

Local Dense Logit Relations for Enhanced Knowledge Distillation

Liuchi Xu¹ Kang Liu² Jinshuai Liu¹ Lu Wang^{1*} Lisheng Xu¹ Jun Cheng^{3,4*}

¹ Northeastern University, China; ² South China Normal University, China;

³ Shenzhen Institutes of Advanced Technology, CAS, China;

⁴The Chinese University of Hong Kong, Hong Kong, China.

{xuliuchi@stumail,liujinshuai@stumail,wanglu@mail}.neu.edu.cn,

2022010207@m.scnu.edu.cn,xuls@bmie.neu.edu.cn,Jun.cheng@siat.ac.cn

Abstract

State-of-the-art logit distillation methods exhibit versatility, simplicity, and efficiency. Despite the advances, existing studies have yet to delve thoroughly into fine-grained relationships within logit knowledge. In this paper, we propose Local Dense Relational Logit Distillation (LDRLD), a novel method that captures inter-class relationships through recursively decoupling and recombining logit information, thereby providing more detailed and clearer insights for student learning. To further optimize the performance, we introduce an Adaptive Decay Weight (ADW) strategy, which can dynamically adjust the weights for critical category pairs using Inverse Rank Weighting (IRW) and Exponential Rank Decay (ERD). Specifically, IRW assigns weights inversely proportional to the rank differences between pairs, while ERD adaptively controls weight decay based on total ranking scores of category pairs. Furthermore, after the recursive decoupling, we distill the remaining non-target knowledge to ensure knowledge completeness and enhance performance. Ultimately, our method improves the student’s performance by transferring fine-grained knowledge and emphasizing the most critical relationships. Extensive experiments on datasets such as CIFAR-100, ImageNet-1K, and Tiny-ImageNet demonstrate that our method compares favorably with state-of-the-art logit-based distillation approaches. The code will be made publicly available.

1. Introduction

Deep Convolutional Neural Networks (CNNs) have been widely applied in various computer vision tasks, including image classification [15, 72], object detection [86, 89], and semantic segmentation [33, 93]. To enhance performance, researchers have designed increasingly complex networks [67] that demand greater computational and mem-

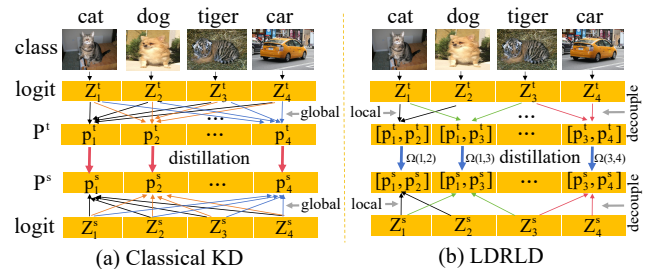


Figure 1. (a) Hinton et al. [18] introduce KD through global softmax, which calculates the probability, leading to information redundancy between classes and diminishing the logit discrimination. In classical KD, the prediction probability difference between “cat” and “dog” derived from teacher’s logit output is calculated as: $\Delta P_{KD} = |p_1^t - p_2^t| = \left| \frac{\exp(Z_1^t) - \exp(Z_2^t)}{\sum_{i=1}^C \exp(Z_i^t)} \right|$. (b) In contrast, our proposed LDRLD uses category pairs and calculates probability difference between “cat” and “dog” as: $\Delta P_{LDRLD} = |p_1^t - p_2^t| = \left| \frac{\exp(Z_1^t) - \exp(Z_2^t)}{\sum_{i=1}^C \exp(Z_i^t)} \right|$. It is obvious that $\Delta P_{LDRLD} > \Delta P_{KD}$, indicating that our approach enhances inter-class differences compared to KD and improves fine-grained logit discrimination.

ory resources. However, deploying these high-performance, parameter-heavy networks on resource-constrained devices, such as smartphones and cameras [10], remains challenging. In response to this issue, researchers have proposed several techniques to decrease the model size and computational cost, such as efficient network design [20, 52, 82], weight pruning [49], low-rank factorization [12], quantization [7], and knowledge distillation (KD) [18]. Among these techniques, KD is particularly efficient and practical. It facilitates the development of lightweight networks suitable for real-time mobile applications without altering the original network architecture.

The core idea of KD is often to transfer the “dark knowledge” from a larger pre-trained teacher model (teacher) to a more compact student model (student) at the logit (or output) layer, providing higher-level semantic information. Generally, KD uses the Kullback-Leibler (\mathcal{KL}) diver-

*Corresponding authors

gence [29] with a temperature coefficient to minimize the difference between the teacher’s and student’s output probability distributions. KD methods are commonly categorized into two main types: feature-based KD [27, 40, 51, 54, 68, 92] and logit-based KD [24, 31, 66, 69, 76, 84, 87]. Feature-based KD methods aim to align knowledge derived from intermediate layers, which contain rich feature representations. Despite their effectiveness, feature-based KD methods demand considerable computational and memory resources, thus posing significant training challenges. Consequently, researchers have shifted their focus to logit-based KD methods to reduce resource consumption. These methods often decouple logit knowledge [31, 66, 69, 70, 76, 84] or employ dynamic temperature techniques [24, 32, 60, 87] to extract critical information more effectively, thus boosting performance, computational efficiency, and versatility.

Considering these advancements, it is important to note that while current logit-based KD methods have achieved impressive performance, they may fail to adequately capture fine-grained inter-class relationships and enhance inter-class distinction, resulting in the limitations specified as follows. Since global softmax focuses on high-probability classes [18], it may reduce the differences between lower-probability classes, thereby limiting the student’s ability to capture their fine-grained logit relationships. Moreover, intuitively, categories with greater semantic differences are generally easier to distinguish. However, classical KD may introduce redundant information from other classes when modeling the two-class relationships. As shown in Fig. 1 (a), the distinction between “cat” and “dog” is reflected through probability distributions, yet the introduction of irrelevant (e.g., “car”) can interfere with and weaken their inter-class discriminability. These limitations collectively hinder the student’s ability to capture unique category characteristics, potentially degrading its performance. Therefore, it is essential to enhance inter-class discriminability while comprehensively leveraging logit relationships.

In this paper, we propose a Local Dense Relational Logit Distillation (LDRLD) method, which recursively decouples and recombines logit knowledge to provide richer information for student learning. By virtue of its effective capture of fine-grained logit relationships, LDRLD can enhance the efficacy of the distillation process. Based on this capability, we observe that categories with closer semantic relationships are more challenging to distinguish, whereas more distant categories are easier to recognize due to their semantic dissimilarities. To further refine our approach and enhance its performance, we introduce an Adaptive Decay Weight (ADW) strategy that dynamically adjusts weights based on ranked category pairs in LDRLD. By employing Inverse Rank Weighting (IRW) and Exponential Rank Decay (ERD), ADW assigns larger weights to closely related category pairs and smaller weights to distant ones. Specifi-

cally, IRW assigns larger weights to pairs with smaller rank differences and smaller weights to those with larger differences, whereas ERD dynamically adjusts weight decay based on the sum of the rankings. This dynamic weighting scheme enhances the student’s ability to optimize performance on visually complex and critical categories. Furthermore, after the recursive decoupling, we also distill the remaining non-target knowledge [84] to ensure completeness. By integrating the ADW strategy into LDRLD, we optimize the utilization of the logit knowledge, which leads to better model decision boundaries. Because of this optimization, the student can capture inter-class relationships and obtain rich contextual knowledge. By enhancing the student’s ability to identify critical categories, the proposed method effectively mitigates classification ambiguities and ensures accurate knowledge transfer.

