

Contrastive Deep Supervision

Linfeng Zhang¹, Xin Chen², Junbo Zhang¹, Runpei Dong³, Kaisheng Ma^{1*}

Tsinghua University¹, Intel Corporation², Xi'an Jiaotong University³

Abstract. The success of deep learning is usually accompanied by the growth in neural network depth. However, the traditional training method only supervises the neural network at its last layer and propagates the supervision layer-by-layer, which leads to hardship in optimizing the intermediate layers. Recently, deep supervision has been proposed to add auxiliary classifiers to the intermediate layers of deep neural networks. By optimizing these auxiliary classifiers with the supervised task loss, the supervision can be applied to the shallow layers directly. However, deep supervision conflicts with the well-known observation that the shallow layers learn low-level features instead of task-biased high-level semantic features. To address this issue, this paper proposes a novel training framework named Contrastive Deep Supervision, which supervises the intermediate layers with augmentation-based contrastive learning. Experimental results on nine popular datasets with eleven models demonstrate its effects on general image classification, fine-grained image classification and object detection in supervised learning, semi-supervised learning and knowledge distillation. Codes have been released in [Github](#).

1 Introduction

Along with the growth in large-scale datasets and computation resources, deep neural networks have become the most dominant models for various tasks [14,52]. However, the increasing depth of neural networks also introduces challenges in their training process. Traditional supervised training method only applies the supervision to the last layer and then propagates the error from the last layer to the shallow layers (Figure 1(a)), which leads to hardship in optimizing the intermediate layers such as gradient vanishing [29].

Recently, deep supervision (*a.k.a. deeply-supervised net*) has been proposed to address this issue by optimizing the intermediate layers directly [38]. As shown in Figure 1(b), deep supervision adds several auxiliary classifiers to the intermediate layers in different depths. During the training phase, these classifiers are optimized with the original final classifier together by the same training loss (*e.g.* cross entropy for classification tasks). Both experimental and theoretical analyses have demonstrated its effectiveness in facilitating model convergence [62].

However, success comes with remaining obstacles. In general, different layers in convolutional neural networks tend to learn features at different levels. Usually, the shallow layers learn low-level features such as colors and edges, while the

* Corresponding author, kaisheng@mail.tsinghua.edu.cn

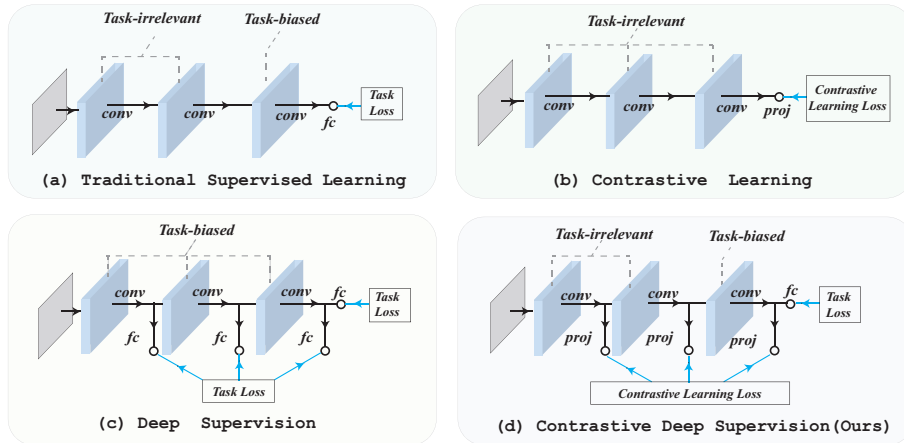


Fig. 1. The overview of the four methods. “ \rightarrow ” and “ \leftarrow ” indicate the path of forward computation and gradients backward computation. “proj” and “fc” indicate the projection heads and the fully connected classifiers, respectively. The gray dash line indicates whether the feature is task-irrelevant or task-biased. (a) Traditional supervised learning only applies supervision to the last layer and propagates it to the previous layers, leading to gradient vanishing. (c) Deep supervision trains both the last layer and the intermediate layers directly, which addresses gradient vanishing but makes all the layers be biased to the task. (d) Our method introduces contrastive learning to supervise the intermediate layer and thus avoid these problems.

last several layers learn more high-level task-related semantic features such as categorical knowledge for classification tasks [82]. However, deep supervision forces the shallow layers to learn the task-related knowledge, which disobeys the original feature extraction process in neural networks. As pointed out in MSDNet [28], this conflict sometimes leads to accuracy degradation in the final classifier. This observation indicates that the supervised task loss is probably not the best supervision for optimizing the intermediate layers.

In this paper, we argue that *contrastive learning can provide better supervision for intermediate layers than the supervised task loss*. Contrastive learning is one of the most popular and effective techniques in representation learning [7,8,34]. Usually, it regards two augmentations from the same image as a positive pair and different images as negative pairs. In the training period, the neural network is trained to minimize the distance of a positive pair while maximizing the distance of a negative pair. As a result, the network can learn the invariance to various data augmentation, such as **Color Jitter** and **Random Gray Scale**. Considering that these data augmentation invariances are usually low-level, task-irrelevant and transferable to various vision tasks [3,64], we argue that they are more beneficial knowledge to be learned by intermediate layers.

Motivated by these observations, we propose a novel training framework named *Contrastive Deep Supervision*. It optimizes the intermediate layers with

contrastive learning instead of traditional supervised learning. As shown in Figure 1(d), several projection heads are attached in the intermediate layers of the neural networks and trained to perform contrastive learning. These projection heads can be discarded in the inference period to avoid additional computation and storage. Different from deep supervision which trains the intermediate layers to learn the knowledge for a specific task, the intermediate layers in our method are trained to learn the invariance to data augmentation, which makes the neural network generalize better. Besides, since contrastive learning can be performed on unlabeled data, the proposed contrastive deep supervision can also be easily extended in the semi-supervised learning paradigm.

Moreover, contrastive deep supervision can be further utilized to boost the performance of another deep learning technique – knowledge distillation. Knowledge distillation (KD) is a popular model compression approach which aims to transfer the knowledge from a cumbersome teacher model to a lightweight student model [2,23,15]. Recently, abundant research finds that distilling the “crucial knowledge” inside the backbone features such as attention and relation [72,49,58] leads to better performance than directly distilling all the backbone features. In this paper, we show that the data augmentation invariances learned by the intermediate layers in contrastive deep supervision are more beneficial knowledge to be distilled. By combining contrastive deep supervision with the naïve feature distillation, the distilled ResNet18 achieves 73.23% accuracy on ImageNet, which outperforms the baseline and the second-best KD method by 4.02% and 2.16%, respectively.

