s inference.

Efficient feature mimicking

On the other hand, when a teacher model trained specifically for the target task is indeed available (i.e., the teacher is no longer a SSL-pretrained one and the student is trained from scratch), we believe that all dimensions in f^t are useful for guiding the student. In this case, feature selection is not needed, and we can use f^t to directly guide the learning of the student.

As we showed earlier, the coordinate system defined by PCA is beneficial for such guidance. Hence, we want \bar{f}^s to mimic $W^T(f^t - \mu)$, because they are the student and teacher features in this coordinate system, respectively.

To make these two feature vectors similar, we follow the LSH loss by Wang, Ge, and Wu (2021). The LSH loss efficiently hashes similar features into the same ‘bins’, which reduces computational costs in high-dimensional data spaces:

$$\mathbf{h} = \text{sign}(W^T(f^t - \mu)V + \mathbf{b}), \quad (10)$$

$$\mathbf{p} = \sigma(W^T\bar{f}^s + \mathbf{b}), \quad (11)$$

where σ is the sigmoid function and \mathbf{h} and \mathbf{p} are (hard and soft) hash codes for the teacher and student, respectively. Suppose there are M hash codes (i.e., $\mathbf{h}, \mathbf{p} \in \mathbb{R}^M$), the loss that forces \bar{f}^s to mimic $(f^t - \mu)V$ is

$$\mathcal{L}_{eLSH} = -\frac{1}{M} \sum_{j=1}^M [h_j \log p_j + (1 - h_j) \log(1 - p_j)], \quad (12)$$

following Wang, Ge, and Wu (2021). We also set all hyperparameters (M and the weight to combine eLSH with \mathcal{L}_{CE}) as the default values in (Wang, Ge, and Wu 2021).

However, a major difference between our eLSH loss and the LSH loss by Wang, Ge, and Wu (2021) is that we do *not* involve data augmentation in computing \mathbf{h} . Hence, only one forward pass is needed to compute \mathbf{h} , while in (Wang,

Teacher	Params	From Scratch		KD Methods								
		T.	S.	Self-distill	KD	DKD	CRD	DIST	OFA	LSHL2 [‡]	TCS-	TCS
ViT-S	21.7M	92.04	76.72	80.50	77.26	78.10	76.60	76.49	80.15	80.38	78.23	81.37
Memory (GB)		-	3.5	22.0	7.5	7.5	3.9	7.5	7.6	7.5	5.6	5.8
Runtime (s/ep)		-	29.34	385.36	49.90	50.83	68.05	50.91	55.07	52.02	29.88	32.14
ConvNeXt-T	27.9M	88.41	76.72	80.50	79.35	80.43	78.87	77.75	80.91	80.24	78.35	81.60
Memory (GB)		-	3.5	22.0	7.5	7.5	3.9	7.5	7.7	7.5	5.6	5.8
Runtime (s/ep)		-	29.34	385.36	55.71	56.66	73.09	57.06	61.07	56.41	30.84	32.71
Mixer-B/16	59.2M	87.29	76.72	80.50	77.79	78.67	76.42	76.36	79.39	79.37	78.80	81.02
Memory (GB)		-	3.5	22.0	7.6	7.6	4.2	7.6	8.1	7.6	5.7	5.8
Runtime (s/ep)		-	29.34	385.36	71.34	71.79	88.94	72.37	77.33	72.44	30.73	31.16

Table 1: Results of KD methods on CIFAR-100. The student is ResNet18, which has 11.6M parameters. Best results are shown in boldface. [‡] indicates results run by us, and other results are directly obtained from respective papers. All metrics were measured on a single NVIDIA 3090 GPU, and all experiments were conducted with the same batch size.

Ge, and Wu 2021) \mathbf{h} needs to be computed in every epoch. Hence, we name our loss as efficient LSH (eLSH).

For traditional KD, our method is also called TCS, and we refer to our method without eLSH as 'TCS-'. When the teacher is SSL-pretrained, eLSH is not used—as aforementioned, it is not suitable in that case.

