

Cross-Architecture Distillation Made Simple with Redundancy Suppression

Weijia Zhang Yuehao Liu Wu Ran Chao Ma*

MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

{weijia.zhang, yuehao.liu, bonjourlemonde, chaoma}@sjtu.edu.cn

Abstract

We describe a simple method for cross-architecture knowledge distillation, where the knowledge transfer is cast into a redundant information suppression formulation. Existing methods introduce sophisticated modules, architecture-tailored designs, and excessive parameters, which impair their efficiency and applicability. We propose to extract the architecture-agnostic knowledge in heterogeneous representations by reducing the redundant architecture-exclusive information. To this end, we present a simple redundancy suppression distillation (RSD) loss, which comprises cross-architecture invariance maximisation and feature decorrelation objectives. To prevent the student from entirely losing its architecture-specific capabilities, we further design a lightweight module that decouples the RSD objective from the student’s internal representations. Our method is devoid of the architecture-specific designs and complex operations in the pioneering method of OFA. It outperforms OFA on CIFAR-100 and ImageNet-1k benchmarks with only a fraction of their parameter overhead, which highlights its potential as a simple and strong baseline to the cross-architecture distillation community. The code and models are released at <https://github.com/VISION-SJTU/RSD>.

1. Introduction

Knowledge distillation (KD) aims to transfer the privileged capability of a pre-trained teacher model to a usually less capable student to improve its performance. Since its introduction by Hinton et al. [25] in 2015, KD has witnessed significant advancements in methodology and performance over a decade. Existing KD methods extract and transfer different kinds of knowledge, such as network outputs [25, 52, 66, 78], intermediate representations [6, 49, 53, 71], or higher-order correlations amongst them [21, 44, 45], and have demonstrated widespread success across different computer vision applications [10, 27, 62, 67–69, 73, 75, 80].

*Corresponding author.

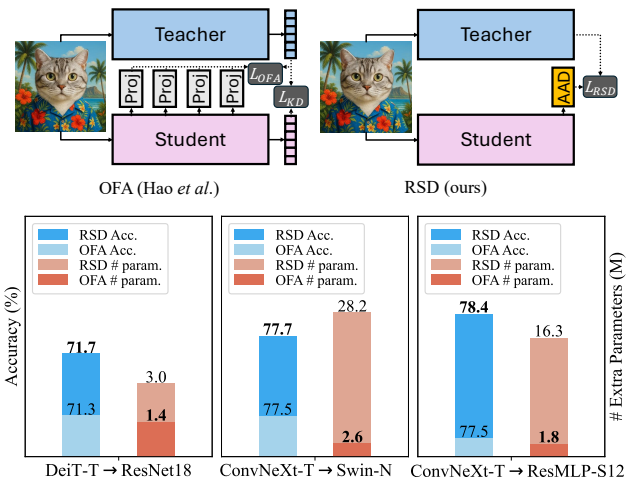


Figure 1. **A comparison of RSD and OFA.** Compared to the pioneering method of OFA, RSD is simpler in design, stronger in performance, and lower in complexity.

Yet, previous works mostly considered distillation between models of the same architectural type (e.g., CNNs [23]).

Recently, with the rise of novel vision architectures such as ViTs [13, 37] and MLPs [56, 58], it becomes increasingly relevant for the community to study knowledge distillation across distinct architectures, a task known as cross-architecture knowledge distillation (CAKD). Distilling across heterogeneous architectures finds broad practical relevance, particularly given that the best-performing models today are often not the most deployment-friendly ones. Nonetheless, contrary to a thriving literature on generic KD, cross-architecture KD remains a less explored setting, with OFA [22] being a pioneering method.

Compared to generic distillation, CAKD is more complicated, which is fundamentally due to the distinct properties of representations from different architectures. First, heterogeneous features have *different dimensionalities*, which requires additional operations to reconcile. They also possess *distinct, even conflicting patterns and characteristics* [18, 43, 48]. Forcing the student to blindly absorb heterogeneous information leads to degraded performance [36, 43]. As a result, generic KD methods show lim-

ited effectiveness or even collapse when directly applied to CAKD, as shown in Tables 1 and 2.

As a seminal method, OFA has to make complex, architecture-tailored designs to unify heterogeneous feature dimensionality, such as depth-wise separable convolution [9] for CNN features, and token operations and attention mechanisms [13] for ViT and MLP features. To circumvent the representation heterogeneity, OFA projects intermediate features into the logit space, which is considered architecture-agnostic. However, useful fine-grained knowledge is lost in this process. OFA attaches separate modules to each of the network stages to retain more information, which amounts to quadrupled parameter complexity. All these designs render OFA complicated and inefficient. For instance, when performing ConvNeXt-T-to-Swin-N distillation, the projectors introduced by OFA are about 3 times larger than the student itself in terms of the number of parameters (*i.e.*, 28.2M *v.s.* 9.6M), shown in Figure 3, making it costly and less practical in real-world scenarios.

In this paper, we present a new method for cross-architecture distillation. Rather than projecting features into the logit space, we learn to explicitly extract the architecture-agnostic knowledge. We achieve this via the proposed redundancy suppression distillation (RSD) objective, rooted in information maximisation [54, 55] and feature decorrelation [3, 11] principles. This allows the student to grasp the underlying architecture-independent knowledge, without being hampered by redundant, distracting patterns specific to the teacher architecture. We also insert a lightweight MLP module to decouple the student’s internal representations from the direct dictation of the RSD objective, which enables the student to preserve some beneficial capabilities unique to its own architecture. A high-level comparison of our method to OFA is made in Figure 1.

The advantages of our method are manifold. First and foremost, it only involves a simple loss function and a lightweight MLP-based gadget. Since RSD is applied to the penultimate-layer embeddings, our method is devoid of the complex and tailored adaptation modules for unifying heterogeneous features in OFA. The embeddings also offer a richer amount of knowledge than the logit space of OFA. RSD does not rely on memory banks [53], assistant networks [40], teacher weight reuse [5, 34], attention mechanisms [33, 36], adversarial training [36], or asymmetric transformation [36, 74]. Besides, it does not require access to the full-stage intermediate features [22, 49], and is shown to also work reasonably well on model logits, adding to its value in practical black-box distillation scenarios where security and privacy concerns are relevant.

In summary, this paper makes several significant contributions to the knowledge distillation literature. We provide a redundancy suppression perspective to cross-architecture distillation, which motivates us to explicitly get rid of the

redundant, teacher-specific information during the distillation process. We describe a simple redundancy suppression criterion based on information maximisation and feature decorrelation objectives. We show that moderate preservation of the student’s inherent characteristics further benefits cross-architecture distillation. We report substantial performance improvements over OFA [22] using a fraction of its parameter costs. RSD may potentially be a strong and efficient baseline for the community and inspire future endeavours to cross-architecture knowledge distillation.