Extensive experiments on nine datasets with eleven neural networks methods have been conducted to evaluate its effectiveness on general image classification, fine-grained image classification, object detection in supervised learning, semi-supervised learning and knowledge distillation, which demonstrates that contrastive deep supervision enables neural networks to learn better visual representation. In the discussion section, we further explain the effectiveness of our method from the perspective of regularization methods, which prevents models from overfitting and leads to better uncertainty estimation. To sum up, the main contributions of our paper can be summarized as follows.

- We propose *contrastive deep supervision*, a neural network training method in which the intermediate layers are directly optimized with contrastive learning. It enables neural networks to learn better visual representation at no expense of additional parameters and computation during inference.
- From the perspective of deep supervision, this paper firstly shows that the intermediate layers can be trained with supervision besides the task loss.
- From the perspective of representation learning, we firstly show that contrastive learning and supervised learning can be combined in a one-stage deep-supervision manner instead of the two-stage “pretrain-finetune” scheme.
- Extensive experiments on nine datasets, eleven neural networks with eleven comparison methods demonstrate the effectiveness of our method on general classification, fine-grained classification and object detection in supervised learning, semi-supervised learning and knowledge distillation.

2 Related Work

2.1 Deep Supervision

Deep neural networks usually contain a large number of layers, which increases the difficulty of optimization. To address this issue, deeply supervised net (*a.k.a.* deep supervision) is proposed to directly supervise the intermediate layers of deep neural networks [38]. Wang *et al.* show that deep supervision can alleviate the vanishing gradient problem and thus leads to significant performance improvements [62]. Usually, deep supervision attaches several auxiliary classifiers at the intermediate layers and supervises these auxiliary classifiers with the task loss (*e.g.* cross-entropy loss in classification). Recently, several methods have been proposed to improve deep supervision with knowledge distillation, which aims to minimize the difference between the prediction of the deepest classifier and the auxiliary classifiers in the intermediate layers [55,40]. Besides classification, abundant research has also demonstrated the effectiveness of deep supervision methods in dynamic neural networks [78], semantic segmentation [81,73,51], object detection [39], knowledge distillation [76] and so on.

2.2 Contrastive Learning

In the last several years, contrastive learning has become the most popular method in representation learning [74,63,32,27,68,18,60,5,24,61]. Oord *et al.* propose the contrastive predictive coding, which aims to predict the low dimension embedding of future signals with an auto-regressive model [47]. He *et al.* propose MoCo, which introduces a dynamic memory bank to record the embeddings of negative samples [19,9,11]. Then, SimCLR is proposed to show the importance of large batch size and long training time in contrastive learning [7,8]. Recently, abundant research has been proposed to study the influence of negative samples further. BYOL is introduced to demonstrate that contrastive learning is effective even without negative samples [16]. SimSiam gives a detailed study on the importance of batch normalization, negative samples, memory bank, and the stop-gradient operation [10]. Besides self-supervised learning, contrastive learning has also shown its power in the traditional supervised learning paradigm. Khosla *et al.* show that state-of-the-art performance can be achieved on ImageNet with the basic contrastive learning in SimCLR by building the positive pairs with label supervision [34,6]. Park *et al.* apply contrastive learning to unpaired image-to-image translation, which breaks the limitation of cycle reconstruction [48].

2.3 Knowledge Distillation

Knowledge distillation, which aims to facilitate the training of a lightweight student model under the supervision of an over-parameterized teacher model, has become one of the most popular methods in model compression. Knowledge distillation is first proposed by Bucilua *et al.* [2] and then expanded by Hinton *et al.* [23], who introduces a temperature-characterized softmax to soften

the distribution of teacher logits. Instead of distilling the knowledge of the logits, more and more techniques are proposed to distill the information in teacher features or its variants, such as attention maps [72,42], negative values [22], task-oriented information [76], relational information [49,58,43], Gram matrix [69], mutual information [1], context information [75] and so on. Besides model compression, knowledge distillation has also achieved significant success in self-supervised learning [30,46], semi-supervised learning [37,56], multi-exit neural network [78,77,70], incremental learning [83] and model robustness [65,79]

3 Methodology

3.1 Deep Supervision

In this subsection, we revisit the formulation of deep supervision methods. Let c be a given backbone classifier, deep supervision introduces several shallow classifiers by using the intermediate features in c . More specifically, assume $c = g \circ f$ where g is the final classifier, f is the feature extractor operator and $f = f_K \circ f_{K-1} \circ \dots \circ f_1$. K denotes the number of convolutional stages in f . At each feature extraction stage i , deep supervision attaches an auxiliary classifier g_i for providing intermediate supervision. Thus, there are K classifiers in total which have the following form:

$$\begin{aligned} c_1(x) &= g_1 \circ f_1(x) \\ c_2(x) &= g_2 \circ f_2 \circ f_1(x) \\ &\dots \\ c_K(x) &= g_K \circ f_K \circ f_{K-1} \circ \dots \circ f_1(x). \end{aligned} \tag{1}$$

Given a set of training samples $\mathcal{X} = \{x_i\}_{i=1}^n$ and its corresponding labels $\mathcal{Y} = \{y_i\}_{i=1}^n$, the training loss of deep supervision \mathcal{L}_{DS} can be formulated as

$$\mathcal{L}_{\text{DS}} = \underbrace{\mathcal{L}_{\text{CE}}(c_K(\mathcal{X}), \mathcal{Y})}_{\text{from standard training}} + \alpha \cdot \sum_{i=1}^{K-1} \underbrace{\mathcal{L}_{\text{CE}}(c_i(\mathcal{X}), \mathcal{Y})}_{\text{from deep supervision}}, \tag{2}$$

where \mathcal{L}_{CE} indicates the cross entropy loss. The first and the second item in the loss function indicate the standard training loss and the additional loss from deep supervision for the intermediate layers, respectively. α is a hyper-parameter to balance the two loss items. Recently, some research has been proposed to apply layer-wise consistency on deep supervision, which additionally minimizes the KL divergence between the prediction of auxiliary classifiers and the final classifier [55,40]. These methods can also be considered as the knowledge distillation which regards the final classifier as the teacher and the auxiliary classifiers as the students. Their training loss can be formulated as

$$\mathcal{L}_{\text{DS}} + \beta \cdot \sum_{i=1}^{K-1} \mathcal{L}_{\text{KL}}(c_i(\mathcal{X}), c_K(\mathcal{X})), \tag{3}$$

where β is a hyper-parameter to balance the two loss functions.