To validate the effectiveness of our method, we conduct comprehensive experiments across various datasets. The contributions of this paper are summarized as follows:

- We propose Local Dense Relational Logit Distillation (LDRLD), a novel method that captures fine-grained logit relationships more effectively and enhances inter-class discriminability.
- We introduce the Adaptive Decay Weight (ADW) strategy, which comprises Inverse Rank Weighting (IRW) and Exponential Rank Decay (ERD). ADW can dynamically adjust the weights of ranked category pairs, thus enabling the student to more effectively optimize the classification of challenging categories.
- Extensive experimental results on diverse datasets, including CIFAR-100, ImageNet-1K, and Tiny-ImageNet, consistently show that our method outperforms existing state-of-the-art logit-based KD methods, and justify its ability to capture and transfer critical inter-class relationships.

2. Related Work

In this section, we will review existing works that relate to feature-based and logit-based knowledge distillation (KD).

2.1. Feature-based Knowledge Distillation

Extensive works [3, 11, 22, 26, 27, 37, 51, 62, 65, 72–75, 79, 86] have shown that the intermediate features contain rich and diverse information, which helps students in acquiring and mastering knowledge more effectively. Specifically, FitNets [51] uses mean squared error (MSE) between the teacher’s and student’s intermediate features to guide the student training. AT [27] adopts activation maps as attention mechanisms to replicate the teacher’s intermediate features more accurately. VID [1] boosts information transfer between the teacher and the student by increasing the mutual information (MI) of intermediate features. CRD [62]

proposes using contrastive learning for knowledge distillation, where contrastive targets are employed to enhance the transfer of the teacher’s representations to the student. ReviewKD [3] treats shallow knowledge as old knowledge and guides the student’s learning through repeated review. FCFD [37] optimizes the functional similarity between teacher and student features, thereby improving the student’s mimicry efficiency. CAT-KD [11] facilitates the student’s learning via the transfer of the teacher’s class activation maps. Although feature-based KD methods provide rich information, their high computational and memory demands have shifted research towards using logit knowledge as a more efficient distillation objective.

2.2. Logit-based Knowledge Distillation

Understanding the Role of Knowledge Distillation. Early research explores the underlying mechanisms of KD, for the purpose of seeking deeper insights [4, 8, 9, 16, 41, 42, 48, 55, 59, 61, 71, 77, 81, 83, 85, 90]. For instance, Tang et al. [61] theoretically demonstrate the potential of the label smoothing (LS) technique for improving student performance. Muller et al. [42] discover that LS can reduce intra-class variations and result in more compact feature representations. Furthermore, Shen et al. [55] show that tight clustering enhances the separability of similar semantic representations across categories. Despite these advancements, the performance of vanilla KD methods, such as deep mutual learning [83], teacher assistants [41, 57], BAN [9], and early stopping [4], may still fall short of practical application demands.

Refining Logit Knowledge for Improving Distillation Performance. Recent studies have achieved better performance by focusing on dynamic temperature and decoupling the logit knowledge to streamline distillation.

1) *Dynamic temperature.* Vanilla KD methods use a fixed temperature, which can limit the effectiveness of transferring logit knowledge. To overcome this limitation, dynamic temperature techniques have been proposed to improve the efficiency of logit knowledge transfer [24, 32, 60, 87]. For example, CTKD [32] adopts an adaptable temperature to modulate task difficulty. MLKD [24] employs Gram matrix distillation with varying temperatures to reduce prediction biases. LSKD [60] proposes adjusting the temperature based on the weighted standard deviation of logits and seamlessly integrates Zero-Score normalization. WTTM [87] uses transformed teacher matching to highlight temperature scaling in knowledge refinement while omitting it on the student side.

2) *Decoupling the logit knowledge.* Recent studies have focused on decoupling logit knowledge to develop more effective KD methods [31, 47, 66, 69, 76, 84]. For example, ATS [31] enhances performance by decomposing the \mathcal{KL} divergence into correction guidance, label smoothing, and

class distinguishability. DKD [84] proposes decoupling the \mathcal{KL} divergence into target and non-target class knowledge distillation, better exploiting information to improve performance. NKD [76] presents normalizing non-target logits to enhance the utilization of soft labels in distillation. ReKD [69] divides logit knowledge into head and tail categories to facilitate knowledge transfer. SDD [66] decouples global and local logit outputs into consistent and complementary terms to enhance the student’s ability to acquire precise logit knowledge. These methods improve student’s performance by absorbing knowledge more efficiently.

3. Method

In this section, we first review the vanilla knowledge distillation (KD) and then introduce our proposed Local Dense Relational Logit Distillation (LDRLD) method. After that, we present the optimization objective of the LDRLD.

Notations. For convenience, we consider the n -th training sample x_n with label y_n . Let C denote the number of categories. The ordered logit value vectors for the teacher and the student are defined as $\mathbf{Z}^t = [\mathbf{Z}_1^t, \mathbf{Z}_2^t, \dots, \mathbf{Z}_C^t] \in \mathbb{R}^{1 \times C}$ and $\mathbf{Z}^s = [\mathbf{Z}_1^s, \mathbf{Z}_2^s, \dots, \mathbf{Z}_C^s] \in \mathbb{R}^{1 \times C}$, respectively. Here, \mathbf{Z}_j^t and \mathbf{Z}_j^s , where $j \neq k$, represent the logits for the non-target classes, and k is the index of y_n , as discussed in [84].

3.1. Vanilla Knowledge Distillation

In KD [18], knowledge is transferred by minimizing the Kullback-Leibler (\mathcal{KL}) divergence \mathcal{L}_{KD} between the student’s and teacher’s output probability distributions. The output probability distributions for the teacher and the student are represented as $P^t = [p_1^t, p_2^t, \dots, p_C^t] \in \mathbb{R}^{1 \times C}$ and $P^s = [p_1^s, p_2^s, \dots, p_C^s] \in \mathbb{R}^{1 \times C}$, respectively. Here, p_j^t and p_j^s represent the predicted probabilities for the j -th class of the teacher and the student, respectively, computed as:

$$p_j^t = \exp(\mathbf{Z}_j^t) / S^t, \quad p_j^s = \exp(\mathbf{Z}_j^s) / S^s, \quad (1)$$

where S^t and S^s are normalization factors computed as $S^t = \sum_{i=1}^C \exp(\mathbf{Z}_i^t)$ and $S^s = \sum_{i=1}^C \exp(\mathbf{Z}_i^s)$, respectively, with temperature omitted for simplicity. Additionally, cross-entropy is taken as incorporate the task loss \mathcal{L}_{Task} to ensure accurate predictions. The total loss is thus formulated as:

$$\begin{aligned} \mathcal{L}_{Total} &= \mathcal{L}_{Task} + \gamma \mathcal{L}_{KD} \\ &= - \sum_{j=1}^C y_j \log p_j^s + \gamma \sum_{j=1}^C p_j^t \log(p_j^t / p_j^s), \end{aligned} \quad (2)$$

where γ serves as the weight to balance the two losses.