2. Related Work

2.1. Vision architectures

Convolutional Neural Networks (CNNs) have been the cornerstone behind modern computer vision. They rely on convolution operations to capture local spatial patterns within images and pooling to reduce dimensionality. The shared weights of convolution equip CNNs with desirable properties such as inductive biases and translational invariance. CNNs are popularised by early pioneers such as AlexNet [12] and VGGNet [51]. ResNet [23] introduces residual connections to enable much deeper and more capable CNNs. Efficient CNNs [50, 77] are also engineered to facilitate their real-world applications. Recent models such as ConvNeXt [38] also incorporate design principles from Vision Transformers for better efficiency and performance.

Vision Transformers (ViTs) are an adaptation of Transformers [60] in natural language processing (NLP) to the vision domain. ViTs [13] treat images as sequences of patches and process them using self-attention mechanisms, which allows them to flexibly capture long-range dependencies and global context within the image. Following ViT’s initial success, Swin Transformer [37] introduces hierarchical window-based attention for improved training speed and performance, whereas DeiT [57] enables effective ViT training on smaller datasets via distillation.

Multi-Layer Perceptrons (MLPs) recently emerged as a competitive alternative to CNNs and ViTs in vision tasks. They convert images into flattened vectors and process them via fully-connected (FC) layers. MLP-Mixer [56] and ResMLP [58] introduced novel mixing operations, such as token-mixing and channel-mixing, to capture spatial dependencies without convolutions or self-attention. The simplicity, scalability, and efficiency of MLPs make them an appealing alternative to more complex models.

2.2. Knowledge distillation

Generic KD. First proposed in [25] as a model compression technique, knowledge distillation (KD) transfers the privileged capability of a usually more capable pretrained

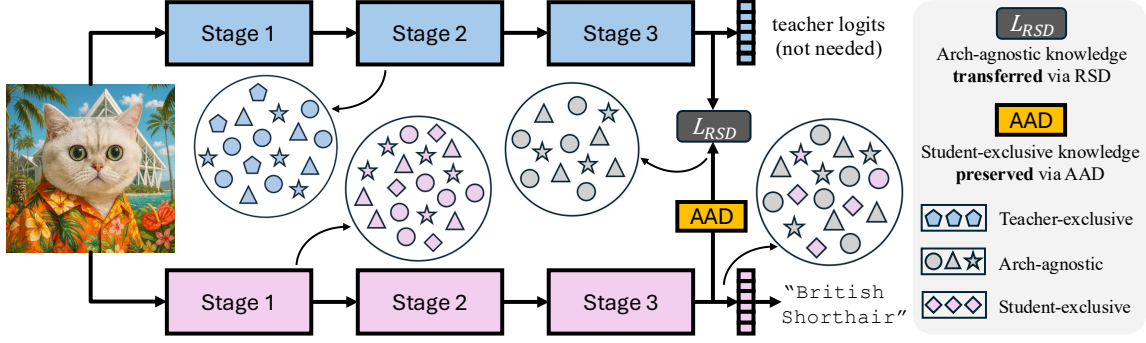


Figure 2. **A schematic diagram of RSD for cross-architecture knowledge distillation.** RSD employs a redundancy suppression objective (*i.e.*, RSD loss) to extract and transfer architecture-agnostic knowledge that is common to both teacher and student architectures. It uses a lightweight Architecture-Agnostic Knowledge Decoupler (AAD) module to align student representation dimension to the teacher’s, while decoupling the RSD dictation to allow the preservation of useful student-exclusive knowledge within the student.

teacher model to a lightweight student model. KD methods can be broadly categorised according to what sort of knowledge is extracted and transferred to the student. A group of methods [25, 32, 42, 52, 65, 66, 70, 74, 78, 79] match the student’s prediction logits to the teacher’s, and are called logits-based methods. By contrast, feature-based methods [4, 6, 20, 24, 33, 49, 53, 64, 71] let the student learn to mimic the intermediate features of the teacher. An orthogonal concept to the above methods is relational KD, where the student learns to replicate the higher-order correlations constructed from the teacher’s features [21, 35, 44, 45, 59] or logits [26, 76]. All these methods consider homogeneous architecture distillation, such as CNN to CNN [6, 25, 33, 49, 52, 53, 65, 70, 74, 78, 79] or ViT to ViT [21, 66, 71].

Cross-architecture KD. While a few works [7, 31, 57] made early attempts at unidirectional CNN-to-ViT distillation under basic settings, it is only until recently that the pioneering work of OFA [22] proposed a universal paradigm alongside rigorous evaluation settings for diverse cross-architecture distillation between CNNs, ViTs, and MLPs. OFA employs a set of projectors to transform intermediate features into the architecture-agnostic logit space, each using a target-knowledge enhancement mechanism. However, OFA introduces significant training overheads due to its complicated designs and mechanisms. In this work, we present a much simpler orthogonal approach that is able to outperform OFA with a fraction of its parameters.

2.3. Redundancy in representation learning

Previous research discovers that redundancy exists in the representations learnt by deep learning models, which motivates a group of works to mitigate the spatial [8, 30, 63] and channel [47] redundancy in CNN features through network or operator redesigns. Meanwhile, redundancy reduction is the core principle in early unsupervised feature learning

theories [1, 3, 11], which have also helped alleviate informational collapse in recent self-supervised learning algorithms [2, 17, 61, 72]. An orthogonal body of research is along domain generalisation [16, 39, 41, 46]. These works aim to learn domain-invariant representations for generalisation to unseen domains, which may be viewed as suppressing the redundant domain-specific information. Yet, these works substantially differ from ours in context and methodology. In this work, we make initial efforts to improve cross-architecture distillation from a representation redundancy perspective.

3. Method

3.1. Preliminaries

Generic KD. In knowledge distillation (KD), we are given a pre-trained teacher model and a to-be-trained student model. During distillation, the student is simultaneously supervised by the ground-truth labels and distillation signals from the teacher. For logit KD, the student is trained to mimic the output predictions given by the teacher, often via the Kullback–Leibler divergence (KLD) loss, and the optimisation objective is given by:

$$\mathcal{L}_{\text{logit-kd}} = \text{CE}(\mathbf{z}^s, y) + \lambda \text{KLD}(\mathbf{z}^s, \mathbf{z}^t), \quad (1)$$

where \mathbf{z}^s and \mathbf{z}^t are the student and teacher logits, respectively, y is the ground-truth label, and $\text{CE}(\cdot)$ is the cross-entropy loss. Whereas for feature KD, the student learns to match its intermediate features to the teacher’s from the corresponding layers. The objective is therefore:

$$\mathcal{L}_{\text{feature-kd}} = \text{CE}(\mathbf{z}^s, y) + \lambda \text{MSE}(\psi(\mathbf{f}^s), \mathbf{f}^t), \quad (2)$$

where the mean-square error (MSE) loss is often used to drive student features \mathbf{f}^s towards teacher features \mathbf{f}^t . A convolutional adaptor network $\psi(\cdot)$ is often used to match the student feature dimension to the teacher’s [49].