3.2 Contrastive Deep Supervision

In this subsection, we first introduce the formulation of contrastive learning. For a minibatch of N images $\{x_1, x_2, \dots, x_N\}$, we apply stochastic data augmentation to each image twice, resulting in a batch of $2N$ images. For convenience, we denote x_i and x_{N+i} images as the two augmentations from the same image, which is regarded as a positive pair. Denote $z = c(x)$ as the normalized projection head outputs, contrastive learning loss (*a.k.a.* NT-Xtrent [7]) can be formulated as

$$\mathcal{L}_{\text{Contra}} = - \sum_{i=1}^N \log \frac{\exp(z_i \cdot z_{i+N})/\tau}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(z_i \cdot z_k)/\tau}, \quad (4)$$

where $\mathbb{1} \in \{0, 1\}$ is an indicator function evaluating to 1 if $k \neq i$ and τ is a temperature hyper-parameter. Intuitively, $\mathcal{L}_{\text{Contra}}$ encourages the encoder network to learn similar representation for different augmentations from the same image while increasing the difference between representations of the augmentations from different images.

The main difference between deep supervision and our method is that deep supervision trains the auxiliary classifiers by the cross entropy loss while our method trains them with the contrastive loss $\mathcal{L}_{\text{Contra}}$. By denoting the contrastive loss at c_i as $\mathcal{L}_{\text{Contra}}(\mathcal{X}; c_i)$, then the training loss of our contrastive deep supervision \mathcal{L}_{CDS} can be formulated as

$$\mathcal{L}_{\text{CDS}} = \underbrace{\mathcal{L}_{\text{CE}}(c_K(\mathcal{X}), \mathcal{Y})}_{\text{from standard training}} + \lambda_1 \sum_{i=1}^{K-1} \underbrace{\mathcal{L}_{\text{Contra}}(\mathcal{X}; c_i)}_{\text{from our method}}, \quad (5)$$

where the first and the second item indicate the standard training loss and the additional loss in our method for the intermediate layers, respectively. λ_1 is a hyper-parameter to balance the two loss items.

Based on the above formulation on supervised learning, we can extend contrastive deep supervision in semi-supervised learning and knowledge distillation.

Semi-supervised Learning In semi-supervised learning, we assume that there is a labeled dataset \mathcal{X}_1 with its labels \mathcal{Y}_1 and an unlabeled dataset \mathcal{X}_2 . On the labeled data, contrastive deep supervision can be applied directly with \mathcal{L}_{CDS} . On the unlabeled data, due to the lack of labels, contrastive deep supervision only optimize the contrastive learning loss $\mathcal{L}_{\text{Contra}}$, which can be formulated as

$$\mathcal{L}_{\text{CDS}}(\mathcal{X}_1, \mathcal{Y}_1) + \mathcal{L}_{\text{Contra}}(\mathcal{X}_2) \quad (6)$$

Knowledge Distillation The intermediate layers in contrastive deep supervision are supervised with contrastive learning and thus they can learn the invariance to different data augmentation. As shown in previous research, these data augmentation invariance is beneficial to various downstream tasks [31]. In this paper, we further propose to improve knowledge distillation with contrastive

deep supervision by transferring the data augmentation invariance learned by the teachers to the students. Denote the student model and the teacher model in knowledge distillation as f^S and f^T respectively, the naïve feature-based knowledge distillation directly minimizes the distance between the backbone features of the student and the teacher, which can be formulated as

$$\sum_{i=1}^K \|f_i^T(\mathcal{X}) - f_i^S(\mathcal{X})\|_2. \quad (7)$$

In contrast, knowledge distillation with contrastive deep supervision minimizes the distance between the embedding vectors (the output of the projection heads) of the student and the teacher, which can be formulated as

$$\mathcal{L}_{\text{CDS for KD}} = \sum_{i=1}^{K-1} \|c_i^T(\mathcal{X}) - c_i^S(\mathcal{X})\|_2. \quad (8)$$

Now we can formulate the overall training loss of the student as

$$\mathcal{L}_{\text{DCDS}} = \mathcal{L}_{\text{CDS}} + \lambda_2 \cdot \mathcal{L}_{\text{CDS for KD}} + \lambda_3 \cdot \mathcal{L}_{\text{KL}}(c_K^T(\mathcal{X}), c_K^S(\mathcal{X})), \quad (9)$$

where λ_2 and λ_3 are the hyper-parameters to balance different loss items. Following previous works in deep supervision, we do not set an individual hyper-parameter for each projection head for convenience in hyper-parameter tuning.

3.3 Other Details and Tricks

Design of Projection Heads In contrastive deep supervision, several projection heads are added to the intermediate layers of neural networks during the training period. These projection heads map the backbone features into a normalized embedding space, where the contrastive learning loss is applied. As discussed in related works, the architecture of the projection head is crucial to model performance [8]. Usually, the projection head is a non-linear projection stacked by two fully connected layers and a ReLU function. However, in contrastive deep supervision, the input feature comes from the intermediate layers instead of the final layer, and thus it is more challenging to project them properly [8]. Hence, we increase the complexity of these projection heads by adding convolutional layers before the non-linear projection.

Contrastive Learning The proposed contrastive deep supervision is a general training framework and does not depend on a specific contrastive learning method. In this paper, we adopt SimCLR [7] and SupCon [34] as the contrastive learning method in most experiments. We argue that the performance of our method can be further improved by using better contrastive learning method.

Negative Samples Previous studies show that the number of negative samples has a vital influence on the performance of contrastive learning. Accordingly, a large batch size, a momentum encoder or a memory bank is usually required [7,19,16]. In contrastive deep supervision, we do not use any of these solutions because the supervised loss (\mathcal{L}_{CE} in Equation 5) is enough to prevent contrastive learning from converging to the collapsing solutions.

Table 1. Comparison experiments (top-1 accuracy / %) with the other deep supervision methods on CIFAR100.