It is also worth mentioning that the KL-divergence loss used for logit-distillation is *not* used in our method.

Experimental Results

We assess our method across two distinct experimental setups: traditional KD, and the pFSL framework applied to few-shot scenarios. We provide essential information about our experiments' data, training and evaluation in this section. For more details, please refer to our supplementary materials.

Experiments in traditional KD

We experimented with CIFAR-100 (Krizhevsky 2009) and ImageNet-1K (Russakovsky et al. 2015). CIFAR-100 contains 50,000 training and 10,000 test samples, each with a resolution of 32×32 pixels. ImageNet-1K is a larger dataset, comprising of 1.2 million training images and 50,000 validation images, all at a resolution of 224×224 pixels. To accommodate the input requirements of Vision Transformers (ViTs) and MLPs, which process image patches, we upsample CIFAR-100 images to 224×224 pixels for all experiments.

In our experiments, we trained the student network on CIFAR-100 using the code and hyperparameters provided by CRD (Tian, Krishnan, and Isola 2019) to ensure consistent and fair comparisons. For ImageNet, we adhered to the standard PyTorch practices, which trains a model for 100 epochs as per CRD guidelines to strike a balance between computational efficiency and learning quality.

To address randomness in the hashing process within the eLSH module, we followed recommendations of Wang, Ge, and Wu (2021) by using the average model from the final 10 epochs of training. This approach helps stabilize the model's

Teacher	Acc	KD Methods					
		KD	DKD	OFA	DIST	LSHL2 [‡]	TCS
DeiT-T	72.17	70.22	69.39	71.34	70.64	70.13	71.40
Memory	-	7.5	7.5	7.6	7.5	7.5	5.8
Runtime	-	415.33	419.02	481.28	419.36	419.69	328.24
Swin-T	81.38	71.14	71.10	71.85	70.91	70.16	71.99
Memory	-	7.5	7.5	7.7	7.5	7.5	5.8
Runtime	-	701.80	701.73	768.68	703.29	703.34	321.27
Mixer-B	76.62	70.89	69.89	71.38	70.66	70.18	71.45
Memory	-	7.6	7.6	8.1	7.6	7.6	5.8
Runtime	-	869.97	868.95	937.92	872.72	877.16	341.88

Table 2: Results on ImageNet-1K. ResNet18 is the student. When trained from scratch, its top-1 accuracy is 69.75%. Self-distillation (Zhang et al. 2019) achieved an accuracy of 70.63%, with a memory usage of 22 GB and runtime of 3207.45 seconds per epoch.

performance by averaging out fluctuations in the late training stages.

We evaluated with a variety of models with heterogeneous architectures. Specifically, we employed ResNet (He et al. 2016) as a representative CNN model, ViT (Dosovitskiy et al. 2020), DeiT (Touvron et al. 2021), Swin (Liu et al. 2021) as ViT models, and MLP-Mixer (Tolstikhin et al. 2021) for MLP-based ones. In our training framework, the features are those from the penultimate layer, which precedes the final classification layer. Specifically, for Vision Transformers (ViT), this corresponds to the CLS token, while for CNN-like models, it refers to the features after global pooling. We compared with classic KD (Hinton, Vinyals, and Dean 2015), DKD (Zhao et al. 2022), CRD (Tian, Krishnan, and Isola 2019), DIST (Huang et al. 2022), OFA (Hao et al. 2024) and LSHL2 (Wang, Ge, and Wu 2021).

Main Results As shown in Tables 1 and 2, our TCS, especially when equipped with the eLSH module, not only achieves clearly higher accuracy than state-of-the-art methods, but also significantly reduces memory and runtime