Algorithm 1 PyTorch-style pseudocode for RSD Loss

```
1: # feat_t: teacher features, feat_s: student features after AAD
2: # B: batch size, D: feature dimension
3: feat_t = Normalize(feat_t) # [B,D]
4: feat_s = Normalize(feat_s) # [B,D]
5: rcc = torch.mm(feat_t.T, feat_s) # [D,D]
6: idt = torch.eye(D) # [D,D]
7: loss_rsd = torch.mse(rcc, idt) # [D,D]
8: loss_rsd[(1 - idt).bool()] *=  $\kappa$  # [D,D]
9: return loss_rsd.mean()
```

OFA. As the first universal method for cross-architecture distillation between CNNs, ViTs, and MLPs, OFA [22] projects intermediate features of the student into a common logit space to be supervised by the teacher logits. Its training objective is given by:

$$\mathcal{L}_{\text{ofa}} = \text{CE}(\mathbf{z}^s, y) + \lambda_1 \sum_{i \in S} \text{OFA}(\psi_{\text{ofa},i}(\mathbf{f}_i^s), \mathbf{z}^t) + \lambda_2 \text{OFA}(\mathbf{z}^s, \mathbf{z}^t), \quad (3)$$

where $\text{OFA}(\cdot)$ is a variant of the KLD loss, and S is the set of stages whose features are used for projection ($|S| = 4$). Each stage i requires a separate projection network $\psi_{\text{ofa},i}$. Each ψ_{ofa} is much bulkier than the original ψ in [49], involving complex arch-specific operations such as depthwise separable convolution [9] or ViT blocks [13]. All these designs in OFA amount to its excessive complexity.

3.2. Redundancy Suppression Distillation

A redundancy suppression perspective to CAKD. In cross-architecture knowledge distillation, the teacher and student representations bear distinct, even conflicting characteristics due to their heterogeneous nature, as have been demonstrated in prior works [18, 19, 43, 48]. Meanwhile, when fed with the same training data, the representations produced by the teacher and the student, albeit very different, also encode the same semantic information corresponding to the very same input data. We expect such information to be fundamentally *architecture-independent* but only input-dependent (denoted by uncoloured circles, triangles, and stars in Figure 2). Yet, this underlying *architecture-agnostic knowledge* is blended with arch-specific patterns (navy and pink colours) or knowledge (pentagons and diamonds) in respective representations. Prior research has shown that blindly dictating these arch-specific patterns over heterogeneous architectures leads to incompatibility and degradation [15, 22, 36, 43]. Intuitively, we would like the student to grasp the arch-independent knowledge, such that it will not be distracted by any teacher-specific patterns when generalising to different heterogeneous teachers. Hence, give a pair of heterogeneous features, we may extract the commonality between them by suppressing the redundant architecture-specific or irrelevant information.

Extracting common knowledge between teacher and student. We leverage the correlations between feature units of heterogeneous representations to learn the architecture-agnostic representation, inspired by classic theories on unsupervised feature extraction [1, 3, 11] and Information Bottleneck (IB) principle [54, 55]. Specifically, we construct matrix $\mathbf{P} \in \mathbb{R}^{D \times D}$ of Pearson correlations between pairs of feature units of the teacher and student representations, *i.e.*, $\mathbf{z}^t, \mathbf{z}^s \in \mathbb{R}^{B \times D}$:

$$\mathbf{P}_{ij} = \frac{\sum_{k=1}^B (\mathbf{z}_{ki}^t - \bar{\mathbf{z}}_i^t)(\mathbf{z}_{kj}^s - \bar{\mathbf{z}}_j^s)}{\sqrt{\sum_{k=1}^B (\mathbf{z}_{ki}^t - \bar{\mathbf{z}}_i^t)^2 \sum_{k=1}^B (\mathbf{z}_{kj}^s - \bar{\mathbf{z}}_j^s)^2}} \quad \text{for } i, j \in \{1, \dots, D\}, \quad (4)$$

where B is the batch size and D is the embedding dimension. Essentially, \mathbf{P} captures the similarity in batch-wise activation patterns between pairs of feature dimensions of the teacher and the student. Each diagonal element in \mathbf{P} represents the correlation between the same feature dimension of teacher and student embeddings. To extract the “commonality” between the pair of heterogeneous representations, we learn \mathbf{z}^s that leads to maximised invariance between itself and \mathbf{z}^t for the same feature unit [3, 11, 61, 72]. This translates to driving the diagonal elements of \mathbf{P} to 1. Hence, the optimisation target for our task can be built via a target matrix $\mathbf{T} \in \mathbb{R}^{D \times D}$, whereby:

$$\mathbf{T}_{ij} = 1 \quad \text{if } i = j. \quad (5)$$

According to early unsupervised feature learning theories, the minimisation of redundancy in features corresponds to extracting statistically independent features [1, 3, 11]. This motivates us to further suppress the redundancy in learnt representation \mathbf{z}^s by decorrelating its feature units. For the cross-architecture distillation problem, we revamp this approach and instead minimise the mutual information between each student feature unit and every other teacher feature unit. This makes the redundancy suppressed within the heterogeneous representation space, which implicitly facilitates the extraction of architecture-agnostic information. Formally, this simply means forcing all off-diagonal elements in \mathbf{P} to zero by:

$$\mathbf{T}_{ij} = 0 \quad \text{if } i \neq j, \quad (6)$$

which suggests \mathbf{T} is an identity matrix, and we may use any distance measure $d(\cdot)$ to compute a loss between \mathbf{P} and \mathbf{T} . We empirically find that this decorrelation objective further improves the performance of our method most of the time. Note that in Equation 5 we have assumed identical teacher and student embedding dimensionality D , which is not necessarily the case. Our decoupler module will make the alignment for arbitrary teacher and student embeddings, which is described next.