Method	RNT18	RNT50	RNT101	RXT50	RXT101	WRN50	WRN101	SET18	SET50	PAT18
Base	77.45	77.81	78.65	79.85	80.67	79.46	79.98	77.46	78.02	76.84
DSN	78.30	78.96	79.37	81.02	81.70	80.98	81.30	78.28	79.46	77.40
DKS	78.96	80.95	81.39	82.27	82.98	81.95	82.58	79.32	80.76	78.96
DHM	78.82	81.12	81.27	82.14	83.27	81.76	82.76	79.14	80.72	78.32
Ours	80.84	81.31	83.12	82.81	83.87	82.28	83.93	80.13	81.51	80.76

Table 2. Comparison experiments (top-1 accuracy / %) with the other deep supervision methods on CIFAR10.

Method	RNT18	RNT50	RNT101	RXT50	RXT101	WRN50	WRN101	SET18	SET50	PAT18
Base	94.96	95.07	95.13	95.09	95.34	95.01	95.27	94.86	95.11	94.78
DSN	95.31	95.41	95.63	95.39	95.70	95.27	95.78	95.21	95.41	95.13
DKS	95.72	95.90	96.21	95.98	96.10	95.50	96.12	95.74	95.72	95.47
DHM	95.61	95.87	96.04	96.10	96.27	95.62	96.31	95.59	95.77	95.38
Ours	96.49	96.78	97.02	96.76	97.05	96.88	97.01	96.50	96.73	96.37

Table 3. Comparison with the other deep supervision methods on ImageNet.

Metric	Model	Baseline	DSN	DKS	DHM	Ours
top-1	RNT18	69.21	69.54	71.32	71.29	72.85
	RNT34	73.17	73.29	74.01	73.89	76.19
	RNT50	75.30	75.37	76.47	76.57	78.25
top-5	RNT18	89.01	88.87	89.20	90.06	91.30
	RNT34	91.24	91.30	91.87	91.66	93.08
	RNT50	92.20	92.49	93.60	93.24	93.99

4 Experiment

4.1 Experiment Setting

Common Image Classification For common image classification, our method has been evaluated on three datasets, including CIFAR10, CIFAR100 and ImageNet [36,13] with kinds of neural networks including ResNet (RNT), ResNeXt (RXT), Wide ResNet (WRN), SENet (SET), PreAct ResNet (PAT), MobileNetv1, MobileNetv2, ShuffleNetv1 and ShuffleNetv2 [20,66,71,26,21,25,54,80].

Fine-grained Image Classification For fine-grained image classification, our method has been evaluated on five popular datasets, including CUB200-

Table 4. Experiments on different object detection models on COCO2017. ResNet50 models are pre-trained on ImageNet with different deep supervision methods and then utilized as the backbones of these detectors.

Model	Method	AP	AP _S	AP _M	AP _L
Faster RCNN	Baseline	37.4	21.2	41.0	48.1
	DSN	37.3 _{-0.1}	21.0 _{-0.2}	40.8 _{-0.2}	48.3 _{-0.2}
	DKS	37.5 _{+0.1}	21.2 _{+0.0}	41.5 _{+0.5}	47.6 _{-0.5}
	DHM	37.6 _{+0.2}	21.3 _{+0.1}	41.3 _{+0.3}	48.2 _{+0.1}
	Ours	38.3_{+0.9}	21.6_{+0.4}	42.0_{+1.0}	50.1_{+2.0}
RetinaNet	Baseline	36.5	20.4	40.3	48.1
	DSN	36.3 _{-0.2}	20.1 _{-0.3}	40.0 _{-0.3}	48.1 _{0.0}
	DKS	36.7 _{+0.2}	20.1 _{-0.3}	40.9 _{+0.6}	48.2 _{+0.1}
	DHM	36.7 _{+0.2}	20.0 _{-0.4}	40.7 _{+0.4}	48.5 _{+0.4}
	Ours	37.3_{+0.8}	21.2_{+0.8}	41.0_{+0.7}	47.9_{-0.2}

Table 5. Comparison (top-1 acc. / %) with deep supervision methods with ResNet50 for fine-grained classification. Models are trained from scratch.

Method	CUB	Cars	Flowers	Dogs	Aircrafts
Baseline	60.65	79.86	87.52	64.00	74.07
DSN	62.37 _{+1.72}	81.04 _{+1.18}	88.54 _{+1.02}	66.32 _{+2.32}	74.49 _{+0.42}
DKS	63.59 _{+2.94}	81.52 _{+1.66}	88.94 _{+0.40}	68.31 _{+4.31}	75.07 _{+1.00}
DHM	64.01 _{+3.36}	81.49 _{+1.63}	89.03 _{+1.51}	68.38 _{+4.38}	75.00 _{+0.93}
Ours	64.65_{+4.00}	82.07_{+2.21}	89.26_{+1.74}	69.02_{+5.02}	75.43_{+1.36}

Table 6. Comparison (top-1 acc. %) with deep supervision methods with ResNet50 for fine-grained classification. Models are finetuned from ImageNet pre-trained weights.

Method	CUB	Cars	Flowers	Dogs	Aircrafts
Baseline	78.50	90.25	97.68	76.47	87.43
DSN	80.14 _{+1.64}	91.32 _{+1.07}	98.64 _{+0.96}	77.21 _{+0.74}	89.31 _{+1.88}
DKS	81.34 _{+2.84}	92.54 _{+2.29}	99.01 _{+1.33}	78.32 _{+1.85}	89.20 _{+1.77}
DHM	81.27 _{+2.77}	92.31 _{+2.06}	98.84 _{+1.16}	78.20 _{+1.73}	89.57 _{+2.14}
Ours	82.10_{+3.60}	92.90_{+2.65}	99.39_{+1.71}	80.99_{+4.52}	90.52_{+3.09}

2011 [59], Stanford Cars [35], Oxford Flowers [45], Stanford Dogs [33] and FGVC Aircraft [44]. ResNet50 is utilized as the classifier for all the experiments.

Object Detection For object detection, our method has been evaluated on MS COCO2017 [41] with Faster RCNN and RetinaNet by MMDetection [4].

Semi-supervised Learning Semi-supervised learning experiments have been conducted on CIFAR100, CIFAR10 with ResNet18. For each dataset, we have evaluated our method with 10%, 20%, 30% and 40% labels.