3.2. Local Dense Relational Logit Distillation

To provide richer information for student learning, we propose the LDRLD method shown in Fig. 2(a). Our method

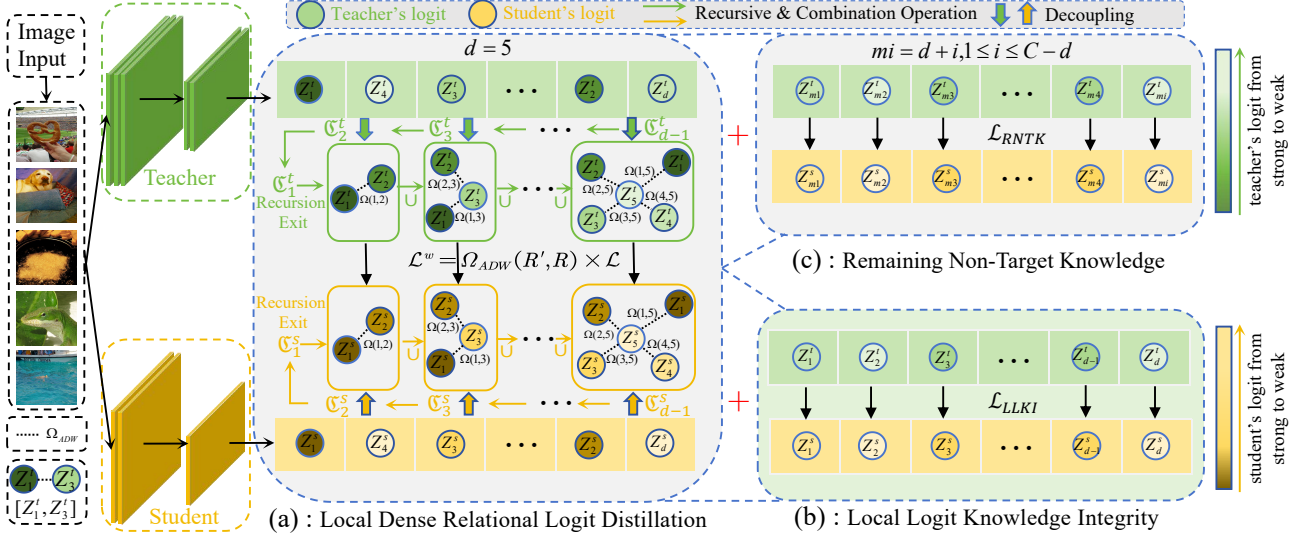


Figure 2. Overview of the proposed LDRLD framework, which includes the three key loss functions: \mathcal{L}^w , \mathcal{L}_{RNTK} , and \mathcal{L}_{LLKI}

recursively decouples and recombines logit knowledge, capturing fine-grained logit relationships with greater inter-class distinction. Specifically, we define the recursion depth d by selecting the top- d logits according to student’s logit arrange for a given sample, with $d \leq C$. The selected logits for the teacher and the student are denoted as $\mathcal{H}^t = [\mathbf{Z}_1^t, \mathbf{Z}_2^t, \dots, \mathbf{Z}_d^t] \subseteq \mathbf{Z}^t$ and $\mathcal{H}^s = [\mathbf{Z}_1^s, \mathbf{Z}_2^s, \dots, \mathbf{Z}_d^s] \subseteq \mathbf{Z}^s$, respectively, with each sequence arranged in descending order. Additionally, the argmax is used to identify the index of the maximum logit from the student’s output, \mathbf{Z}^s . To ensure consistent alignment, a mask is applied to both the student’s and teacher’s outputs. The mask vector $\mathbf{M} \in \mathbb{R}^C$ in the d -th iteration is defined as:

$$\mathbf{M}_{\pi(d)} = \begin{cases} -\infty & \text{if } \pi(d) = i_{\max}, \\ 1 & \text{otherwise,} \end{cases} \quad (3)$$

where $i_{\max} = \arg \max_i \mathbf{Z}_i^s$ ($i = 1, \dots, d$) and $\pi(d)$ denotes the index of the d -th largest student’s logit output. The proposed LDRLD method works as follows:

Step 1: Extraction of the Maximum Logit and the Masking Operation. The maximum logit is extracted from the vectors \mathbf{Z}^t and \mathbf{Z}^s as follows:

$$\mathbf{Z}_1^t = \mathbf{Z}_{\pi(1)}^t, \quad \mathbf{Z}_1^s = \mathbf{Z}_{\pi(1)}^s. \quad (4)$$

Based on Eq. (3), the mask $\mathbf{M}_{\pi(1)}$ is then applied to update the teacher’s and student’s logit vectors as follows:

$$\mathbf{Z}^t \leftarrow \mathbf{M}_{\pi(1)} \odot \mathbf{Z}^t, \quad \mathbf{Z}^s \leftarrow \mathbf{M}_{\pi(1)} \odot \mathbf{Z}^s, \quad (5)$$

where \odot denotes the Hadamard product. The resulting logits are used in **Step 2**.

Step 2: Recursive Decoupling and Combination Operation. We define the concatenation function $\phi: \mathbb{R} \uplus \mathbb{R} \rightarrow \mathbb{R}^{1 \times 2}$, where \uplus denotes an operation that combines two logits into a pair. Similar to **Step 1**, we use Eqs. (3) and (4) to

extract the second-largest student’s logit when $d = 2$, denoted as \mathbf{Z}_2^t and \mathbf{Z}_2^s , respectively. The mask $\mathbf{M}_{\pi(2)}$ based on Eq. (3) is then applied to update these logit vectors:

$$\mathbf{Z}^t \leftarrow \mathbf{M}_{\pi(2)} \odot \mathbf{Z}^t, \quad \mathbf{Z}^s \leftarrow \mathbf{M}_{\pi(2)} \odot \mathbf{Z}^s. \quad (6)$$

After that, the combination generation proceeds recursively as described below:

Base Case. For depth $d = 2$, the first set of combinations is generated for the teacher and the student:

$$\mathfrak{C}_1^t = \phi(\mathbf{Z}_1^t, \mathbf{Z}_2^t), \quad \mathfrak{C}_1^s = \phi(\mathbf{Z}_1^s, \mathbf{Z}_2^s), \quad (7)$$

where the symbol \mathfrak{C} represents the result of the combination, and $\mathfrak{C}_0 = \emptyset$. \mathfrak{C}_1^t and \mathfrak{C}_1^s denote the recursive exit.

Recursive Step. We repeat **Steps 1 and 2** as follows:

- (1) Extraction: Select the d -th largest logit:

$$\mathbf{Z}_d^t = \mathbf{Z}_{\pi(d)}^t, \quad \mathbf{Z}_d^s = \mathbf{Z}_{\pi(d)}^s, \quad (8)$$

where $\pi(d)$ is the index of the d -th largest student’s logit output, \mathbf{Z}^s .

- (2) Masking: Apply the mask $\mathbf{M}_{\pi(d)}$ based on Eq. (3) to update the teacher’s and student’s logit vectors:

$$\mathbf{Z}^t \leftarrow \mathbf{M}_{\pi(d)} \odot \mathbf{Z}^t, \quad \mathbf{Z}^s \leftarrow \mathbf{M}_{\pi(d)} \odot \mathbf{Z}^s. \quad (9)$$

- (3) Combination Generation: For any recursion depth $d \geq 2$, new sets of combinations are generated by applying ϕ to the newly extracted logit \mathbf{Z}_d^t and \mathbf{Z}_d^s , and all previously extracted logits ($\forall i, 2 \leq i \leq d-1$). The combination sets \mathfrak{C}_{d-1}^t and \mathfrak{C}_{d-1}^s are defined as follows:

$$\mathfrak{C}_{d-1}^t = \left\{ \mathfrak{C}_{d-2}^t \cup \bigcup_{i=1}^{d-1} \{ \phi(\mathbf{Z}_i^t, \mathbf{Z}_d^t) \} \right\} \in \mathbb{R}^{\frac{d(d-1)}{2} \times 2}, \quad (10)$$

$$\mathfrak{C}_{d-1}^s = \left\{ \mathfrak{C}_{d-2}^s \cup \bigcup_{i=1}^{d-1} \{ \phi(\mathbf{Z}_i^s, \mathbf{Z}_d^s) \} \right\} \in \mathbb{R}^{\frac{d(d-1)}{2} \times 2}, \quad (11)$$

where \cup represents the union operation.

