

# Multi-Teacher Knowledge Distillation with Reinforcement Learning for Visual Recognition

Chuanguang Yang<sup>1</sup>, Xinqiang Yu<sup>1,2</sup>, Han Yang<sup>1,2</sup>, Zhulin An<sup>1\*</sup>, Chengqing Yu<sup>1,2</sup>, Libo Huang<sup>1</sup>, Yongjun Xu<sup>1</sup>

<sup>1</sup>Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

<sup>2</sup>University of Chinese Academy of Sciences, Beijing, China

{yangchuanguang,yuxinqiang21s,yanghan22s,anzhulin,yuchengqing22b,huanglibo,xyj}@ict.ac.cn

## Abstract

Multi-teacher Knowledge Distillation (KD) transfers diverse knowledge from a teacher pool to a student network. The core problem of multi-teacher KD is how to balance distillation strengths among various teachers. Most existing methods often develop weighting strategies from an individual perspective of teacher performance or teacher-student gaps, lacking comprehensive information for guidance. This paper proposes Multi-Teacher Knowledge Distillation with Reinforcement Learning (MTKD-RL) to optimize multi-teacher weights. In this framework, we construct both teacher performance and teacher-student gaps as state information to an agent. The agent outputs the teacher weight and can be updated by the return reward from the student. MTKD-RL reinforces the interaction between the student and teacher using an agent in an RL-based decision mechanism, achieving better matching capability with more meaningful weights. Experimental results on visual recognition tasks, including image classification, object detection, and semantic segmentation tasks, demonstrate that MTKD-RL achieves state-of-the-art performance compared to the existing multi-teacher KD works.

**Code** — <https://github.com/winycg/MTKD-RL>

## Introduction

Deep neural networks (Yang et al. 2020; Touvron et al. 2021b; Liang et al. 2024b,a) have achieved excellent performance for visual recognition tasks, but accompanied by the growth of computation complexity and memory footprint. Knowledge Distillation (KD) (Hinton, Vinyals, and Dean 2015; Zhang et al. 2023) becomes an effective way to improve a low-complexity student network, and has been widely applied to image classification (Yang et al. 2022b, 2023a), object detection (Li et al. 2024), segmentation (Yang et al. 2022c) and generation (Feng et al. 2024). The conventional KD often investigates a single teacher to transfer knowledge to a student. Compared to the single-teacher KD, multi-teacher KD (Zhang, Chen, and Wang 2022, 2023) provides more comprehensive and diverse knowledge, increasing the upper bound of student performance.

Although multi-teacher KD could lead to more significant improvements, it is more challenging since balancing distillation strengths among various teachers is a non-trivial problem. The vanilla multi-teacher KD assigns equal weights to distill a student (You et al. 2017). However, the equal distillation strength ignores two critical aspects: (1) *teacher performance*: various teachers may perform differently on the same data sample; (2) *teacher-student gaps*: a more powerful teacher may not result in a better student, because a simple student may not have enough capacity to mimic a large teacher, as discussed by Cho *et al.* (Cho and Hariharan 2019). Previous multi-teacher KD methods often explore weight generation from an individual perspective of the aspect (1) *teacher performance* (e.g. information entropy (Kwon et al. 2020; Zhang, Chen, and Wang 2022) and logits (Zhang, Chen, and Wang 2023)) or the aspect (2) *teacher-student gaps*, reflecting in gradient space (Du et al. 2020) and similarity matrix (Liu, Zhang, and Wang 2020). However, these methods only consider a single aspect to generate multi-teacher weights but neglect a comprehensive interaction among the teacher, student, and data samples.

This paper proposes Multi-Teacher Knowledge Distillation with Reinforcement Learning (MTKD-RL) to optimize the multi-teacher weight generation problem. The overview of MTKD-RL is illustrated in Fig.1. We formulate multi-teacher KD training as a Reinforcement Learning (RL) optimization process. We introduce an agent to output the action of multi-teacher weights according to the state information. The state, including teacher information (feature, logit, and cross-entropy loss) and teacher-student gaps (feature similarity and logit KL-divergence), encodes evidence of both the aspect (1) and (2), respectively. We use the student performance and teacher-student gaps as rewards to update the agent by policy gradient algorithm. The framework alternatively performs multi-teacher KD and agent optimization until the student is converged.

Compared to previous multi-teacher KD works, our method constructs teacher performance and teacher-student gaps as prompts, covering the aspect (1) and (2). MTKD-RL reinforces the interaction between the student and teacher using an agent in an RL-based decision mechanism, achieving better matching capability with more meaningful weights. Experimental results on image classification tasks demonstrate that MTKD-RL performs better than previous

\*Corresponding author

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

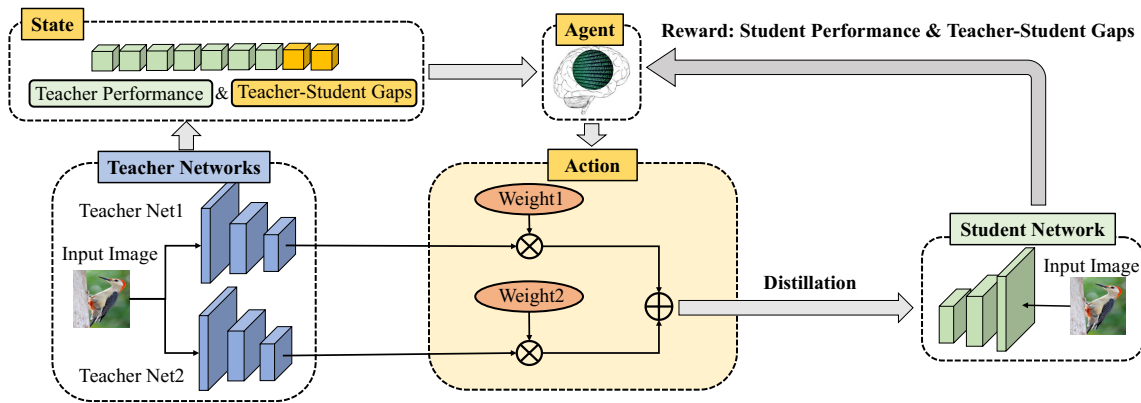


Figure 1: Overview of the basic idea about our proposed MTKD-RL.

state-of-the-art multi-teacher KD works. Extensive experiments on downstream object detection and semantic segmentation show that MTKD-RL guides the student to learn better features for dense prediction tasks.