Dataset	Pretraining Method	Student	Teacher	Method	Shots per Class						
					1	2	3	4	5	8	16
CUB	MoCov3	-	-	baseline	18.5 ± 0.7	27.2 ± 0.2	35.5 ± 1.0	40.0 ± 0.6	45.3 ± 0.9	54.1 ± 0.4	65.3 ± 0.3
		ViT-S	-	IbM2	19.5 ± 0.4	27.6 ± 0.2	35.6 ± 0.7	40.2 ± 0.4	46.2 ± 1.1	55.9 ± 0.4	66.8 ± 0.5
		ViT-B	-	TCS	20.1 ± 0.9	33.1 ± 0.9	44.5 ± 1.0	49.2 ± 0.5	55.2 ± 0.5	63.0 ± 0.4	71.8 ± 0.1
	DINOv2	-	-	baseline	51.0 ± 1.4	66.0 ± 0.5	73.1 ± 0.8	75.1 ± 0.7	78.8 ± 0.5	82.7 ± 0.5	86.3 ± 0.3
		ViT-S	-	IbM2	54.4 ± 1.0	66.8 ± 1.0	75.0 ± 0.6	76.9 ± 0.5	79.8 ± 0.5	83.6 ± 0.5	86.5 ± 0.3
		ViT-B	-	TCS	53.8 ± 0.5	64.5 ± 0.3	74.7 ± 0.5	78.0 ± 0.4	80.9 ± 0.6	84.5 ± 0.2	87.6 ± 0.0
		ViT-L	-	TCS	54.7 ± 1.7	65.3 ± 0.8	75.3 ± 1.2	77.8 ± 0.4	81.2 ± 0.5	84.7 ± 0.3	87.5 ± 0.1
ImageNet-1K	MoCov3	-	-	baseline	33.1 ± 0.6	42.4 ± 0.3	47.2 ± 0.4	49.8 ± 0.4	51.6 ± 0.3	55.2 ± 0.2	59.2 ± 0.3
		ViT-S	-	IbM2	34.4 ± 0.6	43.8 ± 0.3	48.7 ± 0.4	51.5 ± 0.3	53.0 ± 0.2	56.2 ± 0.1	60.3 ± 0.2
		ViT-B	-	TCS	35.1 ± 0.4	46.0 ± 0.3	51.4 ± 0.4	54.2 ± 0.2	55.9 ± 0.3	58.7 ± 0.2	61.6 ± 0.1
	DINOv2	-	-	baseline	37.0 ± 0.7	46.9 ± 0.5	52.9 ± 0.3	57.0 ± 0.1	59.7 ± 0.4	64.0 ± 0.1	67.6 ± 0.0
		ViT-S	-	IbM2	37.2 ± 0.9	47.1 ± 0.5	53.1 ± 0.3	58.6 ± 0.3	61.6 ± 0.2	65.3 ± 0.1	68.9 ± 0.1
		ViT-B	-	TCS	34.2 ± 0.4	47.4 ± 0.6	54.7 ± 0.6	58.9 ± 0.3	61.3 ± 0.2	65.3 ± 0.4	69.7 ± 0.1
		ViT-L	-	TCS	35.8 ± 0.4	48.2 ± 0.6	55.3 ± 0.5	59.7 ± 0.3	61.7 ± 0.1	65.6 ± 0.4	69.7 ± 0.1

Table 3: Average of top-1 accuracy (%) plus standard deviation across 3 random subsets on ImageNet-1K and CUB. The best result in every comparison group is highlighted in boldface.

costs. The training time is almost halved, and the GPU memory footprint is reduced by roughly 25%.

The proposed TCS is versatile to be applied upon various network architectures and sizes, and the teacher-student pair can be of heterogeneous network types. The results show that even when the capacity gap between teacher and student is large (Mirzadeh et al. 2020), TCS copes well with it. Furthermore, TCS requires *only one* forward pass from the teacher, which explains why its training time is roughly half of other KD methods.

Compare to another teacher-free method, self-distillation by Zhang et al. (2019), TCS has clear advantages. Self-distillation requires significantly more parameters, larger GPU memory and computing time, but its accuracy lags behind TCS consistently.

Experiments in few-shot learning

Consistent with the protocols outlined in the pFSL (practical few-shot learning) paper (Fu and Zhu 2024), we freeze the backbones of both the teacher and student networks during the training phase. Despite the recognized challenges associated with distilling knowledge from teacher models trained via unsupervised or self-supervised methods, our approach effectively overcomes these hurdles. In our experiments, we exclusively employed self-supervised models as teachers, demonstrating our method’s capability to extract and distill valuable information from these models.