Step 3: Normalization. To ensure that the combinations in \mathcal{C} satisfy the properties of a probability distribution, we normalize them using the softmax function σ . Let (p_i^t, p_j^t) denote the teacher’s logit pair, and let (p_i^s, p_j^s) denote the student’s logit pair. The probabilities $p_i^t, p_j^t, p_i^s,$ and p_j^s are computed as follows via σ :

$$\begin{cases} p_i^t = \exp(\mathbf{Z}_i^t/\tau) / S_{ij}^t, p_j^t = \exp(\mathbf{Z}_j^t/\tau) / S_{ij}^t, \\ p_i^s = \exp(\mathbf{Z}_i^s/\tau) / S_{ij}^s, p_j^s = \exp(\mathbf{Z}_j^s/\tau) / S_{ij}^s, \end{cases} \quad (12)$$

where S_{ij}^t and S_{ij}^s are the normalization factors for the teacher and student, respectively. Here, $S_{ij}^t = \exp(\mathbf{Z}_i^t/\tau) + \exp(\mathbf{Z}_j^t/\tau)$ and $S_{ij}^s = \exp(\mathbf{Z}_i^s/\tau) + \exp(\mathbf{Z}_j^s/\tau)$. τ represents the temperature coefficient that controls the softness of the output probability distribution.

Building upon **Steps 1 to 3**, we transfer the teacher’s *local dense relational logit knowledge* to the student using the loss defined as follows:

$$\begin{aligned} \mathcal{L} &= \sum_{i=1}^{d-1} \mathcal{KL}(\mathbf{c}_i^t, \mathbf{c}_i^s) \\ &= \sum_{i=1}^{d-1} \sum_{j=i+1}^d \mathcal{KL}(\sigma(\phi(\mathbf{Z}_i^t, \mathbf{Z}_j^t), \tau), \sigma(\phi(\mathbf{Z}_i^s, \mathbf{Z}_j^s), \tau)) \quad (13) \\ &= \sum_{i=1}^{d-1} \sum_{j=i+1}^d [p_i^t \log(p_i^t/p_i^s) + p_j^t \log(p_j^t/p_j^s)]. \end{aligned}$$

Local Logit Knowledge Integrity. Additionally, when the recursion depth d is reached, the d logits are generated sequentially to capture hierarchical relationships, thereby preserving local logit knowledge integrity (LLKI). The resulting teacher’s and student’s logits, \mathcal{H}^t and \mathcal{H}^s , are presented in Fig. 2(b). The loss \mathcal{L}_{LLKI} is defined as follows:

$$\mathcal{L}_{LLKI} = \mathcal{KL}(\mathcal{H}^t, \mathcal{H}^s) = \sum_{i=1}^d p_i^t \log(p_i^t/p_i^s), \quad (14)$$

where the probabilities of the teacher and the student are denoted as $p_i^t = \sigma(\mathbf{Z}_i^t/\tau)$, and $p_i^s = \sigma(\mathbf{Z}_i^s/\tau)$, respectively.

Adaptive Decay Weight strategy. So far, we have uniformly assigned weights to all category pairs. However, this may hinder the optimization of visually complex or critical categories and lead to network performance degradation. To address this, we propose an Adaptive Decay Weight (ADW) strategy based on human perception of category similarity. The ADW integrates Inverse Rank Weighting and Exponential Rank Decay. It dynamically adjusts weights to enhance discrimination by assigning greater weights to category pairs that are more difficult for the model to distinguish (e.g., $[\mathbf{Z}_1^t, \mathbf{Z}_2^t]$), and smaller weights to less related category pairs (e.g., $[\mathbf{Z}_1^t, \mathbf{Z}_4^t]$). This targeted weighting sharpens the student’s focus on key categories while reducing attention to others. Details of the two key components of the ADW strategy are described below:

(1) Inverse Rank Weighting. To emphasize the importance of closely ranked category pairs, we introduce Inverse Rank Weighting (IRW), which is defined as:

$$\Gamma_{IRW}(R', R) = \text{Inv}(|R - R'| + \epsilon), \quad (15)$$

where $\text{Inv}(x)$ denotes the inverse of x . R' and R are the rankings of categories i and j in the pair $[\mathbf{Z}_i, \mathbf{Z}_j]$, respectively. The term $|R - R'|$ represents the difference in category rankings, and $\epsilon = 1.50$ prevents division by zero.

IRW assigns larger weights to pairs with smaller ranking differences and smaller weights to those with larger differences. Thus, IRW directs the student’s focus toward optimizing more similar category pairs.

(2) Exponential Rank Decay. When calculating the pairwise weights according to Eq. (15), note that the pairs $[\mathbf{Z}_1^t, \mathbf{Z}_2^t]$ and $[\mathbf{Z}_{12}^t, \mathbf{Z}_{13}^t]$ receive the same weight. However, intuitively, distinguishing between $[\mathbf{Z}_1^t, \mathbf{Z}_2^t]$ is more difficult due to highly semantic similarity, so the weight for the former should be greater than that for the latter. Based on this understanding, we propose an Exponential Ranking Decay (ERD) that dynamically adjusts weight decay according to the sum of the rankings, defined as:

$$\Phi_{ERD}(R', R) = \delta \times \exp(-\lambda(R + R')), \quad (16)$$

where δ , the weight, is set to 2.0, and λ , the decay rate, is set to 0.05 in our experiments. The decay rate controls how quickly the weight decreases as the sum of rankings ($R' + R$) increase, and a larger value of λ results in faster weight decay.

With IRW and ERD, our ADW strategy is defined as follows:

$$\Omega_{ADW}(R', R) = \Gamma_{IRW}(R', R) \times \Phi_{ERD}(R', R). \quad (17)$$

By combining Eq. (13) and Eq. (17), we formulate the loss of \mathcal{L}^w to facilitate the transfer of local dense relational logit knowledge, defined as follows:

$$\mathcal{L}^w = \sum_{i=1}^{d-1} \sum_{j=i+1}^d \Omega_{ADW}(i, j) [p_i^t \log(p_i^t/p_i^s) + p_j^t \log(p_j^t/p_j^s)]. \quad (18)$$

The complete local logit knowledge \mathcal{L}_{Local} is defined as follows:

$$\mathcal{L}_{Local} = \mathcal{L}^w + \mathcal{L}_{LLKI}. \quad (19)$$

Remaining Non-Target Knowledge. Inspired by the effectiveness of non-target knowledge in previous studies [6, 31, 76, 84], we leverage the remaining non-target classes to ensure comprehensive knowledge transfer and improve student performance. We denote the remaining logits for the teacher and student as $\bar{\mathcal{H}}^t = [\mathbf{Z}_{d+1}^t, \mathbf{Z}_{d+2}^t, \dots, \mathbf{Z}_C^t]$, and $\bar{\mathcal{H}}^s = [\mathbf{Z}_{d+1}^s, \mathbf{Z}_{d+2}^s, \dots, \mathbf{Z}_C^s]$, respectively, in Fig. 2(c). The knowledge is distilled using \mathcal{L}_{RNTK} , defined as follows:

$$\mathcal{L}_{RNTK} = \mathcal{KL}(\bar{\mathcal{H}}^t, \bar{\mathcal{H}}^s) = \sum_{i=d+1}^C \bar{p}_i^t \log(\bar{p}_i^t/\bar{p}_i^s), \quad (20)$$

where the probabilities of the teacher and the student are $\bar{p}_i^t = \sigma(\mathbf{Z}_i^t/\tau)$ and $\bar{p}_i^s = \sigma(\mathbf{Z}_i^s/\tau)$, where $d + 1 \leq i \leq C$, respectively.