The contributions are two-fold: (1) we propose MTKD-RL, an RL-based method to generate meaningful multi-teacher KD weights by reinforcing the interaction between the student and teacher. (2) Experimental results on visual recognition tasks show that MTKD-RL achieves state-of-the-art performance.

## Related Work

### Multi-Teacher Knowledge Distillation

The conventional KD framework (Hinton, Vinyals, and Dean 2015; Yang et al. 2021, 2022a,c, 2024) often improves a student network taught by a high-capacity teacher. However, the student performance may be bounded by a single teacher due to limited knowledge. Many works attempted to rely on multiple teacher networks for knowledge transfer. The core question of multi-teacher KD is how to balance the weights among various teacher networks. A straightforward method (You et al. 2017) is to assign an equal weight to each teacher, but it may ignore the diversity of teachers and lead to suboptimum performance.

Some subsequent works explore better multi-teacher KD algorithms with meaningful balancing weights. AE-KD (Du et al. 2020) formulates multi-teacher KD as a multi-objective optimization problem and derives the weight solutions from the perspective of gradient space. AMTML-KD (Liu, Zhang, and Wang 2020) introduces a latent factor for each teacher and calculates teacher importance by making an element-wise product with the student max-pooled features. EB-KD (Kwon et al. 2020) computes aggregation weights according to information-entropy, since the entropy measures the confidence of a teacher predictive distribution. However, EB-KD may assign an inaccurate weight when the teacher’s predictive category is not correct. To address this problem, CA-MKD (Zhang, Chen, and Wang 2022) introduces cross-entropy complemented by ground-truth labels to guarantee the correct category guidance. Instead of exploiting probability-level information, the recent

MMKD (Zhang, Chen, and Wang 2023) introduces a meta-learning mechanism (Finn, Abbeel, and Levine 2017) with a hard buffer to optimize aggregation weights of both the teachers’ features and logits. Although both MOBA (Kang et al. 2022) and RMTS (Yuan et al. 2021) involve RL and distillation, they have essential distinctions compared to MTKD-RL in *solved tasks*. MOBA boosts RL models in continuous control tasks, *e.g.* games and robotics. RMTS enhances BERT models for natural language processing tasks. Our MTKD-RL improves CNN or ViT models for visual recognition tasks. More references about multi-teacher KD works could refer to this survey (Yang et al. 2023b).

Although previous multi-teacher KD methods achieve good performance gains, they often consider a single level to optimize teacher weights. Moreover, they neglect the interaction between ensemble teachers and the student, especially student accuracy as a critical indicator. Therefore, the generated weights may not reflect the comprehensive capability of teachers and lack compatibility with a specific student. By contrast, our MTKD-RL method constructs both the teacher performance and teacher-student gaps as the state, optimized by the student performance as rewards, leading to more meaningful teacher weights.

### Reinforcement Learning

Reinforcement Learning (RL) has achieved great success in decision-making (Kaelbling, Littman, and Moore 1996; Liu et al. 2020; Chengqing et al. 2023). The basic framework of the RL system (Mnih et al. 2013; Van Hasselt, Guez, and Silver 2016) is to maximize the reward achieved by an agent when interacting with the environment. During RL training, the agent outputs an action using the observed state from the environment. After action execution, the RL system updates the agent according to the returned reward. The RL loop iterates on episodes until the agent converges.

The RL optimization can be divided into value-based (*e.g.* DQN (Mnih et al. 2013) and DDQN (Van Hasselt, Guez, and Silver 2016)) and policy-based (*e.g.* PG (Sutton et al. 1999) and PPO (Schulman et al. 2017)) algorithms. Value-based RL often selects an action with the maximum Q-value. Policy-based RL constructs a probability distribution

to sample an action. Compared with value-based RL, policy-based RL has two advantages: (1) it is more compatible with continuous action space; (2) it avoids policy degradation since it does not have value function errors in value-based RL. Based on the above analyses, policy-based RL is more suitable for optimizing continuous multi-teacher weights as action space. Policy Gradient (PG) (Sutton et al. 1999) uses reward from the observed state to optimize the action policy. In the conventional RL system, the original PG may be difficult to converge due to the large gradient variance. Some advanced variants are proposed to improve PG based on actor-critic (e.g. DPG (Silver et al. 2014) and DDPG (Lillicrap et al. 2015)) and trust region (e.g. PPO (Schulman et al. 2017)).

As shown in Table 6, we found that the multi-teacher decision is not sensitive to PG algorithms. One possible reason is that the gradient optimization of multi-teacher decisions with a fixed visual dataset is more stable than the conventional RL environment. Therefore, we choose the original PG to optimize multi-teacher weights. Notice that this paper focuses on a unified framework of multi-teacher KD with RL instead of proposing a new RL algorithm.

## Methodology

### Preliminary of Multi-Teacher KD

The standard KD transfers knowledge from a teacher network  $T$  to a student network  $S$ . The framework is often formulated as Equ.(1), including a basic task loss, logit KD loss, and feature KD loss.

$$\mathcal{L}_{KD} = \underbrace{\mathcal{H}(\mathbf{y}_i^S, \mathbf{y}_i)}_{\text{task loss}} + \alpha \underbrace{\mathcal{D}_{KL}(\mathbf{y}_i^S, \mathbf{y}_i^T)}_{\text{logit KD loss}} + \beta \underbrace{\mathcal{D}_{dis}(\mathbf{F}_i^S, \mathbf{F}_i^T)}_{\text{feature KD loss}}. \quad (1)$$

$\mathbf{y}_i$  is the ground-truth label of the  $i$ -th sample.  $\mathbf{y}_i^S$  and  $\mathbf{y}_i^T$  are predictive logits from the student and teacher, respectively.  $\mathbf{F}_i^S$  and  $\mathbf{F}_i^T$  are features from the student and teacher, respectively.  $\mathcal{H}$  denotes cross-entropy function.  $\mathcal{D}_{KL}$  indicates Kullback-Leibler divergence to measure the discrepancy of student and teacher probability distributions.  $\mathcal{D}_{dis}$  is a distance metric to measure the similarity of student and teacher features.  $\alpha$  and  $\beta$  are loss weights for logit KD and feature KD losses, respectively.