We experimented on two datasets: ImageNet-1K (Rusakovsky et al. 2015) and CUB-200-2011 (CUB) (Wah et al. 2011). CUB has 11,788 images across 200 bird species. For both datasets, we randomly sampled the training sets following the pFSL protocol by selecting k images per class (i.e., k -shot learning) and formed the novel set with all classes (1000 for ImageNet, 200 for CUB).

The evaluation was on the full validation set. This setting is different from classic FSL, where the number of classes is

very small (mostly 5, in comparison to 1000 or 200). pFSL argues that being able to handle a large number of classes is essential for making few-shot learning practical.

Experiments were conducted using two types of backbones: Vision Transformer (ViT) (Dosovitskiy et al. 2020) and ResNet50 (He et al. 2016). The teacher network’s pre-training methods included DINOv2 (Oquab et al. 2023) and MoCov3 (Chen, Xie, and He 2021). We mainly compare our TCS method with the IbM2 approach proposed by Fu and Zhu (2024).

Main Results Results in Table 3 show that our TCS method improves few-shot learning by comparing the top-1 accuracy of TCS versus IbM2. TCS achieves top accuracy in almost all cases, except in two experiments with either 1- or 2-shot learning. As the number of training shots k increases, TCS’s margin grows to about 1.0% and can exceed 2% in some cases. Another benefit of our TCS is that it is robust—except for when $k = 1, 2$, the standard deviation of TCS is often small.

This trend can be attributed to the PCA computation. When k is too small ($k = 1, 2$), the number of examples used to learn PCA coefficients may even be smaller than its dimensionality. Hence, the coordinate system defined by PCA is low-quality in such scenarios, which leads to relatively worse TCS results.

Notably, unlike previous studies that highlight sensitivity to learning rate variations in few-shot learning scenarios (Hu et al. 2022; Fu and Zhu 2024), our method demonstrates robust performance with a fixed learning rate, thus reducing the necessity for extensive hyperparameter adjustments.

Finally, as shown in Table 3, different pretraining methods for the teacher (MoCov3 or DINOv2) and different teachers (ViT-B or ViT-L) all lead to robust TCS results. A larger teacher with a better pretraining (ViT-L plus DINOv2) in fact lead to the best set of results.

Dataset	Pretraining Method	S	T	Top 1 (%)
ImageNet-1K	MoCov3	ViT-S	-	59.30 ± 0.27
			ViT-B	61.59 ± 0.12
		ViT-B	-	63.76 ± 0.25
			ViT-S	65.73 ± 0.05
CUB	MoCov3	ViT-S	-	65.58 ± 0.30
			ViT-B	71.80 ± 0.05
		ViT-B	-	67.32 ± 0.21
			ViT-S	71.17 ± 0.26
CUB	DINOv2	ViT-S	-	87.00 ± 0.20
			ViT-B	87.47 ± 0.09
		ViT-B	ViT-L	87.52 ± 0.06
			-	88.71 ± 0.19
		ViT-S	88.70 ± 0.22	

Table 4: Impact of model size and capacity gap. Top-1 accuracy (%) along with the standard deviation across three random subsets on ImageNet-1K and CUB are reported.

Ablation studies

We initially investigate the impact of model size discrepancies on teacher-student relationships and the effects of the coordinate system.

Impact of model size: Teacher vs. Student Substantial differences in model size between the teacher and student networks can often be counterproductive, i.e., large capacity gap is harmful (Mirzadeh et al. 2020) in existing KD methods. In contrast, our TCS shows resilience to these discrepancies and even demonstrates improvements when the teacher is small but the student is large. We hypothesize that so long as there is large enough difference between the teacher and student, TCS captures useful dark knowledge from the teacher and helps the student—because the student mainly inherits the coordinate system, not the details of the teacher network.