3.3. Optimization Objective

The proposed method integrates task loss, local relational information, and knowledge from non-target classes. This integration enables precise adjustment of the student’s parameters through the following optimization function:

$$\mathcal{L}_{LDRLD} = \mathcal{L}_{Task} + \alpha\mathcal{L}_{Local} + \beta\mathcal{L}_{RNTK} \quad (21)$$

where α and β are weighting factors that balance the contributions of these losses.

To facilitate the understanding and implementation of our method, we provide a detailed algorithm and specific descriptions in the *Supplementary Material*.

4. Experiments

4.1. Comparison with State-of-the-Art Methods

Image Classification Results on CIFAR-100. We evaluated our method across various network architectures. As shown in Table 1, we first employed an experimental setup where the teacher and the student share the same type of architecture. In contrast, in Table 2, the teacher and the student use different-architectures. Experimental results show that our proposed method outperforms existing state-of-the-art logit-based KD methods, such as SDD, WTTM, and TeKAP. Specifically, in the same-architecture experiment, our method improves performance by 1.08% to 3.87%. In the different-architecture experiment, the improvement ranges from 1.68% to 3.39%. These improvements are mainly due to LDRLD’s ability to extract fine-grained inter-class knowledge from the teacher. While current logit-based KD methods utilize global softmax to capture inter-class relationships, they may obscure subtle differences among categories, thus impairing targeted learning enhancements. Furthermore, the global softmax mechanism often neglects low-probability categories, limiting performance. These differences highlight LDRLD’s effectiveness, which leverages local logit relational techniques to overcome these limitations, enhance inter-class discriminability and ensure better recognition.

Image Classification Results on ImageNet-1K. The results on ImageNet-1K are shown in Table 3. Using the same type of architecture, our method outperforms the vanilla KD by 1.22% in top-1 accuracy. Additionally, when using ResNet50 as the teacher and MobileNetV1 as the student, experimental results show that our method achieves improvements of 2.63% in top-1 and 1.51% in top-5 accuracy. These results further confirm the effectiveness and potential of LDRLD. However, when comparing LDRLD with WTTM using the same architecture, LDRLD’s performance

is slightly lower. This is probably because WTTM trains the student to approximate smoother soft targets rather than peaked soft targets, which helps it absorb more knowledge. In contrast, LDRLD faces challenges caused by the complex category structure in the ImageNet-1K and the limited capacity of the student, which make it less effective in capturing relationships between categories.

Image Classification Results on CIFAR-100 using ViT-based Architectures. Classical KD methods are typically validated within CNN-based frameworks, but their performance in ViT-based frameworks has been shown to be limited, indicating that their generalization ability remains uncertain. To justify the generalization ability of LDRLD, we conducted ViT-based experiments and employed different architectures for distillation, using students based on CNN, ViT, and MLP, and **state-of-the-art** results were obtained, and demonstrating the effectiveness of LDRLD, as shown in Table 4. Evaluating LDRLD on ViT-based models provides additional insights into the method’s versatility and robustness across different model paradigms.

4.2. Ablation Studies

To ensure a fair comparison and validate the effectiveness of component, we conducted ablation experiments on various teacher-student pairs. For more ablation experiments, see the *supplementary materials*.

Impact of the recursion depth d . We studied the impact of the recursion depth d on student performance. When the depth d is small, the limited number of category pairs hinders the student from capturing sufficient inter-class distinctions. Conversely, a large depth d may introduce significant noise, interfere with the learning process, and lead to information redundancy due to the large distances between categories. Thus, excessive recursion depth increases optimization complexity without improving performance. Our findings indicate that a moderate recursion depth, particularly the depth $d = 7$ in Fig. 3, achieves the optimal performance across homogeneous and heterogeneous architectures.




Figure 3. Impact of the depth d on the performance of the student on CIFAR-100.

Impact of the ADW. We conducted an ablation study to evaluate the impact of the component Ω_{ADW} on student performance. As shown in Table 5, using \mathcal{L} and \mathcal{L}_{LLKI} , the performance is quite favorable compared to training the

Distillation Manner	Teacher Accuracy	Publication	ResNet56	ResNet110	ResNet110	ResNet32×4	WRN-40-2	WRN-40-2	VGG13
	Student Accuracy	Year	ResNet20	ResNet20	ResNet32	ResNet8×4	WRN-16-2	WRN-40-1	VGG8
Features	FitNet [51]	ICLR'15	69.21	68.99	71.06	73.50	73.58	72.24	71.02
	RKD [44]	CVPR'19	69.61	69.25	71.82	71.90	73.35	72.22	71.48
	CRD [62]	ICLR'20	71.16	71.46	73.48	75.51	75.48	74.14	73.94
	ReviewKD [3]	CVPR'21	71.89	71.85	73.89	75.63	76.12	75.09	74.84
	CAT-KD [11]	CVPR'23	71.62	71.14	73.62	76.91	75.60	74.82	74.65
	FCFD [37]	ICLR'23	71.68	-	-	76.80	76.34	75.43	74.86
	NORM [38]	ICLR'23	71.35	71.55	73.67	76.49	75.65	74.82	73.95
Logits	KD [18]	NeurIPS'14	70.66	70.67	73.08	73.33	74.92	73.54	72.98
	DKD [84]	CVPR'22	71.97	70.91	74.11	76.32	76.24	74.81	74.68
	CTKD [32]	AAAI'23	71.19	70.99	73.52	73.70	75.45	73.93	73.52
	NKD [76]	ICCV'23	71.18	71.26	73.50	76.35	75.43	74.25	74.86
	LCKA [91]	IJCAI'24	-	70.87	73.64	75.12	75.78	74.63	74.35
	LSKD [60]	CVPR'24	71.24	71.61	73.76	76.16	76.22	74.43	74.23
	SDD [66]	CVPR'24	71.52	71.58	73.97	75.09	75.86	74.53	73.91
	WTTM [87]	ICLR'24	71.92	71.67	74.13	76.06	76.37	74.58	74.44
	TeKAP [19]	ICLR'25	71.32	71.24*	73.42*	74.79	75.21	73.80	74.00
	LDRLD	Ours	72.20	71.98	74.16	77.20	76.35	74.98	75.06
	Δ	Ours	+1.54	+1.31	+1.08	+3.87	+1.43	+1.44	+2.08

Table 1. Evaluation of the top-1 accuracy (%) of student with the same architecture on the CIFAR-100 validation set. Δ denotes the improvement over the vanilla KD method. **We report the average results of three runs for all tables.** * represents our reproduced results.

Distillation Manner	Teacher Accuracy	Publication	ResNet32×4	WRN40-2	ResNet50	VGG13	ResNet32×4	ResNet50
	Student Accuracy	Year	ShuffleNetV1	ShuffleNetV1	MobileNetV2	MobileNetV2	ShuffleNetV2	VGG8
Features	FitNet [51]	ICLR'15	73.54	73.73	63.16	63.16	73.54	70.69
	RKD [44]	CVPR'19	72.28	72.21	64.43	64.43	73.21	71.50
	CRD [62]	ICLR'20	75.11	76.05	69.11	69.11	75.65	74.30
	ReviewKD [3]	CVPR'21	77.45	77.14	69.89	70.37	77.78	75.34
	NORM [38]	ICLR'23	77.42	77.06	70.56	68.94	78.07	75.17
	FCFD [37]	ICLR'23	78.12	77.81	71.07	70.67	78.20	-
	CAT-KD [11]	CVPR'23	78.26	77.35	71.36	69.13	78.41	75.39
Logits	KD [18]	NeurIPS'14	74.07	74.83	67.35	67.37	74.45	73.81
	DKD [84]	CVPR'22	76.45	76.70	70.35	69.71	77.07	75.34
	CTKD [32]	AAAI'23	74.48	75.78	68.47	68.46	75.31	73.63
	NKD [76]	ICCV'23	76.31	76.46	70.22	70.67	76.92	74.01
	LSKD [60]	CVPR'24	-	-	69.02	68.61	75.56	74.42
	SDD [66]	CVPR'24	76.30	76.65	69.55	68.79	76.67	74.89
	WTTM [87]	ICLR'24	74.37	75.42	69.59	69.16	76.55	74.82
	TeKAP [19]	ICLR'25	74.92	76.75	69.00*	67.39*	75.43	74.37*
	LDRLD	Ours	76.46	77.09	70.74	70.11	77.33	75.49
	Δ	Ours	+2.39	+2.26	+3.39	+2.74	+2.88	+1.68

Table 2. Evaluation of the top-1 accuracy (%) of student with the different architecture on the CIFAR-100 validation set.