Teacher	Student	From Scratch		Logits-based				Feature-based				
		T.	S.	KD	DKD	DIST	OFA	FitNets	CC	RKD	CRD	RSD
CNN-based students												
Swin-T	ResNet18	89.26	74.01	78.74	80.26	77.75	<u>80.54</u>	78.87	74.19	74.11	77.63	83.92
ViT-S	ResNet18	92.44	74.01	77.26	78.10	76.49	<u>80.15</u>	77.71	74.26	73.72	76.60	81.50
Mixer-B/16	ResNet18	87.62	74.01	77.79	78.67	76.36	<u>79.39</u>	77.15	74.26	73.75	76.42	81.85
Swin-T	MobileNetV2	89.26	73.68	74.68	71.07	72.89	<u>80.98</u>	74.28	71.19	69.00	79.80	83.68
ViT-S	MobileNetV2	92.44	73.68	72.77	69.80	72.54	<u>78.45</u>	73.54	70.67	68.46	78.14	81.68
Mixer-B/16	MobileNetV2	87.62	73.68	73.33	70.20	73.26	<u>78.78</u>	73.78	70.73	68.95	78.15	81.74
Transformer-based students												
ConvNeXt-T	DeiT-T	88.42	68.00	72.99	74.60	73.55	<u>75.76</u>	60.78	68.01	69.79	65.94	82.46
Mixer-B/16	DeiT-T	87.62	68.00	71.36	73.44	71.67	<u>73.90</u>	71.05	68.13	69.89	65.35	78.50
ConvNeXt-T	Swin-P	88.42	72.63	76.44	76.80	76.41	<u>78.32</u>	24.06	72.63	71.73	67.09	82.21
Mixer-B/16	Swin-P	87.62	72.63	75.93	76.39	75.85	<u>78.93</u>	75.20	73.32	70.82	67.03	81.28
MLP-based students												
ConvNeXt-T	ResMLP-S12	88.42	66.56	72.25	73.22	71.93	<u>81.22</u>	45.47	67.70	65.82	63.35	84.21
Swin-T	ResMLP-S12	89.26	66.56	71.89	72.82	11.05	<u>80.63</u>	63.12	68.37	64.66	61.72	82.67
Average Gain	-	-	-	+3.17	+3.16	-2.31	+7.47	-5.20	-0.33	-1.40	-0.02	+10.69

Table 1. **Cross-architecture distillation results on CIFAR-100.** Best results are in bold and second best results underlined.

Retaining exclusive knowledge of the student. Ideally, the architecture-agnostic representation learnt through the described redundancy suppression objective would suppress any teacher-exclusive or student-exclusive knowledge. Since different architectures have their unique [19, 48], sometimes opposite [43] characteristics and behaviours, we expect that by moderately preserving the student’s unique representation patterns, it may be able to retain some of its exclusive advantages. For instance, a CNN student may exhibit prominent activations to local flurry patterns [18] in a “British shorthair” cat image owing to the locality properties of convolution operations. Such ability is largely architecture-exclusive. It is not possessed by an inductive-bias-deficient ViT teacher [14, 48], and is therefore absent in their common knowledge. In this case, it would be beneficial to let the student retain this ability, instead of being entirely overridden by the architecture-agnostic knowledge.

To this end, we design a decoupler $h(\cdot)$ that buffers the student’s internal representation from arch-agnostic information extraction. It prevents the student embeddings from being entirely exposed to and dictated by the RSD objective. Our architecture-agnostic knowledge decoupling (AAD) module is simple: it consists of two FC layers, $h_e(\cdot)$ and $h_a(\cdot)$, joined by BatchNorm and GeLU activation. $h_e(\cdot)$ is an *expander* that transforms the student embedding to a higher dimensional space. $h_a(\cdot)$ is an *adaptor* that aligns the expanded embedding to the dimension of the teacher embedding. Unlike OFA which uses different projection modules for different architectures, AAD is a *one-size-fits-all module* and is *more parameter-efficient* (see Figure 3). After distillation, it is discarded and there is no

additional overhead introduced at inference.

To summarise, the role of our AAD module is twofold: 1) It aligns the student and teacher embedding dimension to enable the calculation of the RSD loss. 2) It decouples the student internal representation from the one used by the RSD objective for extracting arch-invariant information in order to preserve student-exclusive knowledge.

Full objective. Putting all designs together, we define the redundancy suppression distillation (RSD) loss as:

$$\mathcal{L}_{\text{RSD}} = d(\mathbf{P}(h(\mathbf{z}^s), \mathbf{z}^t), \mathbf{T}) \quad (7)$$

for which we use the MSE for $d(\cdot)$. The full optimisation objective for our framework is simply a combination of the CE loss and our RSD loss:

$$\mathcal{L} = \mathcal{L}_{\text{CE}} + \lambda \mathcal{L}_{\text{RSD}} \quad (8)$$

where λ is a balancing weight. The RSD loss can be implemented in about 8 lines in PyTorch. A pseudocode snippet for it is provided in Algorithm 1.

Discussion. We discuss why we choose the penultimate-layer embeddings for our RSD optimisation. The foremost reason is that we would like to avoid the whole complex operations involved when working with intermediate features, which is precisely what makes OFA complicated. As discussed in Sections 1 and 3.1, OFA had to introduce tailored operations to reconcile heterogeneous features into unified dimensionality, such as depth-wise separable convolution for CNN features, and attention blocks and token merging blocks (and other token operations) for ViT and MLP

Teacher	Student	From Scratch		Logits-based				Feature-based					RSD
		T.	S.	KD	DKD	DIST	OFA	FitNets	CC	RKD	CRD		
CNN-based students													
DeiT-T	ResNet18	72.19	69.75	70.22	69.39	70.64	<u>71.34</u>	70.44	69.77	69.47	69.25	71.70	
Swin-T	ResNet18	81.35	69.75	71.14	71.10	70.91	<u>71.85</u>	71.18	70.07	68.89	69.09	72.13	
Mixer-B/16	ResNet18	76.58	69.75	70.89	69.89	70.66	<u>71.38</u>	70.78	70.05	69.46	68.40	71.41	
DeiT-T	MobileNetV2	72.19	68.87	70.87	70.14	71.08	<u>71.39</u>	70.95	70.69	69.72	69.60	72.18	
Swin-T	MobileNetV2	81.35	68.87	72.05	71.71	71.76	<u>72.32</u>	71.75	70.69	67.52	69.58	72.36	
Mixer-B/16	MobileNetV2	76.55	68.87	<u>71.92</u>	70.93	71.74	72.12	71.59	70.79	69.86	68.89	71.90	
Transformer-based students													
ResNet50	DeiT-T	80.33	72.17	75.10	75.60*	75.13*	76.55*	75.84	72.56	72.06	68.53	<u>76.17*</u>	
ConvNeXt-T	DeiT-T	82.05	72.17	74.00	73.95	74.07	<u>74.41</u>	70.45	73.12	71.47	69.18	74.46	
Mixer-B/16	DeiT-T	76.55	72.17	74.16	72.82	74.22	74.46	<u>74.38</u>	72.82	72.24	68.23	74.26	
ResNet50	Swin-N	80.33	75.53	77.58	78.23*	77.95*	<u>78.64*</u>	78.33	76.05	75.90	73.90	78.78*	
ConvNeXt-T	Swin-N	82.05	75.53	77.15	77.00	77.25	<u>77.50</u>	74.81	75.79	75.48	74.15	77.70	
Mixer-B/16	Swin-N	76.55	75.53	76.26	75.03	76.54	76.63	76.17	75.81	75.52	73.38	<u>76.55</u>	
MLP-based students													
ResNet50	ResMLP-S12	80.33	76.65	77.41	78.23*	77.71*	78.53*	78.13	76.21	75.45	73.23	<u>78.32*</u>	
ConvNeXt-T	ResMLP-S12	82.05	76.65	76.84	77.23	77.24	<u>77.53</u>	74.69	75.79	75.28	73.57	78.41	
Swin-T	ResMLP-S12	81.35	76.65	76.67	76.99	77.25	<u>77.31</u>	76.48	76.15	75.10	73.40	77.61	
Average Gain	-	-	-	+1.55	+1.29	+1.68	+2.20	+1.14	+0.49	-0.37	-1.77	+2.34	