Detailed results are provided in Table 4 for pFSL, with experiments conducted on CUB 16-shot and ImageNet-1K 16-shot. As these results suggest, large capacity gap is not harmful to TCS: All teachers (small or large) are helpful, leading to accuracy higher than the baseline (whose teacher ‘T’ is marked as ‘-’). Furthermore, when the teacher is smaller (ViT-S) but student is larger (ViT-B), TCS even can achieve 2-4% accuracy increase.

Effect of the coordinate system The dark knowledge in TCS resides in the coordinate system computed from PCA. In this study, we replace it with D_t random orthogonal vectors (‘*random*’ in Table 5). We also study using different data to compute PCA. The first group (‘In-domain’) used training data from the target task with 3- or 16-shot or all the training data to compute PCA. ‘N/A’ means TCS is not used. The second group (‘Out-of-domain’) used either random PCA coefficients or computed PCA using 16-shot data *not* in the target task. For example, in the last row of Table 5, PCA was computed using the ‘Diabetic Retinopathy’ dataset although

	Data for PCA	Top 1 (%)
In-domain	N/A	65.30 ± 0.30
	3-shot	70.37 ± 0.11
	16-shot	71.80 ± 0.05
	all	71.90 ± 0.29
Out-of-domain	<i>random</i>	65.71 ± 0.61
	Flowers102	67.56 ± 0.59
	Pets	67.12 ± 0.66
	Diabetic Retinopathy	67.54 ± 0.43

Table 5: Results of computing the PCA coordinate system using different data. Experiments were on the CUB dataset for 16-shot pFSL. The three datasets used in the out-of-domain category all come from VTAB-1K (Zhai et al. 2019).

the target task is CUB.

As shown in Table 5, it is not surprising that random PCA coefficients is not helpful. But, even the coordinate system computed from out-of-domain data consistently improves the accuracy by roughly 2% compared to the baseline not using TCS. It is also not surprising that in-domain PCA is better than out-of-domain PCA by 3-4%. The more data used for computing the coordinate system, the higher the accuracy. These results again verifies that the coordinate system indeed contains essential dark knowledge, or, as we argued, “all you need is a tailored coordinate system” in KD.

Conclusion

We targeted the knowledge distillation (KD) task from a perspective very different from existing methods. In this paper, we argued that the dark knowledge in a teacher network can be captured by knowing its coordinate system (the linear subspace where the teacher’s features reside in), and by tailoring this coordinate system to the target task, the student can be effectively injected the dark knowledge. Hence our method is named TCS.

In TCS, the coordinate system is found by PCA, and tailored to the target task via a proposed iterative feature selection method. TCS was successfully applied in KD and practical few-shot learning. TCS is teacher-free and applies to diverse architectures, allows cross-architecture distillation with large capacity gap. It achieves significantly higher accuracy than state-of-the-art KD methods, while only requiring roughly half of their training time and GPU memory costs.

TCS has its limitations. In this paper, it was tested in KD (classification). Given the versatility of TCS, we plan to extend this method to distillation of object detection and segmentation models. Also, since TCS works well for teacher networks learned through self-supervised learning, it is also interesting and promising to extend TCS to compressing and accelerating large language models (LLMs) or large multi-modal models.

In the future, we are also interested in its theoretical foundation: how is the dark knowledge encoded by the coordinate system?

Acknowledgments

This research was partly supported by the National Natural Science Foundation of China under Grant 62276123.