ResNet34 (teacher): 73.31% Top-1, 91.42% Top-5 accuracy. ResNet18 (student): 69.75% Top-1, 89.07% Top-5 accuracy.															
Features	AT[27]	CRD[62]	ReviewKD[3]	FCFD[37]	CAT-KD[11]	Logits	KD[18]	CTKD[32]	DKD[84]	LSKD[60]	SDD[66]	WTTM[87]	TeKAP[19]	LDRLD	Δ
Top-1	70.69	71.17	71.61	72.24	71.26	Top-1	70.66	71.32	71.70	71.42	71.44	72.19	71.35*	71.88	+1.22
Top-5	90.01	90.13	90.51	90.74	90.45	Top-5	89.88	90.27	90.31	90.29	90.05	-	90.54*	90.58	+0.70
ResNet50 (teacher): 76.16% Top-1, 92.87% Top-5 accuracy. MobileNetV1 (student): 68.87% Top-1, 88.76% Top-5 accuracy.															
Features	AT[27]	CRD[62]	ReviewKD[3]	FCFD[37]	CAT-KD[11]	Logits	KD[18]	IPWD[43]	DKD[84]	LSKD[60]	SDD[66]	WTTM[87]	TeKAP[19]	LDRLD	Δ
Top-1	70.18	71.32	72.56	73.37	72.24	Top-1	70.49	72.65	72.05	72.18	72.24	73.09	72.87*	73.12	+2.63
Top-5	89.68	90.41	91.00	91.35	91.13	Top-5	89.92	91.08	91.05	90.80	90.71	-	91.05*	91.43	+1.51

Table 3. Evaluate the top-1 and top-5 accuracy (%) of student using same and different architectures on the ImageNet-1K validation set.

Teacher	Student	From Scratch		Feature-based				Logit-based						
		T:Accuracy	S:Accuracy	FitNet [51]	RKD [44]	CRD [62]	FOFA [34]	KD [18]	DKD [84]	DIST [21]	OFA [14]	TeKAP [19]	LDRLD	Δ
Swin-T	ResNet18	89.26	74.01	78.87	74.11	77.63	81.22	78.74	80.26	77.75	80.54	81.38*	82.17	+3.43
ViT-S	ResNet18	92.04	74.01	77.71	73.72	76.60	80.11	77.26	78.10	76.49	80.15	79.06*	80.36	+3.10
Mixer-B/16	ResNet18	87.29	74.01	77.15	73.75	76.42	80.07	77.79	78.67	76.36	79.39	80.05*	80.69	+2.90
Swin-T	MobileNetV2	89.26	73.68	74.28	69.00	79.80	78.78	74.68	71.07	72.89	80.98	80.23*	81.64	+6.96
ViT-S	MobileNetV2	92.04	73.68	73.54	68.46	78.14	78.87	72.77	69.80	72.54	78.45	78.41*	79.21	+6.44
Mixer-B/16	MobileNetV2	87.29	73.68	73.78	68.95	78.15	78.62	73.33	70.20	73.26	78.78	79.89*	80.64	+7.31
ConvNeXt-T	DeiT-T	88.41	68.00	60.78	69.79	65.94	79.59	72.99	74.60	73.55	75.76	76.32*	77.46	+4.47
Mixer-B/16	DeiT-T	87.29	68.00	71.05	69.89	65.35	74.66	71.36	73.44	71.67	73.90	74.83*	75.31	+3.95
ConvNeXt-T	Swin-P	88.41	72.63	24.06	71.73	67.09	80.74	76.44	76.80	76.41	78.32	79.18*	80.71	+4.27
Mixer-B/16	Swin-P	87.29	72.63	75.20	70.82	67.03	78.44	75.93	76.39	75.85	78.93	78.97*	80.52	+4.59
ConvNeXt-T	ResMLP-S12	88.41	66.56	45.47	65.82	63.35	83.50	72.25	73.22	71.93	75.21	81.14*	79.28	+7.03
Swin-T	ResMLP-S12	89.26	66.56	63.12	64.66	61.72	80.94	71.89	72.82	11.05	73.58	80.22*	80.54	+8.65

Table 4. Evaluation of the top-1 accuracy (%) of student using ViT-based heterogeneous architectures on the CIFAR100 dataset.

student from scratch. After incorporating Ω_{ADW} , the performance improves further. The results show that the component Ω_{ADW} , which adjusts category pair weights, can enhance discriminability and effectively improve the performance of the student.

\mathcal{L}_{Local}			VGG13	VGG13	ResNet50	ResNet32×4
\mathcal{L}	Ω_{ADW}	\mathcal{L}_{LLKI}	VGG8	MobileNetV2	MobileNetV2	ShuffleNetV2
-	-	-	70.50	64.60	64.60	71.82
✓	-	✓	74.25 (+3.75)	68.95 (+4.35)	69.23 (+4.63)	76.92 (+5.10)
✓	✓	✓	74.46 (+3.96)	69.27 (+4.67)	69.52 (+4.92)	77.07 (+5.25)

Table 5. Impact of ADW on student’s performance on CIFAR100.

4.3. Generalization Performance Evaluation

To validate the generalization ability of our method, we extended the task to person Re-ID. As presented in Table 6, the results indicate that our method exhibits superior performance compared to other distillation methods, such as AT and DKD, confirming its effectiveness.

Methods	Rank@1	Rank@5	Rank@10	mAP
ResNet50(teacher)	88.03	95.01	96.67	71.77
ResNet18(student)	85.04	94.06	96.23	65.30
AT [27]	86.71	94.88	96.89	68.42
KD [18]	88.21	94.88	96.80	71.79
DKD [84]	87.50	94.74	96.71	71.83
LDRLD	88.31	95.05	97.03	72.13
Δ	+3.27	+0.99	+0.80	+6.83

Table 6. Comparison with different methods on Market-1501.

4.4. Visualizations

We evaluate our experimental results using Grad-CAM [53] and Correlation Matrices for visualization and the latter is provided in the *Supplementary Material*.

Visualization of Grad-CAM. To demonstrate the effectiveness of our proposed method, we use Grad-CAM to visualize the feature map. As shown in Figs. 4 (c) and (e), our method accurately identifies target objects by focusing on the detailed features of primary categories. In contrast, Figs. 4 (b) and (d) reveal that vanilla KD tends to focus on non-critical regions, causing it to overlook crucial information, and leading to performance degradation. Instead,

Our method’s targeted attention can retain important details, thereby enhancing overall performance.