When extending to multi-teacher KD, each teacher network transfers knowledge to the same student network, resulting in the total loss as Equ.(2).

$$\begin{aligned} \mathcal{L}_{MTKD} = & \underbrace{\mathcal{H}(\mathbf{y}_i^S, \mathbf{y}_i)}_{\text{task loss}} + \alpha \underbrace{\sum_{m=1}^M w_{l,i}^m \mathcal{D}_{KL}(\mathbf{y}_i^S, \mathbf{y}_i^{T_m})}_{\text{logit KD loss}} \\ & + \beta \underbrace{\sum_{m=1}^M w_{f,i}^m \mathcal{D}_{dis}(\mathbf{F}_i^S, \mathbf{F}_i^{T_m})}_{\text{feature KD loss}}. \end{aligned} \quad (2)$$

Here,  $M$  is the number of networks in the teacher pool  $\{T_m\}_{m=1}^M$ . This paper focuses on optimizing the teacher weights  $\{\mathbf{w}_i^m = [w_{l,i}^m, w_{f,i}^m]\}_{m=1}^M$  using reinforcement

learning, as shown in the next section. Without bells and whistles, we use the traditional KD (Hinton, Vinyals, and Dean 2015) as the logit KD loss and FitNet (Romero et al. 2014) as the feature KD loss by default. We set  $\alpha = 1$  and  $\beta = 5$ , as analysed in Fig.2. More advanced logit KD losses (e.g. DIST (Huang et al. 2022)) and feature KD losses (e.g. ND (Wang et al. 2025)) could also be applied in Equ.(2), which are orthogonal to our multi-teacher KD framework.

### RL-based Multi-Teacher Weights Optimization

In this section, we propose to optimize sample-wise teacher weights dynamically using reinforcement learning. As shown in Fig.1, we introduce an agent over the training graph. The agent interacts with the multi-teacher KD environment. It receives the state information to generate teacher weights as the action. We apply teacher weights into Equ.(2) to distill the student network. After an episode of training samples, the student performance is regarded as the reward to optimize agent parameters. The reinforced loop proceeds until the student network converges. The instantiation of (state, action, reward) in the context of multi-teacher KD is shown as follows.

**State** In our RL-based framework, the state representation of the input sample  $x_i$  is formulated as an embedding  $s_i$ . The state  $s_i$  is a concatenation of teacher performance (i.e. feature representation, logit vector, and cross-entropy loss) and teacher-student gaps (i.e. feature similarity and probability KL-divergence). The concrete details are formulated as follows.

**(1) Teacher feature representation.** Given an input sample  $x_i$ , the  $m$ -th teacher's feature representation after the penultimate layer is defined as  $\mathbf{f}_i^{T_m} \in \mathbb{R}^{d_m}$ , where  $d_m$  is the embedding size of the  $m$ -th teacher. The feature representation often encodes semantic information from the input sample.

**(2) Teacher logit vector.** Given an input sample  $x_i$ , the  $m$ -th teacher's logit vector after the final output layer is defined as  $\mathbf{z}_i^{T_m} \in \mathbb{R}^C$ , where  $C$  is the number of classes. The logit vector represents the direct prediction of a class confidence distribution over the complete class space.

**(3) Teacher cross-entropy loss.** Given an input sample  $x_i$  with the label  $\mathbf{y}_i$ , the  $m$ -th teacher's cross-entropy loss is defined as  $\mathcal{L}_{CE}^{T_m} = \mathcal{H}(\mathbf{y}_i^{T_m}, \mathbf{y}_i)$ . The cross-entropy loss  $\mathcal{L}_{CE}^{T_m}$  measures the fitting performance between the teacher's prediction and the ground-truth label.

**(4) Teacher-student feature similarity.** Given an input sample  $x_i$ , the cosine similarity between the student's feature representation  $\mathbf{f}_i^S \in \mathbb{R}^d$  and the teacher's feature representation  $\mathbf{f}_i^{T_m} \in \mathbb{R}^{d_m}$  is formulated as cosine similarity:

$$\cos_i^{T_m} = r(\mathbf{f}_i^S) \cdot \mathbf{f}_i^{T_m} / \|r(\mathbf{f}_i^S)\| \|\mathbf{f}_i^{T_m}\|, \quad (3)$$

where  $r(\cdot)$  is a linear regressor to transform  $\mathbf{f}_i^S$  to match the embedding size of  $\mathbf{f}_i^{T_m}$ .  $\cdot$  is a dot product and  $\|\cdot\|$  denotes  $l_2$  norm. The cosine similarity measures the distance gap between the student and teacher in feature space.

**(5) Teacher-student probability KL-divergence.** Given an input sample  $x_i$ , the KL-divergence between the student's logit vector  $\mathbf{y}_i^S \in \mathbb{R}^C$  and the teacher's logit vector

$\mathbf{y}_i^{T_m} \in \mathbb{R}^C$  is formulated as  $KL_i^{T_m} = \mathcal{D}_{KL}(\mathbf{y}_i^S, \mathbf{y}_i^{T_m})$ . The KL-divergence measures the discrepancy between the student and teacher in class probability space.

We concatenate (1)~(5) to construct the state embedding. Given an input sample  $\mathbf{x}_i$ , the  $m$ -th teacher's state embedding is formulated as  $\mathbf{s}_i^m$ :

$$\mathbf{s}_i^m = [ \underbrace{\mathbf{f}_i^{T_m} \parallel \mathbf{z}_i^{T_m} \parallel \mathcal{L}_{CE}^{T_m}}_{\text{Teacher Performance}} \parallel \underbrace{\cos_i^{T_m} \parallel KL_i^{T_m}}_{\text{Teacher-Student Gaps}} ], \quad (4)$$

where  $\parallel$  denotes the concatenation operator. Based on comprehensive state information, the agent can make dynamic decisions by referring to teacher performance and teacher-student gaps among various input samples.

**Action** We construct an agent  $\pi_{\theta_m}(\mathbf{s}_i^m)$  for each teacher network  $T_m$ , where  $\theta_m$  denotes the trainable agent parameters. The agent model  $\pi_{\theta_m}(\mathbf{s}_i^m)$  contains several linear layers with a middle ReLU activation function, and finished by a softmax function. The output is the teacher weight vector  $\mathbf{w}_i^m$ , each weight is within  $(0, 1)$  in continuous action space. The process is formulated as  $\mathbf{w}_i^m = \pi_{\theta_m}(\mathbf{s}_i^m)$ . We also combine the confidence- and divergence-aware weight generation strategies in the action space. The detailed action construction is shown in appendix.