Table 2. **Cross-architecture distillation results on ImageNet-1k.** Best results are in bold and second best results underlined. * denotes results achieved by combining with FitNets [49] following OFA [22].

features. By contrast, the penultimate-layer representations are always 1-D and are neither feature-map-like nor token-like, saving us from all the architecture-specific operations and designs. In this sense, RSD is a *more universal cross-architectural distillation method* than OFA. These embeddings are also smaller compared to intermediate features, which are resource efficient. Obtained near the network output, they are also less arch-specific than features from earlier layers, which puts us in a better position to extract more arch-invariant information.

4. Experiments

4.1. Experimental setup

Implementations. We conduct cross-architecture distillation using different pairs of CNN, Transformer, and MLP models. We use ResNet [23], MobileNet [50], and ConvNeXt [38] for CNN models, ViT [13], Swin Transformer [37], and DeiT [57] for Transformer, and ResMLP [58] and MLP-Mixer [56] for MLP. Experiments are conducted on the CIFAR-100 [29] and ImageNet-1k [12] datasets, following the configurations of OFA [22]. More details are provided in the supplementary material.

Baselines. Under different setups, we compare our method against established generic KD methods (*i.e.*, KD [25], DKD [78], DIST [26], FitNets [49], CC [45], RKD [44], and CRD [53]) and, notably, the cross-architectural pioneer OFA [22].

4.2. CIFAR-100 results

Table 1 presents the evaluation results for a total of 12 cross-architecture teacher-student pairs on CIFAR-100. It can be seen that the proposed method consistently outperforms all prior knowledge distillation methods, including OFA [22]. It is noteworthy that over several distillation pairs, the margin RSD leads OFA by is several times that of OFA over prior arts. For example, RSD leads OFA by 3.38% for Swin-T-to-ResNet18 distillation, where OFA is only 0.28% higher than DKD [78]. For ViT-S-to-MobileNetV2 distillation, where OFA brings a mere 0.31% improvement over CRD [53], RSD achieves an impressive 3.23% advantage. The highest performance gain is observed on ConvNeXt-T-to-DeiT-T distillation, where RSD leads OFA by a surprising margin of 6.70% – nearly the gap between OFA and an un-distilled student. These results highlight the effectiveness of RSD, especially given that RSD is orthogonal to OFA and uses only a fraction of its parameters.

4.3. ImageNet-1k results

We evaluate our method on the large-scale ImageNet-1k dataset, using 15 heterogenous teacher-student pairs. As shown in Table 2, RSD achieves state-of-the-art performance on a majority of the teacher-student pairs. Notably, RSD yields particularly large gains on certain pairs. For instance, when distilling from ConvNeXt-T to ResMLP-S12, RSD is 0.88% higher than OFA, given that OFA only leads the earliest generic distillation method KD [25] by 0.69%.

Design	Swin-T→ ResNet18	ConvNeXt-T→ ResMLP-S12
Baseline	74.01	76.65
+ RSD-corr	80.65	83.40
+ RSD-decorr	83.92	84.21

Table 3. Effect of RSD.

Design	CIFAR-100 ViT-S→ ResMLP-S12	ImageNet-1k ConvNeXt-T→ Mixer-B/16
RSD	82.94	80.73
w/o AAD	82.26	79.93

Table 4. Effect of AAD.

Logit loss	Swin-T→ ResNet18	Mixer-B/16→ DeiT-T	ConvNeXt-T→ ResMLP-S12
w/ KD [25]	78.74	71.36	72.25
w/ DKD [78]	80.26	73.44	73.22
w/ OFA [22]	80.60	70.69	78.87
w/ RSD	83.23	77.22	81.15

Table 5. RSD as a strong logit distiller.

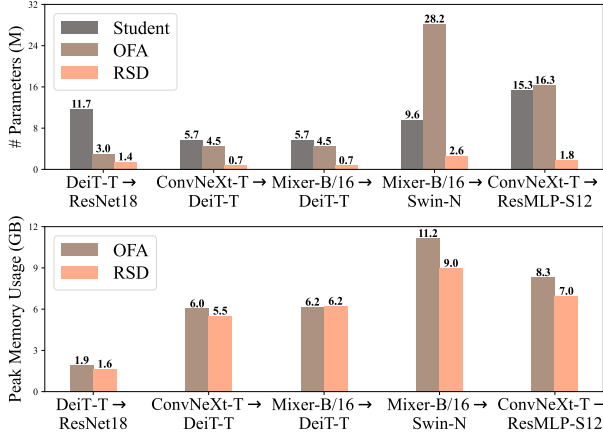


Figure 3. A comparison of the computational overheads.

For ConvNeXt-T-to-Swin-N distillation, RSD has a 0.20% advantage over OFA, when OFA is only 0.25% higher than DIST [26], by using only 9% as many extra parameters as OFA. RSD achieves an average gain of 2.34% across the 15 heterogeneous model pairs, highest amongst all methods.

4.4. Further analysis

Effect of proposed designs. We conduct ablation experiments to gauge the individual effectiveness of the proposed components in Tables 3 and 4. It is noteworthy when there is a dimensionality mismatch between the teacher and student embeddings, it is impossible to ablate the AAD block to assess the sole effect of the RSD objective. To enable such analysis, we set up teacher-student pairs with identical penultimate embedding dimension, such that the unadapted student representation may get directly exposed to the architecture-agnostic knowledge. The ablation results demonstrate the effectiveness of the core designs of redundancy suppression distillation and architecture-agnostic knowledge decoupling.