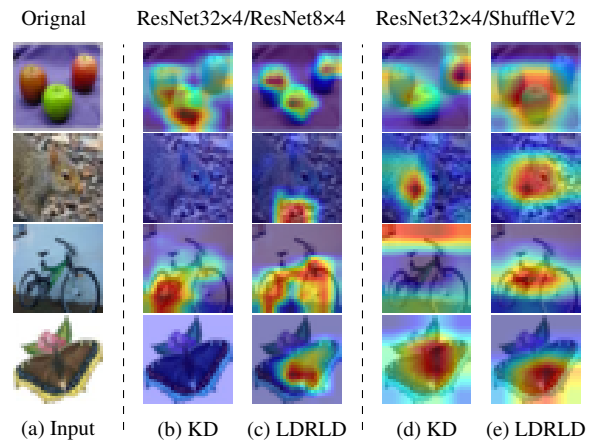


Figure 4. Feature map visualization of the student’s penultimate layers on CIFAR-100 dataset using vanilla KD and LDRLD.

5. Conclusion

We propose the Local Dense Relational Logit Distillation method and further incorporate the Adaptive Decay Weight strategy, including Inverse Rank Weighting and Exponential Rank Decay, to effectively capture fine-grained features within logit knowledge. Experimental results demonstrate that our method outperforms state-of-the-art logit-based KD methods on the CIFAR-100, Tiny-ImageNet, MSCOCO-2017, and ImageNet-1K datasets, validating its effectiveness.

Limitations and future work. Despite these achievements, our method still requires manual selection of the recursion depth d . Since the optimal choice of d depends on the number of classes and the complexity of the task, it is necessary to adjust d accordingly. Therefore, future work will explore adaptive approaches to address this limitation.

Acknowledgement

The work was supported in part by the National Natural Science Foundation of China (U21A20487, 62273082), in part by the Guangdong Technology Project (2023TX07Z126, 2022B1515120067), in part by the Shenzhen Technology Project (JCYJ20220818101206014, JCYJ20220818101211025, GJHZ20240218112504008), in part by Yunnan Science & Technology Project (202305AF150152, 202302AD080008).

References

- [1] Sungsoo Ahn, Shell Xu Hu, Andreas Damianou, Neil D Lawrence, and Zhenwen Dai. Variational information distillation for knowledge transfer. In *CVPR*, pages 9163–9171, 2019. 2, 1, 5
- [2] Zhiqiang Bao, Zihao Chen, Chang-Dong Wang, Wei-Shi Zheng, Zhenhua Huang, and Yunwen Chen. Post-distillation via neural resuscitation. *IEEE Transactions on Multimedia*, 26:3046–3060, 2024. 2
- [3] Pengguang Chen, Shu Liu, Hengshuang Zhao, and Jiaya Jia. Distilling knowledge via knowledge review. In *CVPR*, pages 5008–5017, 2021. 2, 3, 7, 1, 4
- [4] Jang Hyun Cho and Bharath Hariharan. On the efficacy of knowledge distillation. In *ICCV*, pages 4794–4802, 2019. 3
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 1
- [6] Qianggang Ding, Sifan Wu, Tao Dai, Hao Sun, Jiadong Guo, Zhang-Hua Fu, and Shutao Xia. Knowledge refinery: Learning from decoupled label. In *AAAI*, pages 7228–7235, 2021. 5
- [7] Zhihao Duan, Ming Lu, Jack Ma, Yuning Huang, Zhan Ma, and Fengqing Zhu. Qarv: Quantization-aware resnet vae for lossy image compression. *TPAMI*, 46(1):436–450, 2023. 1
- [8] Luyang Fang, Yongkai Chen, Wenxuan Zhong, and Ping Ma. Bayesian knowledge distillation: A Bayesian perspective of distillation with uncertainty quantification. In *International Conference on Machine Learning*, pages 12935–12956. PMLR, 2024. 3
- [9] Tommaso Furlanello, Zachary Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. Born again neural networks. In *ICML*, pages 1607–1616, 2018. 3
- [10] Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. Knowledge distillation: A survey. *IJCV*, 129: 1789–1819, 2021. 1
- [11] Ziyao Guo, Haonan Yan, Hui Li, and Xiaodong Lin. Class attention transfer based knowledge distillation. In *CVPR*, pages 11868–11877, 2023. 2, 3, 7, 1
- [12] Benjamin D Haeffele and Rene Vidal. Structured low-rank matrix factorization: Global optimality, algorithms, and applications. *TPAMI*, 42(6):1468–1482, 2019. 1
- [13] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *NeurIPS*, 31, 2018. 4
- [14] Zhiwei Hao, Jianyuan Guo, Kai Han, Yehui Tang, Han Hu, Yunhe Wang, and Chang Xu. One-for-all: Bridge the gap between heterogeneous architectures in knowledge distillation. *NeurIPS*, 36:79570–79582, 2023. 8
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 1
- [16] Zhou Helong, Song Liangchen, Chen Jiajie, Zhou Ye, Wang Guoli, Yuan Junsong, and Qian Zhang. Rethinking soft labels for knowledge distillation: a bias-variance tradeoff perspective. In *ICLR*, 2021. 3
- [17] Byeongho Heo, Jeesoo Kim, Sangdoo Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *ICCV*, pages 1921–1930, 2019. 1
- [18] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. In *NeurIPS Workshop*, 2014. 1, 2, 3, 7, 8, 4, 5
- [19] Md Imtiaz Hossain, Sharmen Akhter, Choong Seon Hong, and Eui-Nam Huh. Single teacher, multiple perspectives: Teacher knowledge augmentation for enhanced knowledge distillation. In *ICLR*, 2025. 7, 8, 5
- [20] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. In *CVPR*, pages 484–492, 2017. 1
- [21] Tao Huang, Shan You, Fei Wang, Chen Qian, and Chang Xu. Knowledge distillation from a stronger teacher. *NeurIPS*, 35: 33716–33727, 2022. 8
- [22] Tao Huang, Yuan Zhang, Mingkai Zheng, Shan You, Fei Wang, Chen Qian, and Chang Xu. Knowledge diffusion for distillation. In *NeurIPS*, pages 65299–65316, 2023. 2, 1
- [23] Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv preprint arXiv:1707.01219*, 2017. 1, 5
- [24] Ying Jin, Jiaqi Wang, and Dahua Lin. Multi-level logit distillation. In *CVPR*, pages 24276–24285, 2023. 2, 3
- [25] Jangho Kim, SeongUk Park, and Nojun Kwak. Paraphrasing complex network: Network compression via factor transfer. *NeurIPS*, 31, 2018. 1, 5
- [26] Jungeun Kim, Junwon You, Dongjin Lee, Ha Young Kim, and Jae-Hun Jung. Do topological characteristics help in knowledge distillation? In *ICML*, 2024. 2
- [27] Nikos Komodakis and Sergey Zagoruyko. Paying more attention to attention: improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2017. 2, 7, 8, 1, 5
- [28] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Toronto, ON, Canada*, 2009. 1
- [29] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951. 2
- [30] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015. 1
- [31] Xin-Chun Li, Wen-Shu Fan, Shaoming Song, Yinchuan Li, Shao Yunfeng, De-Chuan Zhan, et al. Asymmetric tem-