**Reward** We define an episode as one training batch  $\{\mathbf{x}_i\}_{i=1}^B$  including  $B$  samples. For each sample  $\mathbf{x}_i$ , we construct the state embedding  $\mathbf{s}_i^m$  for each teacher network  $T_m$  and guide the agent  $\pi_{\theta_m}(\mathbf{s}_i^m)$  to produce the teacher weight  $\mathbf{w}_i^m$ . We perform weighted multi-teacher KD (Equ.(2)) with the generated weights  $\{\mathbf{w}_i^m\}_{m=1}^M$  to train the student network. The reward function should be relevant to the trained student performance after multi-teacher KD weighted by the agent. A better student network should have a lower classification loss and more similar class probability and feature distributions with the teacher. Therefore, we utilize cross-entropy loss between the student and ground-truth labels, probability KL-divergence loss and feature Mean Squared Error (MSE) loss between the student and teacher as evidence to construct the reward function as Equ.(5):

$$R_i^m = -\mathcal{H}(\mathbf{y}_i^S, \mathbf{y}_i) - \alpha \mathcal{D}_{KL}(\mathbf{y}_i^S, \mathbf{y}_i^{T_m}) - \beta \mathcal{D}_{dis}(\mathbf{F}_i^S, \mathbf{F}_i^{T_m}), \quad (5)$$

where  $R_i^m$  denotes the reward of the  $m$ -th agent over the  $i$ -th input sample. Here, the reward is defined by the negative loss, since a lower loss value indicates a better student.

**Optimize agent with policy gradient** Unlike traditional supervised learning, RL-based optimization often does not know whether the chosen action is correct or not. Therefore, policy-based RL often optimizes the agent according to the returned reward after making a decision. If an agent earns a larger reward after choosing an action, Policy Gradient (PG) (Sutton et al. 1999) would increase the corresponding gradient derived from this action. We apply PG to optimize the agent's parameters, as shown in Equ.(6).

$$\theta_m \leftarrow \theta_m - \eta \sum_{i=1}^B \bar{R}_i^m \nabla_{\theta_m} \pi_{\theta_m}(\mathbf{s}_i^m), \quad m = 1, 2, \dots, M. \quad (6)$$

---

#### Algorithm 1: Overall MTKD-RL Training Procedure

---

1. Pre-train the student network  $S$  by multi-teacher KD following  $\mathcal{L}_{MTKD}$  (Equ.(2)) with equal weights, i.e.  $\{\mathbf{w}_i^m = \mathbf{1}\}_{m=1}^M$ , using all training samples  $\mathbf{x}_i \in \mathcal{D}$ .
  2. Pre-train the agent models  $\{\pi_{\theta_m}\}_{m=1}^M$  using the returned rewards  $\{R_i^m\}_{m=1}^M$  with  $\{\mathbf{w}_i^m = \mathbf{1}\}_{m=1}^M$  as the action for all training samples  $\mathbf{x}_i \in \mathcal{D}$ .
  3. Run Algorithm 2 to alternatively perform multi-teacher KD and agent optimization until convergence.
- 

---

#### Algorithm 2: Alternative Multi-Teacher KD and Agent Optimization

---

**Input:** Training dataset  $\mathcal{D}$ . Pre-trained teacher networks  $\{T_m\}_{m=1}^M$ . A student network  $S$ . Agent models  $\{\pi_{\theta_m}\}_{m=1}^M$ .  
**Output:** Trained student  $S$ .

- 1: **while** the student network  $S$  is not converged **do**
  - 2: Randomly Shuffle  $\mathcal{D}$  to produce a new batch sequence.
  - 3: **for** each batch  $\mathcal{B} \in \mathcal{D}$  **do**
  - 4: Construct state information  $\{\mathbf{s}_i^m\}_{m=1}^M$  following Equ.(4) for each sample  $\mathbf{x}_i \in \mathcal{B}$ , where the batch size is  $B$ .
  - 5: Generate multi-teacher weights  $\{\mathbf{w}_i^m\}_{m=1}^M$  by  $\mathbf{w}_i^m = \pi_{\theta_m}(\mathbf{s}_i^m)$
  - 6: Compute the multi-teacher KD loss  $\mathcal{L}_{MTKD}$  following Equ.(2) with generated weights  $\{\mathbf{w}_i^m\}_{m=1}^M$ .
  - 7: Update the student network  $S$  using  $\mathcal{L}_{MTKD}$ .
  - 8: Compute rewards  $\{R_i^m\}_{m=1}^M$  following Equ.(5).
  - 9: Save  $(\{\mathbf{s}_i^m\}_{m=1}^M, \{\mathbf{w}_i^m\}_{m=1}^M, \{R_i^m\}_{m=1}^M)_{i=1}^B$  to the episode history  $\mathcal{H}$ .
  - 10: **end for**
  - 11: **for** each  $(\{\mathbf{s}_i^m\}_{m=1}^M, \{\mathbf{w}_i^m\}_{m=1}^M, \{R_i^m\}_{m=1}^M)_{i=1}^B \in \mathcal{H}$  **do**
  - 12: Update the agent models  $\{\pi_{\theta_m}\}_{m=1}^M$  following Equ.(6).
  - 13: **end for**
  - 14: **end while**
- 

$\bar{R}_i^m$  is the normalized reward formulated as Equ.(7).

$$\bar{R}_i^m = \frac{R_i^m - \min_k R_i^k}{\max_k R_i^k - \min_k R_i^k} - \frac{1}{M} \sum_{k=1}^M R_i^k. \quad (7)$$

Here,  $k \in \{1, \dots, M\}$ . We apply min-max normalization to rescale the reward  $R_i^m$  to  $(0, 1)$ , and then subtract the mean value of  $M$  teacher rewards  $\{R_i^k\}_{k=1}^M$ . The PG would lead to a positive update when  $\bar{R}_i^m > 0$ , or a negative update when  $\bar{R}_i^m < 0$ . Benefiting from PG optimization, the agent prefers adaptively assigning teacher weights according to the reward.