RSD as a strong logit distiller. While we have chosen to employ RSD between the informative feature embeddings from the teacher and the student, we further discovered that RSD is also a strong logit distiller. We conducted experiments to directly apply the proposed RSD loss over the logits. For head-to-head comparisons, we also apply the logit distillation losses from KD, DKD, and OFA independently. For OFA, we adopt the set of losses computed from

Logit loss	CIFAR-100 Swin-T→ ResNet18	ImageNet-1k ConvNeXt-T→ DeiT-T→ ResNet18
OFA	80.54	70.64 [†]
OFA+RSD	83.35	71.21

Table 6. Integrating the RSD objective into OFA.

Stage	Swin-T→ ResNet18	Mixer-B/16→ DeiT-T
\emptyset	83.92	76.45
{4}	84.39	76.83
{3,4}	84.81	76.37
{2,3,4}	84.63	77.17
{1,2,3,4}	85.45	76.90

Table 7. Integrating OFA into RSD at different network stages.

projected logits to ablate any impact of the logit KD loss of [25]. From Table 5, RSD’s logit variant itself (without AAD or any projectors) is a strong logit distillation objective. It outperforms individual existing logit distillation objectives with considerable margins, particularly OFA. This also reveals that OFA’s performance heavily relies on the logit KD loss (*i.e.*, 3rd term in Equation 3). By contrast, *RSD is free from such constraint and does not rely on an additional logit matching loss*. Noticeably, this one single RSD loss still surpasses the full formulations of OFA in Table 1 that encompass costly and tailored projector branches, adaptive target knowledge enhancer, and the KD loss.

Compatibility with OFA. We investigate whether our spirit of redundancy suppression is complementary to the method of OFA. To make RSD integrated with OFA, we simply replace all OFA losses in it with our RSD loss. This involves using the proposed RSD loss to replace OFA’s KD loss at logit level and the OFA losses at each network stage. Effectively, RSD now reduces the redundancy between full-stage intermediate features of the teacher and the student, using OFA’s projected logits as a proxy. The results are provided in Table 6, where [†] denotes the result reproduced by ourselves for a fair analysis. We find that the RSD objective can be seamlessly integrated with OFA and leads to substantial performance gains. In addition, we perform another set of experiments to extrapolate RSD to OFA’s multi-stage features. In Table 7, we start with our RSD baseline where no intermediate features are utilised, denoted by \emptyset (*i.e.*, only

	T.: ResNet34 S.: ResNet18	T.: ResNet50 S.: MobileNetV1
KD [25]	70.66	68.58
OFD [24]	70.81	71.25
CRD [53]	71.17	71.37
RKD [44]	71.34	71.32
CAT-KD [20]	71.26	72.24
SimKD [5]	71.59	72.25
ReviewKD [6]	71.61	72.56
DKD [78]	71.70	72.05
SDD [65]	71.14	72.24
DIST [26]	72.07	73.24
OFA [22]	72.10	-
RSD	72.18	<u>73.08</u>

Table 8. **Same-architecture distillation results on ImageNet-1K.** The best result is in bold and second best underlined.

penultimate-layer embeddings are used). We then gradually incorporate intermediate features from various network stages, and apply our RSD objective over OFA-projected pseudo logits. As can be seen, RSD is compatible and generally benefited by additional intermediate features.

Computational cost. We compare the computational costs of our method to OFA in terms of the number of extra parameters introduced and peak GPU memory usage on ImageNet-1k. The measurements are visualised in Figure 3 for multiple teacher-student pairs. Where OFA introduces an amount of extra network parameters on par with or more than that of the student itself, *our method saves up to 10× in parameter count to achieve better performance*. In terms of peak GPU memory usage, RSD is also better off on average. RSD’s overhead benefit persists on the CIFAR-100 dataset, yet with even larger performance gains.

Same-Architecture Evaluation. Following OFA, we also perform evaluation under the same-architecture distillation set-up. We additionally compare with established generic KD methods OFD [24], ReviewKD [6], CAT-KD [20], SimKD [5], and the recent SDD [65]. As shown in Table 8, our method delivers competitive or even better results compared to OFA and leading same-architecture distillation methods, including SDD, DKD, and DIST. We notice that the advantage of RSD is more prominent on cross-architecture distillation. This observation reinforces our initial motivation, where our designs are rooted in theories for learning architecture-invariant knowledge from heterogeneous representations. Overall, the same-architecture distillation results highlight the significance of our findings to a broader knowledge distillation community.

Visualisation of cross-architectural feature similarity. We employ centered kernel alignment (CKA) [28] to assess the effectiveness of the proposed method. CKA is a similarity metric that can accommodate inputs of different dimensions. It enables analysis of cross-architectural feature

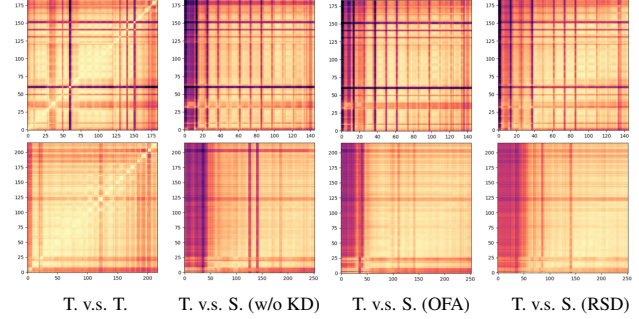


Figure 4. **Cross-architectural feature similarities measured by CKA.** Brighter colours indicate higher similarity and vice versa. Top: ConvNeXt-T teacher and ResMLP-S12 student. Bottom: ViT-S teacher and MobileNetV2 student.

similarity between different layers and stages in heterogeneous network pairs. Figure 4 visualises the CKA scores between layers of different teacher-student pairs. We observe that under the first set-up, RSD significantly increases the representation similarity between heterogeneous architectures at various stages. Whereas under the second set-up, some shallow feature dissimilarities in OFA are slightly alleviated, while some others are enlarged. This is justified because unlike OFA, our method does not directly access or process those intermediate features. Overall, RSD notably increases feature similarity at middle and deep layers.

Limitations & future work. Despite its simple formulation, RSD’s performance can sometimes be sensitive to hyperparameters λ and κ . Compared to CIFAR-100, its advantages are less prominent on larger-scale datasets such as ImageNet-1k, which is in part due to the unused 2-D features. RSD in its current design, albeit efficient and universal, only utilises the 1-D embeddings. While this offers benefits in efficiency and design, it fails to harness the rich spatial context associated with the 2-D feature maps, and is therefore unable to be directly extended to tasks that heavily demand spatial information such as object detection. We leave the extension to such tasks for future investigation.

5. Conclusion

We introduced RSD, a simple approach for cross-architecture knowledge distillation based on redundant knowledge suppression. RSD employs invariance maximisation and feature decorrelation objectives to extract architecture-agnostic knowledge common to heterogeneous architectures. It also allows student-exclusive patterns to be retained rather than entirely overridden by RSD through a lightweight decoupling module. RSD achieves superior performance over the recent cross-architectural baseline of OFA with a fraction of its parameter overhead, while avoiding architecture-tailored operations in OFA.