References

- Bengio, Y.; Courville, A.; and Vincent, P. 2013. Representation Learning: A Review and New Perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8): 1798–1828.
- Chen, G.; Choi, W.; Yu, X.; Han, T.; and Chandraker, M. 2017. Learning efficient object detection models with knowledge distillation. In *NIPS*, 742–751.
- Chen, L.; Wang, D.; Gan, Z.; Liu, J.; Henao, R.; and Carin, L. 2021a. Wasserstein contrastive representation distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16296–16305.
- Chen, P.; Liu, S.; Zhao, H.; and Jia, J. 2021b. Distilling knowledge via knowledge review. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5008–5017.
- Chen, W.-Y.; Liu, Y.-C.; Kira, Z.; Wang, Y.-C. F.; and Huang, J.-B. 2019. A closer look at few-shot classification. In *International Conference on Learning Representations*, 1–16.
- Chen, X.; Xie, S.; and He, K. 2021. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9640–9649.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Fu, M.; and Zhu, K. 2024. Instance-based Max-margin for Practical Few-shot Recognition. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hao, Z.; Guo, J.; Han, K.; Tang, Y.; Hu, H.; Wang, Y.; and Xu, C. 2024. One-for-All: Bridge the Gap Between Heterogeneous Architectures in Knowledge Distillation. *Advances in Neural Information Processing Systems*, 36.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Heo, B.; Kim, J.; Yun, S.; Park, H.; Kwak, N.; and Choi, J. Y. 2019. A comprehensive overhaul of feature distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1921–1930.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Hospedales, T.; Antoniou, A.; Micaelli, P.; and Storkey, A. 2022. Meta-learning in neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9): 5149–5169.
- Hu, S. X.; Li, D.; Stühmer, J.; Kim, M.; and Hospedales, T. M. 2022. Pushing the limits of simple pipelines for few-shot learning: External data and fine-tuning make a difference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9068–9077.
- Huang, T.; You, S.; Wang, F.; Qian, C.; and Xu, C. 2022. Knowledge distillation from a stronger teacher. *Advances in Neural Information Processing Systems*, 35: 33716–33727.
- Huang, Z.; and Wang, N. 2017. Like What You Like: Knowledge Distill via Neuron Selectivity Transfer. *arXiv preprint arXiv:1707.01219*.
- Jin, X.; Peng, B.; Wu, Y.; Liu, Y.; Liu, J.; Liang, D.; Yan, J.; and Hu, X. 2019. Knowledge Distillation via Route Constrained Optimization. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 1345–1354.
- Kim, J.; Park, S.; and Kwak, N. 2018. Paraphrasing complex network: Network compression via factor transfer. *Advances in neural information processing systems*, 31.
- Kim, K.; Ji, B.; Yoon, D.; and Hwang, S. 2021. Self-Knowledge Distillation with Progressive Refinement of Targets. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 6547–6556.
- Koohpayegani, S. A.; Tejankar, A.; and Pirsiavash, H. 2020. Compress: Self-supervised learning by compressing representations. *Advances in Neural Information Processing Systems*, 33: 12980–12992.
- Krizhevsky, A. 2009. Learning multiple layers of features from tiny images. Technical report, University of Toronto.
- Li, X.; Wu, J.; Fang, H.; Liao, Y.; Wang, F.; and Qian, C. 2020. Local correlation consistency for knowledge distillation. In *European Conference on Computer Vision*, 18–33. Springer.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Mangla, P.; Singh, M.; Sinha, A.; Kumari, N.; Balasubramanian, V. N.; and Krishnamurthy, B. 2020. Charting the right manifold: Manifold Mixup for few-shot learning. In *IEEE Winter Conference on Applications of Computer Vision*, 2207–2216.
- Miles, R.; Elezi, I.; and Deng, J. 2024. vkD: Improving Knowledge Distillation using Orthogonal Projections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15720–15730.
- Mirzadeh, S. I.; Farajtabar, M.; Li, A.; Levine, N.; Matsukawa, A.; and Ghasemzadeh, H. 2020. Improved knowledge distillation via teacher assistant. In *Proceedings of the 34th AAAI conference on artificial intelligence*, 5191–5198.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H. V.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; Howes, R.; Huang, P.-Y.; Xu, H.; Sharma, V.; Li, S.-W.; Galuba, W.; Rabbat, M.; Assran, M.; Ballas, N.; Synnaeve, G.; Misra, I.; Jegou, H.; Mairal, J.; Labatut, P.; Joulin, A.; and Bojanowski, P. 2023. DINOv2: Learning robust visual features without supervision.