#### Overall MTKD-RL Framework

Algorithm 1 illustrates the overall MTKD-RL training procedure. At first, we pre-train the student network  $S$  by multi-teacher KD following  $\mathcal{L}_{MTKD}$  (Equ.(2)) with equal weights for one training epoch. Afterwards, we use the collected (state, action, reward) information during multi-teacher KD

to pre-train the agent models  $\{\pi_{\theta_m}\}_{m=1}^M$ . After pre-training, the RL-loop would start from a good initialization and avoid the training collapse problem. Then we run Algorithm 2 to alternatively perform multi-teacher KD and agent optimization until convergence.

As shown in Algorithm 2, for each iterative epoch, we freeze the agent models  $\{\pi_{\theta_m}\}_{m=1}^M$  and train the student network  $S$  using multi-teacher KD with the generated weights from agent models. Then, we utilize the collected (state, action, reward) information during multi-teacher KD to optimize the agent models  $\{\pi_{\theta_m}\}_{m=1}^M$ . The alternative training proceeds until the student network  $S$  is converged.

## Experiments

### Experimental Setup

**Dataset.** We use CIFAR-100 (Krizhevsky, Hinton et al. 2009) and ImageNet (Deng et al. 2009) datasets for image classification experiments. The object detection experiments adopt COCO-2017 (Lin et al. 2014) dataset. We utilize Cityscapes (Cordts et al. 2016), ADE20K (Zhou et al. 2017) and COCO-Stuff-164K (Caesar, Uijlings, and Ferrari 2018) datasets for semantic segmentation.

**Compared methods.** We compare several methods in this experimental section. ‘Baseline’ denotes the traditional training without distillation. ‘AVER (KD+FitNet)’ means the multi-teacher KD with equal weights for the traditional Hinton’s logit KD loss (Hinton, Vinyals, and Dean 2015) and Mean Squared Error (MSE)-based feature KD loss (Romero et al. 2014). Other compared methods, such as AMTML-KD (Liu, Zhang, and Wang 2020), AEKD (Du et al. 2020), CA-MKD (Zhang, Chen, and Wang 2022), and MMKD (Zhang, Chen, and Wang 2023), are advanced strategies to compute multi-teacher KD weights.

**Evaluated networks.** Experiments are conducted over Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs). CNNs includes ResNet (He et al. 2016), WRN (Zagoruyko S 2016), ResNeXt (Xie et al. 2017), MobileNetV2 (Sandler et al. 2018), ShuffleNetV2 (Ma et al. 2018), and RegNet (Radosavovic et al. 2020). ViTs includes CaiT (Touvron et al. 2021b) and DeiT (Touvron et al. 2021a). *More training details are shown in appendix.*

### Experimental Results on Image Classification

**Experiments on CIFAR-100.** In Table 1, we conduct multi-teacher KD experiments on CIFAR-100. Compared to the baseline without distillation, even the vanilla AVER (KD) with equal weights achieves significant improvements over four student networks with an average gain of 2.46%. AVER (KD+FitNet) complements feature-level distillation and obtains an average improvement of 0.68% over the logit-level AVER (KD). Previous multi-teacher weight generation strategies, such as AMTML-KD, AEKD, CA-MKD, and MMKD, generally outperform the AVER version, indicating that dynamic sample-wise weights in a data-driven manner are better than simple static weights. Our MTKD-RL achieves the best performance and surpasses the state-of-the-art CA-MKD and MMKD methods with average gains of 0.31% and 0.33%, respectively. The results demonstrate

Network	Params	FLOPs	Method	Acc@1
Teacher pool				
RegNetY-400MF	5.8M	467.1M	Pretrained	78.87
RegNetX-400MF	4.8M	466.2M		79.15
ResNet-32x4	7.4M	1083.0M		79.59
WRN-28-4	5.9M	845.6M		79.37
Student				
RegNetX-200MF	2.4M	223.4M	Baseline	77.38 $\pm$ 0.23
			AVER (KD)	78.64 $\pm$ 0.31
			AVER (KD+FitNet)	79.12 $\pm$ 0.28
			AMTML-KD	79.46 $\pm$ 0.45
			AEKD	79.22 $\pm$ 0.37
			CA-MKD	80.28 $\pm$ 0.26
			MMKD	80.15 $\pm$ 0.41
			MTKD-RL (Ours)	<b>80.58</b> $\pm$ 0.29
MobileNetV2	2.4M	22.4M	Baseline	69.17 $\pm$ 0.18
			AVER (KD)	71.85 $\pm$ 0.27
			AVER (KD+FitNet)	72.67 $\pm$ 0.26
			AMTML-KD	72.89 $\pm$ 0.36
			AEKD	72.82 $\pm$ 0.42
			CA-MKD	74.16 $\pm$ 0.32
			MMKD	74.35 $\pm$ 0.34
			MTKD-RL (Ours)	<b>74.63</b> $\pm$ 0.35
ShuffleNetv2	1.4M	44.5M	Baseline	72.84 $\pm$ 0.15
			AVER (KD)	75.77 $\pm$ 0.18
			AVER (KD+FitNet)	76.83 $\pm$ 0.34
			AMTML-KD	76.93 $\pm$ 0.26
			AEKD	77.08 $\pm$ 0.21
			CA-MKD	78.09 $\pm$ 0.17
			MMKD	77.87 $\pm$ 0.20
			MTKD-RL (Ours)	<b>78.39</b> $\pm$ 0.14
ResNet-56	0.9M	125.8M	Baseline	72.52 $\pm$ 0.27
			AVER (KD)	73.57 $\pm$ 0.29
			AVER (KD+FitNet)	73.93 $\pm$ 0.16
			AMTML-KD	74.21 $\pm$ 0.33
			AEKD	74.02 $\pm$ 0.39
			CA-MKD	75.17 $\pm$ 0.24
			MMKD	75.26 $\pm$ 0.27
			MTKD-RL (Ours)	<b>75.35</b> $\pm$ 0.33

Table 1: **Comparison of accuracy among various multi-teacher KD methods over CNNs on CIFAR-100.** We use the pre-trained teacher pool to distill student networks. The **bold** and underline numbers denotes the best and second-best results for each student. Experiments perform 3 runs.

that our proposed reinforcement learning framework is better than the entropy-based and meta-learning-based mechanisms for multi-teacher weight optimization.