Table 7. Comparison experiments (top-1 and top-5 accuracy / %) with the other eight knowledge distillation methods on ImageNet with ResNet. Numbers in bold indicate the highest. Results marked with \dagger come from the paper of SSKD [67].

Metric	Model	Base	KD	AT	RKD	SP	CRD	CC \dagger	OKD \dagger	SSKD \dagger	Ours
top-1	RNT18	69.21	70.52	70.74	70.63	70.61	71.07	69.96	70.55	71.62	73.23
	RNT34	73.17	74.44	74.69	74.61	74.60	74.99	–	–	–	76.65
	RNT50	75.30	76.62	76.79	76.92	76.88	77.21	–	–	–	78.68
top-5	RNT18	89.01	89.88	90.00	89.71	89.80	91.06	89.17	89.59	90.67	91.56
	RNT34	91.24	92.07	92.18	92.14	92.10	92.58	–	–	–	93.38
	RNT50	92.20	93.36	93.51	93.60	93.58	93.88	–	–	–	94.42

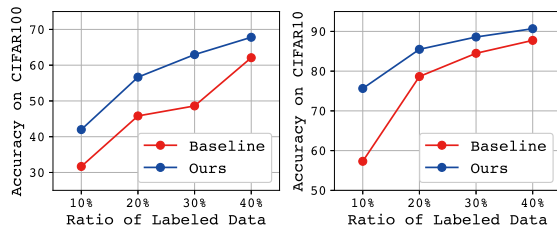


Fig. 2. Experimental results of semi-supervised training on CIFAR100 and CIFAR10 with ResNet18.

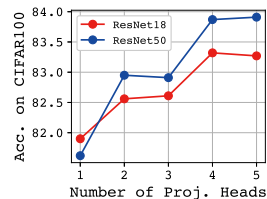


Fig. 3. Influence from the number of projection heads.

Comparison Methods Three previous deep supervision methods are utilized for comparison, including DSN [38], DKS [55] and DHM [40]. In knowledge distillation experiments, we have evaluated our method with nine knowledge distillation methods, including KD [23], FitNet [53], AT [72], RKD [49], SP [58] and CRD [57]. Besides, we also cite results on ImageNet of CC [50], OKD [84], and SSKD [67] from the paper of SSKD.

4.2 Experimental Results

Image Classification Experimental results on CIFAR100, CIFAR10 and ImageNet are shown in Table 1, Table 2 and Table 3, respectively. It is observed that: **(a)** Our method achieves 3.44% and 1.70% top-1 accuracy improvements on CIFAR100 and CIFAR10 on average, respectively. It consistently outperforms the second-best deep supervision method by 1.05% and 0.90% on the two datasets, respectively. **(b)** On ImageNet, contrastive deep supervision leads to 3.64%, 3.02% and 2.95% top-1 accuracy improvements on ResNet18, ResNet34 and ResNet50, respectively. On average, it outperforms the baseline and the second-best method by 3.20% and 1.83% top-1 accuracy, respectively.

Object Detection Table 4 shows the the performance of our method on object detection. In these experiments, We firstly pretrain the ResNets on ImageNet with standard training (Baseline), three deep supervision methods, and

Table 8. Comparison with the other knowledge distillation methods on CIFAR.

CIFAR100								
Model	Base	KD	FitNet	AT	RKD	SP	CRD	Ours
ResNet18	77.45	78.68	78.15	78.09	78.21	78.19	81.41	83.31
ResNet50	77.81	79.19	78.42	78.34	78.94	78.81	82.45	83.53
ResNet101	78.65	80.40	80.78	80.97	81.24	80.94	82.57	84.80
ResNeXt50	79.85	81.41	82.67	82.59	83.71	82.67	83.41	84.41
ResNeXt101	80.67	82.03	82.51	82.43	83.01	82.64	84.50	85.37
WRNet50	79.46	81.02	81.29	81.16	82.06	82.07	82.94	84.27
WRNet101	79.98	81.82	82.07	82.16	82.54	82.49	83.07	85.04
SENet18	77.46	78.92	79.09	79.15	79.41	79.31	81.22	82.68
SENet50	78.02	79.78	80.13	80.45	80.69	80.71	81.79	83.36
SENet101	78.92	80.31	80.54	80.53	80.74	80.52	82.75	84.15
MobileNetV1	68.32	70.04	70.25	70.17	70.89	70.19	72.68	73.79
MobileNetV2	69.34	70.58	70.64	70.51	70.83	70.68	71.82	72.61
ShuffleNetV1	72.46	74.08	74.19	74.11	74.56	74.68	75.11	75.77
ShuffleNetV2	72.81	74.39	74.47	74.51	74.82	74.67	75.62	76.11
PreActNet18	76.84	78.25	78.34	78.67	79.01	79.12	81.62	82.83
PreActNet50	77.31	79.04	79.27	79.54	79.82	79.76	81.27	83.42
CIFAR10								
Model	Base	KD	FitNet	AT	RKD	SP	CRD	Ours
ResNet18	94.96	95.24	95.31	95.26	95.31	95.27	95.81	96.84
ResNet50	95.07	95.31	95.45	95.47	95.33	95.29	96.21	97.08
ResNet101	95.13	95.39	95.71	95.49	95.43	95.18	96.37	97.40
ResNeXt50	95.09	95.27	95.36	95.68	95.59	95.37	96.49	97.15
ResNeXt101	95.34	95.68	95.92	95.78	95.81	95.38	96.51	97.40
WRNet50	95.01	95.34	95.38	95.34	95.61	95.73	96.17	97.37
WRNet101	95.27	95.51	95.48	95.71	95.99	95.82	96.34	97.39
SENet18	94.86	95.21	95.30	95.47	95.34	95.41	96.00	96.96
SENet50	95.11	95.39	95.44	95.64	95.57	95.47	96.21	97.19
SENet101	95.30	95.64	95.81	95.78	95.81	95.77	96.19	97.36
MobileNetV1	90.24	91.27	92.59	92.87	93.01	92.90	93.27	93.94
MobileNetV2	90.76	91.09	91.57	91.75	91.82	91.83	92.17	92.87
ShuffleNetV1	91.57	91.99	92.30	92.19	92.47	92.38	93.08	94.04
ShuffleNetV2	91.19	91.87	92.23	92.41	92.30	92.54	92.90	93.16
PreActNet18	94.78	95.08	95.28	95.39	95.51	95.69	96.07	96.70
PreActNet50	94.89	95.21	95.57	95.49	95.37	95.48	96.11	96.93

our method, and then finetuning them as the backbone for object detection models, including RetinaNet and Faster RCNN on COCO2017 datasets. It is observed that with backbones pre-trained with our method, there are 0.9 and 0.8 AP improvements on Faster RCNN and RetinaNet respectively, which outperforms the second-best method by 0.6 AP, indicating that the representation learned with our method are more beneficial to downstream tasks.