- Park, W.; Kim, D.; Lu, Y.; and Cho, M. 2019. Relational knowledge distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3967–3976.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115: 211–252.
- Schmitt, S.; Hudson, J.; Zidek, A.; and et. al. 2018. Kickstarting Deep Reinforcement Learning. *CoRR*, abs/1803.03835.
- Son, W.; Na, J.; Choi, J.; and Hwang, W. 2021. Densely Guided Knowledge Distillation using Multiple Teacher Assistants. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 9375–9384.
- Tian, Y.; Krishnan, D.; and Isola, P. 2019. Contrastive representation distillation. *arXiv preprint arXiv:1910.10699*.
- Tolstikhin, I. O.; Houlsby, N.; Kolesnikov, A.; Beyer, L.; Zhai, X.; Unterthiner, T.; Yung, J.; Steiner, A.; Keysers, D.; Uszkoreit, J.; et al. 2021. Mlp-mixer: An all-mlp architecture for vision. *Advances in neural information processing systems*, 34: 24261–24272.
- Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; and Jégou, H. 2021. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, 10347–10357. PMLR.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The Caltech-UCSD Birds-200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology.
- Wang, G.-H.; Ge, Y.; and Wu, J. 2021. Distilling knowledge by mimicking features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11): 8183–8195.
- Yang, Z.; Cui, Y.; Chen, Z.; Che, W.; Liu, T.; Wang, S.; and Hu, G. 2020. TextBrewer: An Open-Source Knowledge Distillation Toolkit for Natural Language Processing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 9–16. Association for Computational Linguistics.
- Yim, J.; Joo, D.; Bae, J.; and Kim, J. 2017. A Gift from Knowledge Distillation: Fast Optimization, Network Minimization and Transfer Learning. In *CVPR*, 7130–7138.
- Yu, H.; and Wu, J. 2023. Compressing Transformers: Features Are Low-Rank, but Weights Are Not! In *Proceedings of the AAAI Conference on Artificial Intelligence*, 11007–11015.
- Yu, R.; Li, A.; Morariu, V. I.; and Davis, L. 2017. Visual Relationship Detection With Internal and External Linguistic Knowledge Distillation. In *ICCV*.
- Yuan, L.; Tay, F. E.; Li, G.; Wang, T.; and Feng, J. 2020. Revisiting Knowledge Distillation via Label Smoothing Regularization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3902–3910.
- Yun, S.; Park, J.; Lee, K.; and Shin, J. 2020. Regularizing Class-Wise Predictions via Self-Knowledge Distillation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13873–13882.
- Zagoruyko, S.; and Komodakis, N. 2016. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*.
- Zhai, X.; Puigcerver, J.; Kolesnikov, A.; Ruysen, P.; Riquelme, C.; Lucic, M.; Djolonga, J.; Pinto, A. S.; Neumann, M.; Dosovitskiy, A.; et al. 2019. A large-scale study of representation learning with the Visual Task Adaptation Benchmark. *arXiv preprint arXiv:1910.04867*.
- Zhang, L.; Song, J.; Gao, A.; Chen, J.; Bao, C.; and Ma, K. 2019. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3713–3722.
- Zhao, B.; Cui, Q.; Song, R.; Qiu, Y.; and Liang, J. 2022. Decoupled knowledge distillation. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 11953–11962.
- Zhou, H.; Song, L.; Chen, J.; Zhou, Y.; Wang, G.; Yuan, J.; and Zhang, Q. 2021. Rethinking soft labels for knowledge distillation: A bias-variance tradeoff perspective. *arXiv preprint arXiv:2102.00650*.
- Zhu, K.; He, Y.-Y.; and Wu, J. 2023a. Coarse is better? a new pipeline towards self-supervised learning with uncurated images. *arXiv preprint arXiv:2306.04244*.
- Zhu, K.; He, Y.-Y.; and Wu, J. 2023b. Quantized feature distillation for network quantization. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence*, 11452–11460.
- Zhu, K.; Zhao, L.; Ge, Z.; and Zhang, X. 2024. Self-supervised visual preference alignment. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 291–300.
- Zhu, W.; Liu, J.; and Huang, Y. 2023. HNSSL: Hard Negative-Based Self-Supervised Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 4778–4787.