**Experiments on ImageNet.** As shown in Table 2, we conduct multi-teacher KD experiments over CNNs on ImageNet. MTKD-RL enhances the baseline by 2.47% and 3.13% accuracy gains on ResNet-18 and ResNet-34, respectively. It outperforms the AVER (KD+FitNet) with an average gain of 1.24%, manifesting that our method generates meaningful multi-teacher weights on large-scale ImageNet. Compared to state-of-the-art MMKD, MTKD-RL achieves 0.49% and 0.71% accuracy improvements on ResNet-18 and ResNet-34, respectively. The results show the superiority of MTKD-RL for applying to the large-scale dataset compared to other multi-teacher KD strategies. As shown in Table 2, we further conduct multi-teacher KD experiments over ViTs. MTKD-RL increases the baseline by 2.91% and 1.90% on DeiT-Tiny and CaiT-XXS24, respectively. It also exceeds the best competitor MMKD by an average gain of

Network	Params	FLOPs	Method	Acc@1
Teacher pool				
ResNet-50	25.6M	4.1G	Pretrained	76.13
ResNet-101	44.5M	7.8G		77.37
Wide ResNet-50-2	68.9M	11.4G		78.47
ResNeXt-50 (32×4d)	25.0M	4.2G		77.62
Student				
ResNet-18	11.7M	1.8G	Baseline	70.35
			AVER (KD)	71.52
			AVER (KD+FitNet)	71.56
			AMTML-KD	71.89
			AEKD	71.67
			CA-MKD	72.38
			MMKD	72.33
		MTKD-RL (Ours)	<b>72.82</b>	
ResNet-34	21.8M	3.7G	Baseline	73.64
			AVER (KD)	75.32
			AVER (KD+FitNet)	75.55
			AMTML-KD	75.68
			AEKD	75.66
			CA-MKD	75.87
			MMKD	76.06
		MTKD-RL (Ours)	<b>76.77</b>	

Table 2: Comparison of accuracy among various multi-teacher KD methods over CNNs on ImageNet.

Network	Params	FLOPs	Method	Acc@1
Teacher pool				
CaiT-S24	46.9M	9.4G	Pretrained	83.36
DeiT-Small	22.1M	4.6G		79.82
DeiT-Base	86.6M	17.5G		81.80
Student				
DeiT-Tiny	5.7M	1.3G	Baseline	72.23
			AVER (KD)	73.87
			AVER (KD+FitNet)	73.98
			AMTML-KD	73.68
			AEKD	73.72
			CA-MKD	74.12
			MMKD	<u>74.35</u>
	MTKD-RL (Ours)	<b>75.14</b>		
CaiT-XXS24	12.0M	2.5G	Baseline	77.32
			AVER	78.36
			AVER (KD+FitNet)	78.27
			AMTML-KD	78.57
			AEKD	78.44
			CA-MKD	<u>78.65</u>
			MMKD	78.42
	MTKD-RL (Ours)	<b>79.22</b>		

Table 3: Comparison of accuracy among various multi-teacher KD methods over ViT on ImageNet.

0.80% on the two networks. The results reveal that MTKD-RL can also work well on ViT.

## Experimental Results on Object Detection

As shown in Table 4, we transfer ImageNet-pretrained ResNet backbones for downstream object detection on COCO-2017. The ResNet backbones pretrained by our MTKD-RL obtain consistent performance improvements over various detectors (Mask-RCNN (He et al. 2017), Cascade-RCNN (Cai and Vasconcelos 2018), Reti-

Backbone	Detector	Method	mAP
ResNet-18	Mask-RCNN	Baseline	34.1
		MTKD-RL (Ours)	<b>35.1</b>
	Cascade-RCNN	Baseline	36.5
		MTKD-RL (Ours)	<b>37.7</b>
	RetinaNet	Baseline	31.6
		MTKD-RL (Ours)	<b>32.7</b>
	Faster-RCNN	Baseline	33.5
		MTKD-RL (Ours)	<b>34.7</b>
ResNet-34	Mask-RCNN	Baseline	37.6
		MTKD-RL (Ours)	<b>39.0</b>
	Cascade-RCNN	Baseline	39.6
		MTKD-RL (Ours)	<b>40.8</b>
	RetinaNet	Baseline	35.2
		MTKD-RL (Ours)	<b>36.9</b>
	Faster-RCNN	Baseline	37.0
		MTKD-RL (Ours)	<b>38.5</b>

Table 4: Comparison of downstream object detection based on ImageNet-pretrained ResNet backbones over various detectors on COCO-2017.

Segmentor	Method	Cityscapes	ADE20K	COCO-Stuff
DeepLabV3	Baseline	76.34	36.08	29.97
	MTKD-RL (Ours)	<b>77.42</b>	<b>37.07</b>	<b>31.82</b>
PSPNet	Baseline	74.60	36.84	31.25
	MTKD-RL (Ours)	<b>75.89</b>	<b>37.78</b>	<b>32.39</b>

Table 5: Comparison of downstream semantic segmentation tasks based on an ImageNet-pretrained ResNet-34 on Cityscapes, ADE20K and COCO-Stuff-164K.

naNet (Lin et al. 2017) and Faster-RCNN (Ren et al. 2016)) than the baseline backbones. MTKD-RL outperforms the baseline with 1.1% and 1.5% mAP gains on average for ResNet-18 and ResNet-34, respectively. The results indicate that MTKD-RL can guide the network to learn better feature representations for downstream object detection.

## Experimental Results on Semantic Segmentation

As shown in Table 5, we transfer the ImageNet-pretrained ResNet-34 backbone for downstream semantic segmentation datasets. Our MTKD-RL consistently surpasses the baseline across various datasets by equipping with DeepLabV3 (Chen et al. 2018) and PSPNet (Zhao et al. 2017) segmentation heads. The average improvements are 1.19%, 0.97% and 1.50% on Cityscapes, ADE20K and COCO-Stuff-164K datasets, respectively. The results show that the backbone pretrained by MTKD-RL can be well extended to semantic segmentation tasks and generate better pixel-wise feature representations.

## Ablation Study and Parameter Analysis

**Ablation study of various RL methods.** As shown in Table 6(a), we investigate various RL methods for multi-




Figure 2: Parameter analyses and ablation study over ShuffleNetV2 on CIFAR-100.