Fine-grained Image Classification Experiments on fine-grained image classification are shown in Table 6. It is observed that: **(a)** Contrastive deep supervision leads to consistent and significant accuracy improvements on the five datasets. On average, it leads to 3.80%, 2.43%, 1.73%, 4.77% and 2.25% accuracy improvements on the five datasets, respectively. **(b)** Besides, the benefits of our method in “*finetuning from ImageNet*” and “*training from scratch*” are very similar (except on Aircraft), which indicates that the effectiveness of our method is consistent in different training settings.

Semi-supervised Learning Experiments on semi-supervised learning with ResNet18 on CIFAR10 and CIFAR100 are shown in Figure 2. It is observed that: **(a)** Our method leads to consistent accuracy improvements at all the ratios of labeled data. **(b)** The benefits of our method become larger when there is less labeled data, which indicates that our method is effective in using the unlabeled data to optimize the intermediate layers.

Knowledge Distillation Knowledge distillation experiments on ImageNet and CIFAR are shown in Table 7 and Table 8, respectively. It is observed that: **(a)** Our method achieves 5.07% and 2.20% top-1 accuracy improvements on CIFAR100 and CIFAR10 on average, outperforming the second-best KD method by 1.40% and 0.87% on the two datasets, respectively. **(b)** The similar results can also be observed in ImageNet experiments. Our method leads to 4.02%/2.55%, 3.48%/2.14% and 3.38%/2.22% top-1/top-5 accuracy improvements on ResNet18, ResNet34 and ResNet50, respectively. On average, it outperforms the baseline and the second-best method by 3.62% and 1.76% top-1 accuracy, respectively.

5 Discussion

5.1 Contrastive Deep Supervision as a Regularizer

Loss Curves Regularization methods in deep learning are usually utilized to avoid model overfitting by introducing additional penalties or loss. In this subsection, we show that the contrastive learning loss introduced by our method in the intermediate layers works as a regularizer. Figure 4 shows the cross entropy loss between predicted results and labels during the training period from two ResNet18 models trained by the standard method and our method, respectively. It is observed that at most of epochs, the baseline model has lower cross entropy loss than our model. When both models are converged (epoch 280-300), the baseline model has only 0.005 loss while our model still has 0.025 loss. These observations indicate that there is serve overfitting in the baseline model while deep contrastive supervision can alleviate overfitting and thus improve the accuracy.

Uncertainty Estimation Besides, the comparison on expected calibrated error (ECE) of models trained with the standard method and our method has been shown in Figure 5. A lower ECE indicates that the predicted probability of a neural network estimates representative of the true correctness likelihood better [17]. It is observed that compared with the baseline model, our method leads to a lower ECE, indicating better uncertainty estimation and interpretability.

Table 9. Comparison between our method and contrastive learning methods with ResNet50 on ImageNet. Baseline¹⁻²: Two baselines trained with and without AutoAugmentation [12]. SupCon¹⁻³: Three models trained by supervised contrastive learning with different hyper-parameters. BYOL: ResNet50 unsupervisedly pre-trained by 1000 epochs and then supervisedly finetuned. BYOL+DSN: ResNet50 pre-trained with BYOL and then finetuned with deep supervision. Ours^{1,3}: ResNet50 trained with contrastive deep supervision in different settings. Ours²: ResNet50 trained with contrastive deep supervision+ knowledge distillation.

Method	Batchsize	Epoch	AutoAug	top-1 acc. (%)
Baseline ¹	256	90	×	75.3
Baseline ²	4096	270	✓	77.6
SupCon ¹	6144	350	✓	78.7
SupCon ²	512	350	✓	74.5
SupCon ³	6144	100	✓	77.0
BYOL	1024	1080	×	77.7
BYOL+DSN	1024	1080	×	78.2
Ours ¹	256	90	×	78.3
Ours ²	256	90	×	78.7
Ours ³	256	350	✓	79.8

5.2 Comparison with Contrastive Learning

Comparison between our method and two “pretrain & finetune” contrastive learning methods is shown in Table 9. It is observed that without a large batch size and the advanced data augmentation policy (AutoAugment), contrastive deep supervision (Ours¹) with only 25% training time achieves 0.4% lower accuracy than SupCon³. Besides, contrastive deep supervision with the same training time and data augmentation (Ours³) achieves 1.1% and 1.6% higher accuracy than SupCon³ and BYOL+DSN, respectively, which demonstrates the advantage of our method over the traditional contrastive learning methods.

5.3 Ablation Study on Knowledge Distillation

The main difference between the naïve feature distillation and feature distillation with our contrastive deep supervision is “*what to distill*”. Naïve feature distillation distills the backbone features while our method distills the embedding learned by contrastive deep supervision. To further demonstrate its effectiveness, we have trained a ResNet50 model on CIFAR100 with both contrastive deep supervision and distillation on backbone features. Experimental results show this model achieves 82.26% accuracy, which is 1.27% lower than distilling the embedding. These results demonstrate that distilling the embedding learned by contrastive deep supervision is more beneficial.

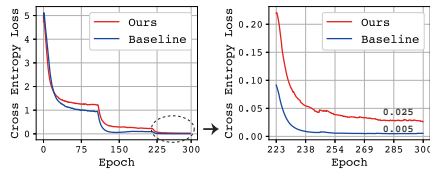


Fig. 4. Comparison on the cross entropy loss between predicted results and labels during the training period. Note that our method also leads to better accuracy (80.84% vs 77.45%)

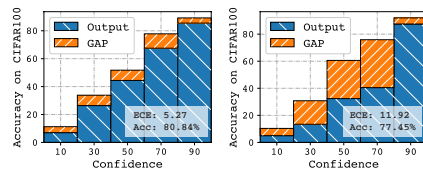


Fig. 5. Comparison on reliability diagrams. “GAP” indicates the difference between confidence and accuracy. “Output” indicates accuracy. ECE: Expected Calibrated Error (lower is better).