Acknowledgements. This work was supported by NSFC (62322113, 62376156), Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), and the Fundamental Research Funds for the Central Universities.

References

- [1] Joseph J Atick and A Norman Redlich. Towards a theory of early visual processing. *Neural computation*, 2(3):308–320, 1990. 3, 4
- [2] Adrien Bardes, Jean Ponce, and Yann LeCun. Vi-creg: Variance-invariance-covariance regularization for self-supervised learning. *ICLR*, 2022. 3
- [3] Horace B Barlow. Unsupervised learning. *Neural computation*, 1(3):295–311, 1989. 2, 3, 4
- [4] Defang Chen, Jian-Ping Mei, Yuan Zhang, Can Wang, Zhe Wang, Yan Feng, and Chun Chen. Cross-layer distillation with semantic calibration. In *AAAI*, 2021. 3
- [5] Defang Chen, Jian-Ping Mei, Hailin Zhang, Can Wang, Yan Feng, and Chun Chen. Knowledge distillation with the reused teacher classifier. In *CVPR*, 2022. 2, 8
- [6] Pengguang Chen, Shu Liu, Hengshuang Zhao, and Jiaya Jia. Distilling knowledge via knowledge review. In *CVPR*, 2021. 1, 3, 8
- [7] Xianing Chen, Qiong Cao, Yujie Zhong, Jing Zhang, Shenghua Gao, and Dacheng Tao. Dearth: data-efficient early knowledge distillation for vision transformers. In *CVPR*, 2022. 3
- [8] Yunpeng Chen, Haoqi Fan, Bing Xu, Zhicheng Yan, Yanis Kalantidis, Marcus Rohrbach, Shuicheng Yan, and Jiashi Feng. Drop an octave: Reducing spatial redundancy in convolutional neural networks with octave convolution. In *ICCV*, 2019. 3
- [9] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *CVPR*, 2017. 2, 4
- [10] Kaiwen Cui, Yingchen Yu, Fangneng Zhan, Shengcai Liao, Shijian Lu, and Eric P Xing. Kd-dlgan: Data limited image generation via knowledge distillation. In *CVPR*, 2023. 1
- [11] Gustavo Deco and L Parra. Non-linear feature extraction by redundancy reduction in an unsupervised stochastic neural network. *Neural networks*, 10(4):683–691, 1997. 2, 3, 4
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 2, 6, 1
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 1, 2, 4, 6
- [14] Stéphane d’Ascoli, Hugo Touvron, Matthew L Leavitt, Ari S Morcos, Giulio Biroli, and Levent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. In *ICML*, 2021. 5
- [15] Jiawei Fan, Chao Li, Xiaolong Liu, and Anbang Yao. Scalekd: Strong vision transformers could be excellent teachers. *NeurIPS*, 2024. 4
- [16] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 17(59):1–35, 2016. 3
- [17] Quentin Garrido, Yubei Chen, Adrien Bardes, Laurent Najman, and Yann Lecun. On the duality between contrastive and non-contrastive self-supervised learning. *ICLR*, 2023. 3, 2
- [18] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *ICLR*, 2018. 1, 4, 5
- [19] Jianyuan Guo, Kai Han, Han Wu, Yehui Tang, Xinghao Chen, Yunhe Wang, and Chang Xu. Cmt: Convolutional neural networks meet vision transformers. In *CVPR*, 2022. 4, 5
- [20] Ziyao Guo, Haonan Yan, Hui Li, and Xiaodong Lin. Class attention transfer based knowledge distillation. In *CVPR*, 2023. 3, 8
- [21] Zhiwei Hao, Jianyuan Guo, Ding Jia, Kai Han, Yehui Tang, Chao Zhang, Han Hu, and Yunhe Wang. Learning efficient vision transformers via fine-grained manifold distillation. *NeurIPS*, 2022. 1, 3
- [22] Zhiwei Hao, Jianyuan Guo, Kai Han, Yehui Tang, Han Hu, Yunhe Wang, and Chang Xu. One-for-all: Bridge the gap between heterogeneous architectures in knowledge distillation. *NeurIPS*, 2024. 1, 2, 3, 4, 6, 7, 8
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 2, 6
- [24] Byeongho Heo, Jeessoo Kim, Sangdoo Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *ICCV*, 2019. 3, 8
- [25] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *arXiv:1503.02531*, 2015. 1, 2, 3, 6, 7, 8
- [26] Tao Huang, Shan You, Fei Wang, Chen Qian, and Chang Xu. Knowledge distillation from a stronger teacher. In *NeurIPS*, 2022. 3, 6, 7, 8
- [27] Yuge Huang, Jiaxiang Wu, Xingkun Xu, and Shouhong Ding. Evaluation-oriented knowledge distillation for deep face recognition. In *CVPR*, 2022. 1
- [28] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *ICML*, 2019. 8
- [29] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009. 6, 1
- [30] Jiafeng Li, Ying Wen, and Lianghua He. Seconv: Spatial and channel reconstruction convolution for feature redundancy. In *CVPR*, 2023. 3
- [31] Kehan Li, Runyi Yu, Zhennan Wang, Li Yuan, Guoli Song, and Jie Chen. Locality guidance for improving vision transformers on tiny datasets. In *ECCV*, 2022. 3
- [32] Zheng Li, Xiang Li, Lingfeng Yang, Borui Zhao, Renjie Song, Lei Luo, Jun Li, and Jian Yang. Curriculum temperature for knowledge distillation. In *AAAI*, 2023. 3