RL methods	Acc	Component	Acc	Method	Time	Mem	Acc
PG	<b>78.39</b> $\pm 0.14$	AVER	76.83 $\pm 0.34$	Baseline	29s	2.3G	72.84 $\pm 0.15$
DPG	78.16 $\pm 0.11$	$f_i^{T_m} \parallel z_i^{T_m} \parallel \mathcal{L}_{CE}^{T_m}$	77.84 $\pm 0.17$	AVER	41s	2.8G	76.83 $\pm 0.34$
DDPG	78.05 $\pm 0.22$	$\cos \parallel \mathcal{D}_{KL}$	77.46 $\pm 0.21$	CA-MKD	54s	2.9G	78.09 $\pm 0.17$
PPO	78.28 $\pm 0.08$	All	<b>78.39</b> $\pm 0.14$	MTKD-RL	47s	3.2G	<b>78.39</b> $\pm 0.14$
(a) RL methods		(b) Component analysis		(c) Training costs			

Table 6: Analysis of MTKD-RL over ShuffleNetV2 on CIFAR-100.

teacher weight optimization. We find that MTKD-RL is not sensitive to RL optimization methods. Therefore, we use the original PG due to its simplicity and effectiveness.

**Component analysis for RL-based optimization.** In Table 6(b), we analyse various components to construct state and action information. Using the teacher performance, including feature  $f_i^{T_m}$ , logit  $z_i^{T_m}$ , and cross-entropy loss  $\mathcal{L}_{CE}^{T_m}$ , leads to 1.01% gain over AVER. Using the teacher-student gaps, including feature similarity  $\cos(r(f_i^S), f_i^{T_m})$  and logit KL-divergence  $\mathcal{D}_{KL}(y_i^S, y_i^{T_m})$ , results in 0.63% gain. Gathering the comprehensive knowledge of teacher performance and teacher-student gaps for RL-based optimization achieves the best 1.56% improvement.

**Training costs.** As shown in Table 6(c), we compare training complexity in one epoch with other methods over NVIDIA RTX 4090. Compared to the AVER without multi-teacher weight optimization, MTKD-RL needs extra 15% time and 14% memory costs, because our method requires agent training and storing (state, action, reward) as episode history. In contrast to state-of-the-art CA-MKD, MTKD-RL uses 13% less time and extra 10% memory footprint but achieves a better accuracy gain. In summary, our method leads to significant performance improvements with little increase in training complexity.

**Parameter analyses of loss weights.** As shown in Fig.2, we investigate the impact of logit loss weight  $\alpha$  and feature loss weight  $\beta$ , where  $\alpha = 1$  and  $\beta = 5$  achieve the best performance.

**Impact of the number of teacher networks.** Fig.2 shows the performance curve caused by the number of teacher networks  $M$ . As the teacher number  $M$  increases, the performance generally improves, since more teachers could provide richer ensemble knowledge but saturates at a certain

Method	Baseline	MTKD-RL	+DIST	+ND
Acc@1	70.35	72.82	73.29	<b>73.61</b>

Table 7: Combined with single-teacher KD methods to distill ResNet-18 on ImageNet.

capacity of  $M = 4$ .

**Ablation study of loss terms.** As shown in Fig.2, we conduct an ablation study of loss terms in MTKD-RL. Compared to the baseline without distillation, logit-level and feature-level KD lead to 4.14% and 3.61% accuracy improvements, respectively. One of the reasons may be that logit-level KD is architecture-unaware distillation and more robust than feature-level KD. Summarizing the logit-level and feature-level KD provides comprehensive information for the student, maximizing the accuracy gain by 5.55%.

**Combined with single-teacher KD.** In Table 7, Combining MTKD-RL with advanced single-teacher KD methods of DIST (Huang et al. 2022) and ND (Wang et al. 2025) further achieves 0.47% and 0.79% accuracy gains, respectively.

## Conclusion

This paper proposes MTKD-RL, formulating multi-teacher KD as an RL-based decision process. Compared to existing works, our method uses more comprehensive *teacher performance* and *teacher-student gaps* to construct the input evidence. It reinforces the interaction between teacher and student by an agent. Experimental results on visual recognition tasks show that MTKD-RL achieves state-of-the-art performance among existing multi-teacher KD methods. We hope our work could inspire future research to explore more advanced RL strategies for multi-teacher KD.



## Acknowledgments

This work is partially supported by the National Natural Science Foundation of China (No.62476264 and No.62406312), China National Postdoctoral Program for Innovative Talents (No.BX20240385) funded by China Postdoctoral Science Foundation, Beijing Natural Science Foundation (No.4244098), and Science Foundation of the Chinese Academy of Sciences.

## References

- Caesar, H.; Uijlings, J.; and Ferrari, V. 2018. Coco-stuff: Thing and stuff classes in context. In *CVPR*, 1209–1218.
- Cai, Z.; and Vasconcelos, N. 2018. Cascade r-cnn: Delving into high quality object detection. In *CVPR*, 6154–6162.
- Chen, L.-C.; Zhu, Y.; Papandreou, G.; Schroff, F.; and Adam, H. 2018. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 801–818.
- Chengqing, Y.; Guangxi, Y.; Chengming, Y.; Yu, Z.; and Xiwei, M. 2023. A multi-factor driven spatiotemporal wind power prediction model based on ensemble deep graph attention reinforcement learning networks. *Energy*, 263: 126034.
- Cho, J. H.; and Hariharan, B. 2019. On the efficacy of knowledge distillation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4794–4802.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3213–3223.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *IEEE conference on computer vision and pattern recognition*, 248–255. IEEE.
- Du, S.; You, S.; Li, X.; Wu, J.; Wang, F.; Qian, C.; and Zhang, C. 2020. Agree to disagree: Adaptive ensemble knowledge distillation in gradient space. *advances in neural information processing systems*, 33: 12345–12355.
- Feng, W.; Yang, C.; An, Z.; Huang, L.; Diao, B.; Wang, F.; and Xu, Y. 2024. Relational diffusion distillation for efficient image generation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 205–213.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Huang, T.; You, S.; Wang, F.; Qian, C.; and Xu, C. 2022. Knowledge distillation from a stronger teacher. *Advances in Neural Information Processing Systems*, 35: 33716–33727.
- Kaelbling, L. P.; Littman, M. L.; and Moore, A. W. 1996. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4: 237–285.
- Kang, J.; Chen, X.; Wang, J.; Hu, C.; Liu, X.; and Dudek, G. 2022. MOBA: Multi-teacher Model Based Reinforcement Learning. *openreview.net*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. *Technical Report*.
- Kwon, K.; Na, H.; Lee, H.; and Kim, N. S. 2020. Adaptive knowledge distillation based on entropy. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7409–7413. IEEE.
- Li, L.; Bao, Y.; Dong, P.; Yang, C.; Li, A.; Luo, W.; Liu, Q.; Xue, W.; and Guo, Y. 2024. DetKDS: Knowledge Distillation Search for Object Detectors. In *Forty-first International Conference on Machine Learning*.
- Liang, K.; Meng, L.; Liu, M.; Liu, Y.; Tu, W.; Wang, S.; Zhou, S.; Liu, X.; Sun, F.; and He, K. 2024a. A survey of knowledge graph reasoning on graph types: Static, dynamic, and multi-modal. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Liang, K.; Meng, L.; Liu, Y.; Liu, M.; Wei, W.; Liu, S.; Tu, W.; Wang, S.; Zhou, S.; and Liu, X. 2024b. Simple Yet Effective: Structure Guided Pre-trained Transformer for Multi-modal Knowledge Graph Reasoning. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 1554–1563.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Liu, H.; Yu, C.; Wu, H.; Duan, Z.; and Yan, G. 2020. A new hybrid ensemble deep reinforcement learning model for wind speed short term forecasting. *Energy*, 202: 117794.
- Liu, Y.; Zhang, W.; and Wang, J. 2020. Adaptive multi-teacher multi-level knowledge distillation. *Neurocomputing*, 415: 106–113.
- Ma, N.; Zhang, X.; Zheng, H.-T.; and Sun, J. 2018. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European conference on computer vision (ECCV)*, 116–131.



- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Radosavovic, I.; Kosaraju, R. P.; Girshick, R.; He, K.; and Dollár, P. 2020. Designing network design spaces. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10428–10436.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2016. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6): 1137–1149.
- Romero, A.; Ballas, N.; Kahou, S. E.; Chassang, A.; Gatta, C.; and Bengio, Y. 2014. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*.
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4510–4520.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Silver, D.; Lever, G.; Heess, N.; Degris, T.; Wierstra, D.; and Riedmiller, M. 2014. Deterministic policy gradient algorithms. In *International conference on machine learning*, 387–395. Pmlr.
- Sutton, R. S.; McAllester, D.; Singh, S.; and Mansour, Y. 1999. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12.
- Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; and Jégou, H. 2021a. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, 10347–10357. PMLR.
- Touvron, H.; Cord, M.; Sablayrolles, A.; Synnaeve, G.; and Jégou, H. 2021b. Going deeper with image transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 32–42.
- Van Hasselt, H.; Guez, A.; and Silver, D. 2016. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30.
- Wang, Y.; Cheng, L.; Duan, M.; Wang, Y.; Feng, Z.; and Kong, S. 2025. Improving Knowledge Distillation via Regularizing Feature Direction and Norm. In *European Conference on Computer Vision*, 20–37. Springer.
- Xie, S.; Girshick, R.; Dollár, P.; Tu, Z.; and He, K. 2017. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1492–1500.
- Yang, C.; An, Z.; Cai, L.; and Xu, Y. 2021. Hierarchical self-supervised augmented knowledge distillation. In *International Joint Conference on Artificial Intelligence*, 1217–1223.
- Yang, C.; An, Z.; Cai, L.; and Xu, Y. 2022a. Knowledge distillation using hierarchical self-supervision augmented distribution. *IEEE Transactions on Neural Networks and Learning Systems*, 35(2): 2094–2108.
- Yang, C.; An, Z.; Huang, L.; Bi, J.; Yu, X.; Yang, H.; Diao, B.; and Xu, Y. 2024. CLIP-KD: An Empirical Study of CLIP Model Distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15952–15962.
- Yang, C.; An, Z.; Zhou, H.; Cai, L.; Zhi, X.; Wu, J.; Xu, Y.; and Zhang, Q. 2022b. Mixskd: Self-knowledge distillation from mixup for image recognition. In *European Conference on Computer Vision*, 534–551. Springer.
- Yang, C.; An, Z.; Zhou, H.; Zhuang, F.; Xu, Y.; and Zhang, Q. 2023a. Online knowledge distillation via mutual contrastive learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8): 10212–10227.
- Yang, C.; An, Z.; Zhu, H.; Hu, X.; Zhang, K.; Xu, K.; Li, C.; and Xu, Y. 2020. Gated convolutional networks with hybrid connectivity for image classification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 12581–12588.
- Yang, C.; Yu, X.; An, Z.; and Xu, Y. 2023b. Categories of response-based, feature-based, and relation-based knowledge distillation. In *Advancements in Knowledge Distillation: Towards New Horizons of Intelligent Systems*, 1–32. Springer.
- Yang, C.; Zhou, H.; An, Z.; Jiang, X.; Xu, Y.; and Zhang, Q. 2022c. Cross-image relational knowledge distillation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12319–12328.
- You, S.; Xu, C.; Xu, C.; and Tao, D. 2017. Learning from multiple teacher networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1285–1294.
- Yuan, F.; Shou, L.; Pei, J.; Lin, W.; Gong, M.; Fu, Y.; and Jiang, D. 2021. Reinforced multi-teacher selection for knowledge distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 14284–14291.
- Zagoruyko S, K. N. 2016. Wide residual networks. In *Proceedings of the British Machine Vision Conference*.
- Zhang, H.; Chen, D.; and Wang, C. 2022. Confidence-aware multi-teacher knowledge distillation. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4498–4502. IEEE.
- Zhang, H.; Chen, D.; and Wang, C. 2023. Adaptive Multi-Teacher Knowledge Distillation with Meta-Learning. *arXiv preprint arXiv:2306.06634*.
- Zhang, T.; Xue, M.; Zhang, J.; Zhang, H.; Wang, Y.; Cheng, L.; Song, J.; and Song, M. 2023. Generalization Matters: Loss Minima Flattening via Parameter Hybridization for Efficient Online Knowledge Distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20176–20185.
- Zhao, H.; Shi, J.; Qi, X.; Wang, X.; and Jia, J. 2017. Pyramid scene parsing network. In *CVPR*, 2881–2890.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. In *CVPR*, 633–641.