5.4 Sensitivity Study

Where to Apply Projection Heads We study the influence from the position of projection heads with the following four schemes: (1) *uniform scheme* - applying projection heads into different depths uniformly; (2) *downsampling scheme* - applying projection heads into the layers before downsampling; (3) *shallow scheme* - applying projection heads into only the shallower layers; (4) *deep scheme* - applying projection heads to only the deeper layers; Experimental results on CIFAR100 with ResNet50 show that the four schemes achieves 81.23%, 81.31%, 81.07% and 80.99% accuracy, respectively. It is observed that both *uniform* and *downsampling* schemes leads to excellent performance, indicating our method is not sensitive to where to apply projection heads.

The Number of Projection Heads We have studied the influence from the number of projection heads in Figure 3. It is observed that when there are less than five projection heads, more projection heads tend to achieve better performance. The fifth projection head does not leads to more accuracy improvements.

6 Conclusion

This paper proposes *contrastive deep supervision*, a novel training methodology that directly optimizes the intermediate layers of deep neural networks with contrastive learning. It enables the neural network to learn better visual representation without additional computation and storage in inference. Experiments on nine datasets with eleven neural networks have demonstrated its effectiveness in general image classification, fine-grained image classification and object detection for traditional supervised learning, semi-supervised learning and knowledge distillation. It outperforms the previous deep supervision methods, knowledge distillation methods, and contrastive learning methods by a clear margin. Besides, we also show that contrastive deep supervision works as a regularizer to prevent models from overfitting, and thus leads to better uncertainty estimation.

References

1. Ahn, S., Hu, S.X., Damianou, A., Lawrence, N.D., Dai, Z.: Variational information distillation for knowledge transfer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9163–9171 (2019)
2. Buciluă, C., Caruana, R., Niculescu-Mizil, A.: Model compression. In: Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 535–541. ACM (2006)
3. Chaitanya, K., Erdil, E., Karani, N., Konukoglu, E.: Contrastive learning of global and local features for medical image segmentation with limited annotations. *Advances in Neural Information Processing Systems* **33** (2020)
4. Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., et al.: Mmdetection: Open mmlab detection toolbox and benchmark. arXiv preprint arXiv:1906.07155 (2019)
5. Chen, L., Wang, D., Gan, Z., Liu, J., Henao, R., Carin, L.: Wasserstein contrastive representation distillation. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 16296–16305. Computer Vision Foundation / IEEE (2021)
6. Chen, L., Wang, D., Gan, Z., Liu, J., Henao, R., Carin, L.: Wasserstein contrastive representation distillation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16296–16305 (2021)
7. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
8. Chen, T., Kornblith, S., Swersky, K., Norouzi, M., Hinton, G.E.: Big self-supervised models are strong semi-supervised learners. *Advances in neural information processing systems* **33**, 22243–22255 (2020)
9. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297 (2020)
10. Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15750–15758 (2021)
11. Chen, X., Xie, S., He, K.: An empirical study of training self-supervised visual transformers. arXiv e-prints pp. arXiv–2104 (2021)
12. Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V.: Autoaugment: Learning augmentation strategies from data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
13. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR. pp. 248–255 (2009)
14. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: NAACL (2018)
15. Furlanello, T., Lipton, Z.C., Tschannen, M., Itti, L., Anandkumar, A.: Born again neural networks. In: ICML (2018)
16. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al.: Bootstrap your own latent—a new approach to self-supervised learning. *Advances in Neural Information Processing Systems* **33**, 21271–21284
17. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: International Conference on Machine Learning. pp. 1321–1330. PMLR (2017)

18. Han, Z., Fu, Z., Chen, S., Yang, J.: Contrastive embedding for generalized zero-shot learning. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 2371–2381. Computer Vision Foundation / IEEE (2021)
19. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9729–9738 (2020)
20. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
21. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: European conference on computer vision. pp. 630–645. Springer (2016)
22. Heo, B., Kim, J., Yun, S., Park, H., Kwak, N., Choi, J.Y.: A comprehensive overhaul of feature distillation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1921–1930 (2019)
23. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. In: NeurIPS (2014)
24. Hou, J., Graham, B., Nießner, M., Xie, S.: Exploring data-efficient 3d scene understanding with contrastive scene contexts. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 15587–15597. Computer Vision Foundation / IEEE (2021)
25. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. In: CVPR (2017)
26. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: CVPR. pp. 7132–7141 (2018)
27. Hu, Q., Wang, X., Hu, W., Qi, G.: Adco: Adversarial contrast for efficient learning of unsupervised representations from self-trained negative adversaries. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 1074–1083 (2021)
28. Huang, G., Chen, D., Li, T., Wu, F., van der Maaten, L., Weinberger, K.Q.: Multi-scale dense networks for resource efficient image classification. In: ICLR (2018)
29. Huang, G., Sun, Y., Liu, Z., Sedra, D., Weinberger, K.Q.: Deep networks with stochastic depth. In: ECCV. pp. 646–661 (2016)
30. Hyun Lee, S., Ha Kim, D., Cheol Song, B.: Self-supervised knowledge distillation using singular value decomposition. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 335–350 (2018)
31. Jaiswal, A., Babu, A.R., Zadeh, M.Z., Banerjee, D., Makedon, F.: A survey on contrastive self-supervised learning. *Technologies* **9**(1), 2 (2021)
32. Jeon, S., Min, D., Kim, S., Sohn, K.: Mining better samples for contrastive learning of temporal correspondence. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 1034–1044 (2021)
33. Khosla, A., Jayadevaprakash, N., Yao, B., Fei-Fei, L.: Novel dataset for fine-grained image categorization. In: First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition. Colorado Springs, CO (June 2011)
34. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. *Advances in Neural Information Processing Systems* **33**, 18661–18673 (2020)
35. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: 4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13). Sydney, Australia (2013)

36. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Tech. rep., Citeseer (2009)
37. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: International Conference on Learning Representations (ICLR). vol. 4, p. 6 (2017)
38. Lee, C.Y., Xie, S., Gallagher, P., Zhang, Z., Tu, Z.: Deeply-supervised nets. In: Artificial Intelligence and Statistics. pp. 562–570 (2015)
39. Li, C., Zeeshan Zia, M., Tran, Q.H., Yu, X., Hager, G.D., Chandraker, M.: Deep supervision with shape concepts for occlusion-aware 3d object parsing. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5465–5474 (2017)
40. Li, D., Chen, Q.: Dynamic hierarchical mimicking towards consistent optimization objectives. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7642–7651 (2020)
41. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
42. Liu, M., Chen, X., Zhang, Y., Li, Y., Rehg, J.M.: Attention distillation for learning video representations. In: BMVC (2020)
43. Liu, Y., Shu, C., Wang, J., Shen, C.: Structured knowledge distillation for dense prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020)
44. Maji, S., Kannala, J., Rahtu, E., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. Tech. rep. (2013)
45. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: 2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing. pp. 722–729. IEEE (2008)
46. Noroozi, M., Vinjimoor, A., Favaro, P., Pirsiavash, H.: Boosting self-supervised learning via knowledge transfer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9359–9367 (2018)
47. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
48. Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: European Conference on Computer Vision. pp. 319–345. Springer (2020)
49. Park, W., Kim, D., Lu, Y., Cho, M.: Relational knowledge distillation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3967–3976 (2019)
50. Peng, B., Jin, X., Liu, J., Li, D., Wu, Y., Liu, Y., Zhou, S., Zhang, Z.: Correlation congruence for knowledge distillation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 5007–5016 (2019)
51. Reiß, S., Seibold, C., Freytag, A., Rodner, E., Stiefelhagen, R.: Every annotation counts: Multi-label deep supervision for medical image segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19–25, 2021. pp. 9532–9542. Computer Vision Foundation / IEEE (2021)
52. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in neural information processing systems. pp. 91–99 (2015)
53. Romero, A., Ballas, N., Kahou, S.E., Chassang, A., Gatta, C., Bengio, Y.: Fitnets: Hints for thin deep nets. In: ICLR (2015)
54. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4510–4520 (2018)

55. Sun, D., Yao, A., Zhou, A., Zhao, H.: Deeply-supervised knowledge synergy. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6997–7006 (2019)
56. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: Advances in neural information processing systems. pp. 1195–1204 (2017)
57. Tian, Y., Krishnan, D., Isola, P.: Contrastive representation distillation. In: ICLR (2020)
58. Tung, F., Mori, G.: Similarity-preserving knowledge distillation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1365–1374 (2019)
59. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset. Tech. Rep. CNS-TR-2011-001, California Institute of Technology (2011)
60. Wang, F., Liu, H.: Understanding the behaviour of contrastive loss. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 2495–2504. Computer Vision Foundation / IEEE (2021)
61. Wang, L., Huang, J., Li, Y., Xu, K., Yang, Z., Yu, D.: Improving weakly supervised visual grounding by contrastive knowledge distillation. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 14090–14100. Computer Vision Foundation / IEEE (2021)
62. Wang, L., Lee, C.Y., Tu, Z., Lazebnik, S.: Training deeper convolutional networks with deep supervision. arXiv preprint arXiv:1505.02496 (2015)
63. Wang, P., Han, K., Wei, X., Zhang, L., Wang, L.: Contrastive learning based hybrid networks for long-tailed image classification. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 943–952 (2021)
64. Xie, E., Ding, J., Wang, W., Zhan, X., Xu, H., Li, Z., Luo, P.: Detco: Unsupervised contrastive learning for object detection. arXiv preprint arXiv:2102.04803 (2021)
65. Xie, Q., Luong, M.T., Hovy, E., Le, Q.V.: Self-training with noisy student improves imagenet classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10687–10698 (2020)
66. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: CVPR. pp. 5987–5995 (2017)
67. Xu, G., Liu, Z., Li, X., Loy, C.C.: Knowledge distillation meets self-supervision. In: European Conference on Computer Vision. pp. 588–604. Springer (2020)
68. Yang, M., Li, Y., Huang, Z., Liu, Z., Hu, P., Peng, X.: Partially view-aligned representation learning with noise-robust contrastive loss. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 1134–1143. Computer Vision Foundation / IEEE (2021)
69. Yim, J., Joo, D., Bae, J., Kim, J.: A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4133–4141 (2017)
70. Yu, J., Huang, T.S.: Universally slimmable networks and improved training techniques. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
71. Zagoruyko, S., Komodakis, N.: Wide residual networks. In: BMVC (2016)
72. Zagoruyko, S., Komodakis, N.: Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In: ICLR (2017)

73. Zeng, G., Yang, X., Li, J., Yu, L., Heng, P.A., Zheng, G.: 3d u-net with multi-level deep supervision: fully automatic segmentation of proximal femur in 3d mr images. In: International workshop on machine learning in medical imaging. pp. 274–282. Springer (2017)
74. Zhang, H., Koh, J.Y., Baldrige, J., Lee, H., Yang, Y.: Cross-modal contrastive learning for text-to-image generation. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 833–842. Computer Vision Foundation / IEEE (2021)
75. Zhang, L., Kaisheng, M.: Improve object detection with feature-based knowledge distillation: Towards accurate and efficient detectors. In: ICLR (2021)
76. Zhang, L., Shi, Y., Shi, Z., Ma, K., Bao, C.: Task-oriented feature distillation. In: NeurIPS (2020)
77. Zhang, L., Song, J., Gao, A., Chen, J., Bao, C., Ma, K.: Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In: arXiv preprint:1905.08094 (2019)
78. Zhang, L., Tan, Z., Song, J., Chen, J., Bao, C., Ma, K.: Scan: A scalable neural networks framework towards compact and efficient models. ArXiv **abs/1906.03951** (2019)
79. Zhang, L., Yu, M., Chen, T., Shi, Z., Bao, C., Ma, K.: Auxiliary training: Towards accurate and robust models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 372–381 (2020)
80. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6848–6856 (2018)
81. Zhang, Y., Chung, A.C.: Deep supervision with additional labels for retinal vessel segmentation task. In: International conference on medical image computing and computer-assisted intervention. pp. 83–91. Springer (2018)
82. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Object detectors emerge in deep scene cnns. arXiv preprint arXiv:1412.6856 (2014)
83. Zhou, P., Mai, L., Zhang, J., Xu, N., Wu, Z., Davis, L.S.: M2kd: Multi-model and multi-level knowledge distillation for incremental learning. arXiv preprint arXiv:1904.01769 (2019)
84. Zhu, X., Gong, S., et al.: Knowledge distillation by on-the-fly native ensemble. *Advances in neural information processing systems* **31** (2018)