- [33] Sihao Lin, Hongwei Xie, Bing Wang, Kaicheng Yu, Xiaojun Chang, Xiaodan Liang, and Gang Wang. Knowledge distillation via the target-aware transformer. In *CVPR*, 2022. 2, 3
- [34] Dongyang Liu, Meina Kan, Shiguang Shan, and Xilin Chen. Function-consistent feature distillation. In *ICLR*, 2023. 2
- [35] Li Liu, Qingle Huang, Sihao Lin, Hongwei Xie, Bing Wang, Xiaojun Chang, and Xiaodan Liang. Exploring inter-channel correlation for diversity-preserved knowledge distillation. In *ICCV*, 2021. 3
- [36] Yufan Liu, Jiajiong Cao, Bing Li, Weiming Hu, Jingting Ding, and Liang Li. Cross-architecture knowledge distillation. In *ACCV*, 2022. 1, 2, 4
- [37] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 1, 2, 6
- [38] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *CVPR*, 2022. 2, 6
- [39] Wang Lu, Jindong Wang, Haoliang Li, Yiqiang Chen, and Xing Xie. Domain-invariant feature exploration for domain generalization. *TMLR*, 2022. 3
- [40] Seyed-Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. Improved knowledge distillation via teacher assistant. In *AAAI*, 2020. 2
- [41] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *ICML*, 2013. 3
- [42] Amin Parchami-Araghi, Moritz Böhle, Sukrut Rao, and Bernt Schiele. Good teachers explain: Explanation-enhanced knowledge distillation. In *ECCV*, 2024. 3
- [43] Namuk Park and Songkuk Kim. How do vision transformers work? *ICLR*, 2022. 1, 4, 5
- [44] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *IEEE TMM*, 2019. 1, 3, 6, 8
- [45] Baoyun Peng, Xiao Jin, Jiaheng Liu, Shunfeng Zhou, Yichao Wu, Yu Liu, Dongsheng Li, and Zhaoning Zhang. Correlation congruence for knowledge distillation. In *CVPR*, 2019. 1, 3, 6
- [46] Vihari Piratla, Praneeth Netrapalli, and Sunita Sarawagi. Efficient domain generalization via common-specific low-rank decomposition. In *ICML*, 2020. 3
- [47] Jiaxiong Qiu, Cai Chen, Shuaicheng Liu, Heng-Yu Zhang, and Bing Zeng. Slimconv: Reducing channel redundancy in convolutional neural networks by features recombining. *IEEE TIP*, 30:6434–6445, 2021. 3
- [48] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? *NeurIPS*, 2021. 1, 4, 5
- [49] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *ICLR*, 2015. 1, 2, 3, 4, 6
- [50] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018. 2, 6
- [51] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 2
- [52] Shangquan Sun, Wenqi Ren, Jingzhi Li, Rui Wang, Rui Wang, and Xiaochun Cao. Logit standardization in knowledge distillation. In *CVPR*, 2024. 1, 3
- [53] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *ICLR*, 2020. 1, 2, 3, 6, 8
- [54] Naftali Tishby and Noga Zaslavsky. Deep learning and the information bottleneck principle. In *IEEE Information Theory Workshop*, 2015. 2, 4
- [55] Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000. 2, 4
- [56] Ilya O Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, et al. Mlp-mixer: An all-mlp architecture for vision. *NeurIPS*, 2021. 1, 2, 6
- [57] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers & distillation through attention. In *ICML*, 2021. 2, 3, 6
- [58] Hugo Touvron, Piotr Bojanowski, Mathilde Caron, Matthieu Cord, Alaaeldin El-Nouby, Edouard Grave, Gautier Izacard, Armand Joulin, Gabriel Synnaeve, Jakob Verbeek, et al. Resmlp: Feedforward networks for image classification with data-efficient training. *IEEE TPAMI*, 45(4):5314–5321, 2022. 1, 2, 6
- [59] Frederick Tung and Greg Mori. Similarity-preserving knowledge distillation. In *ICCV*, 2019. 3
- [60] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 2
- [61] Runqian Wang and Kaiming He. Diffuse and disperse: Image generation with representation regularization. *arXiv:2506.09027*, 2025. 3, 4
- [62] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. In *CVPR*, 2019. 1
- [63] Yulin Wang, Kangchen Lv, Rui Huang, Shiji Song, Le Yang, and Gao Huang. Glance and focus: a dynamic approach to reducing spatial redundancy in image classification. *NeurIPS*, 2020. 3
- [64] Yuzhu Wang, Lechao Cheng, Manni Duan, Yongheng Wang, Zunlei Feng, and Shu Kong. Improving knowledge distillation via regularizing feature norm and direction. In *ECCV*, 2024. 3
- [65] Shicai Wei, Chunbo Luo, and Yang Luo. Scale decoupled distillation. In *CVPR*, 2024. 3, 8
- [66] Kan Wu, Jinnian Zhang, Houwen Peng, Mengchen Liu, Bin Xiao, Jianlong Fu, and Lu Yuan. Tinyvit: Fast pretraining distillation for small vision transformers. In *ECCV*. Springer, 2022. 1, 3

- [67] Jiaxin Xie, Hao Ouyang, Jingtian Piao, Chenyang Lei, and Qifeng Chen. High-fidelity 3d gan inversion by pseudo-multi-view optimization. In *CVPR*, 2023. [1](#)
- [68] Chuanguang Yang, Helong Zhou, Zhulin An, Xue Jiang, Yongjun Xu, and Qian Zhang. Cross-image relational knowledge distillation for semantic segmentation. In *CVPR*, 2022.
- [69] Zhendong Yang, Zhe Li, Xiaohu Jiang, Yuan Gong, Zehuan Yuan, Danpei Zhao, and Chun Yuan. Focal and global knowledge distillation for detectors. In *CVPR*, 2022. [1](#)
- [70] Zhendong Yang, Ailing Zeng, Zhe Li, Tianke Zhang, Chun Yuan, and Yu Li. From knowledge distillation to self-knowledge distillation: A unified approach with normalized loss and customized soft labels. In *ICCV*, 2023. [3](#)
- [71] Zhendong Yang, Zhe Li, Ailing Zeng, Zexian Li, Chun Yuan, and Yu Li. Vitkd: Feature-based knowledge distillation for vision transformers. In *CVPR Workshop*, 2024. [1](#), [3](#)
- [72] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *ICML*, 2021. [3](#), [4](#), [1](#)
- [73] Bo Zhang, Jiacheng Sui, and Li Niu. Foreground object search by distilling composite image feature. In *ICCV*, 2023. [1](#)
- [74] Weijia Zhang, Dongnan Liu, Weidong Cai, and Chao Ma. Cross-view consistency regularisation for knowledge distillation. In *ACM MM*, 2024. [2](#), [3](#)
- [75] Weijia Zhang, Dongnan Liu, Chao Ma, and Weidong Cai. Alleviating foreground sparsity for semi-supervised monocular 3d object detection. In *WACV*, 2024. [1](#)
- [76] Weijia Zhang, Fei Xie, Weidong Cai, and Chao Ma. Vrm: Knowledge distillation via virtual relation matching. *arXiv:2502.20760*, 2025. [3](#)
- [77] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*, 2018. [2](#)
- [78] Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *CVPR*, 2022. [1](#), [3](#), [6](#), [7](#), [8](#)
- [79] Kaixiang Zheng and En-Hui Yang. Knowledge distillation based on transformed teaching matching. In *ICLR*, 2024. [3](#)
- [80] Shengchao Zhou, Weizhou Liu, Chen Hu, Shuchang Zhou, and Chao Ma. Unidistill: A universal cross-modality knowledge distillation framework for 3d object detection in bird’s-eye view. In *CVPR*, 2023. [1](#)