- perature scaling makes larger networks teach well again. *NeurIPS*, 35:3830–3842, 2022. 2, 3, 5
- [32] Zheng Li, Xiang Li, Lingfeng Yang, Borui Zhao, Renjie Song, Lei Luo, Jun Li, and Jian Yang. Curriculum temperature for knowledge distillation. In *AAAI*, pages 1504–1512, 2023. 2, 3, 7, 1, 4
- [33] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yinan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *CVPR*, pages 7061–7070, 2023. 1
- [34] Jhe-Hao Lin, Yi Yao, Chan-Feng Hsu, Hongxia Xie, Hong-Han Shuai, and Wen-Huang Cheng. Feature-based one-for-all: A universal framework for heterogeneous knowledge distillation. *AAAI*, 2025. 8
- [35] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755, 2014. 1
- [36] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. 4
- [37] Dongyang Liu, Meina Kan, Shiguang Shan, and Xilin CHEN. Function-consistent feature distillation. In *ICLR*, 2023. 2, 3, 7, 1
- [38] Xiaolong Liu, LUKING LI, Chao Li, and Anbang Yao. NORM: Knowledge distillation via n-to-one representation matching. In *ICLR*, 2023. 7, 1
- [39] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *ECCV*, pages 116–131, 2018. 1, 2
- [40] Roy Miles and Krystian Mikolajczyk. Understanding the role of the projector in knowledge distillation. In *AAAI*, pages 4233–4241, 2024. 2
- [41] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. Improved knowledge distillation via teacher assistant. In *AAAI*, pages 5191–5198, 2020. 3
- [42] Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? *NeurIPS*, 32, 2019. 3
- [43] Yulei Niu, Long Chen, Chang Zhou, and Hanwang Zhang. Respecting transfer gap in knowledge distillation. *NeurIPS*, 35:21933–21947, 2022. 7, 1
- [44] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *CVPR*, pages 3967–3976, 2019. 7, 8, 1, 5
- [45] Nikolaos Passalis and Anastasios Tefas. Learning deep representations with probabilistic knowledge transfer. In *ECCV*, pages 268–284, 2018. 1, 5
- [46] Baoyun Peng, Xiao Jin, Jiaheng Liu, Dongsheng Li, Yichao Wu, Yu Liu, Shunfeng Zhou, and Zhaoning Zhang. Correlation congruence for knowledge distillation. In *ICCV*, pages 5007–5016, 2019. 1, 5
- [47] Bo Peng, Zhen Fang, Guangquan Zhang, and Jie Lu. Knowledge distillation with auxiliary variable. In *ICML*, pages 40185–40199. PMLR, 2024. 3
- [48] Mary Phuong and Christoph Lampert. Towards understanding knowledge distillation. In *ICML*, pages 5142–5151, 2019. 3
- [49] Ziheng Qin, Kai Wang, Zangwei Zheng, Jianyang Gu, Xiangyu Peng, Xu Zhao Pan, Daquan Zhou, Lei Shang, Baigui Sun, Xuansong Xie, and Yang You. Infobatch: Lossless training speed up by unbiased dynamic data pruning. In *ICLR*, 2024. 1
- [50] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *TPAMI*, 39(6):1137–1149, 2017. 4
- [51] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *ICLR*, 2015. 2, 7, 8, 1, 4, 5
- [52] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, pages 4510–4520, 2018. 1, 2
- [53] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, pages 618–626, 2017. 8
- [54] Chaomin Shen, Yaomin Huang, HaoKun Zhu, Jinsong Fan, and Guixu Zhang. Student-oriented teacher knowledge refinement for knowledge distillation. In *ACM MM*, 2024. 2
- [55] Zhiqiang Shen, Zechun Liu, Dejjia Xu, Zitian Chen, Kwang-Ting Cheng, and Marios Savvides. Is label smoothing truly incompatible with knowledge distillation: An empirical study. *ICLR*, 2021. 3
- [56] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 1
- [57] Wonchul Son, Jaemin Na, Junyong Choi, and Wonjun Hwang. Densely guided knowledge distillation using multiple teacher assistants. In *ICCV*, pages 9395–9404, 2021. 3
- [58] Jie Song, Ying Chen, Jingwen Ye, and Mingli Song. Spot-adaptive knowledge distillation. *TIP*, 31:3359–3370, 2022. 1, 5
- [59] Samuel Stanton, Pavel Izmailov, Polina Kirichenko, Alexander A Alemi, and Andrew G Wilson. Does knowledge distillation really work? *NeurIPS*, 34:6906–6919, 2021. 3
- [60] Shangquan Sun, Wenqi Ren, Jingzhi Li, Rui Wang, and Xiaochun Cao. Logit standardization in knowledge distillation. In *CVPR*, pages 15731–15740, 2024. 2, 3, 7, 1
- [61] Jiaxi Tang, Rakesh Shivanna, Zhe Zhao, Dong Lin, Anima Singh, Ed H Chi, and Sagar Jain. Understanding and improving knowledge distillation. *arXiv preprint arXiv:2002.03532*, 2020. 3
- [62] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *ICLR*, 2020. 2, 7, 8, 1, 5
- [63] Frederick Tung and Greg Mori. Similarity-preserving knowledge distillation. In *ICCV*, pages 1365–1374, 2019. 1, 5
- [64] Lu Wang, Liuchi Xu, Xiong Yang, Zhenhua Huang, and Jun Cheng. Debiased distillation for consistency regularization. In *AAAI*, pages 7799–7807, 2025. 4

- [65] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. In *CVPR*, pages 4933–4942, 2019. 2, 4
- [66] Shicai Wei, Chunbo Luo, and Yang Luo. Scaled decoupled distillation. In *CVPR*, pages 15975–15983, 2024. 2, 3, 7, 1
- [67] Sanghyun Woo, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In So Kweon, and Saining Xie. Convnext v2: Co-designing and scaling convnets with masked autoencoders. In *CVPR*, pages 16133–16142, 2023. 1
- [68] Liu Xiaolong, Li Lujun, Li Chao, and Anbang Yao. Norm: Knowledge distillation via n-to-one representation matching. In *ICLR*, 2023. 2
- [69] Liuchi Xu, Jin Ren, Zhenhua Huang, Weishi Zheng, and Yunwen Chen. Improving knowledge distillation via head and tail categories. *TCSVT*, 34(5):3465–3480, 2024. 2, 3
- [70] Mengqi Xue, Jie Song, Xinchao Wang, Ying Chen, Xingen Wang, and Mingli Song. Kdexplainer: A task-oriented attention model for explaining knowledge distillation. In *International Joint Conference on Artificial Intelligence*, pages 3228–3234, 2021. 2
- [71] Chenglin Yang, Lingxi Xie, Chi Su, and Alan L Yuille. Snapshot distillation: Teacher-student optimization in one generation. In *CVPR*, pages 2859–2868, 2019. 3
- [72] Chuanguang Yang, Helong Zhou, Zhulin An, Xue Jiang, Yongjun Xu, and Qian Zhang. Cross-image relational knowledge distillation for semantic segmentation. In *CVPR*, pages 12319–12328, 2022. 1, 2
- [73] Jing Yang, Brais Martinez, Adrian Bulat, Georgios Tzimiropoulos, et al. Knowledge distillation via softmax regression representation learning. In *ICLR*, 2021.
- [74] Zhendong Yang, Zhe Li, Xiaohu Jiang, Yuan Gong, Zehuan Yuan, Danpei Zhao, and Chun Yuan. Focal and global knowledge distillation for detectors. In *CVPR*, pages 4643–4652, 2022.
- [75] Zhendong Yang, Zhe Li, Mingqi Shao, Dachuan Shi, Zehuan Yuan, and Chun Yuan. Masked generative distillation. In *ECCV*, pages 53–69, 2022. 2
- [76] Zhendong Yang, Ailing Zeng, Zhe Li, Tianke Zhang, Chun Yuan, and Yu Li. From knowledge distillation to self-knowledge distillation: A unified approach with normalized loss and customized soft labels. In *ICCV*, pages 17185–17194, 2023. 2, 3, 5, 7, 1
- [77] Linfeng Ye, Shayan Mohajer Hamidi, Renhao Tan, and EN-HUI YANG. Bayes conditional distribution estimation for knowledge distillation based on conditional mutual information. In *ICLR*, 2024. 3
- [78] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven C. H. Hoi. Deep learning for person re-identification: A survey and outlook. *TPAMI*, 44(6):2872–2893, 2022. 2
- [79] Kaiyu Yue, Jiangfan Deng, and Feng Zhou. Matching guided distillation. In *ECCV*, pages 312–328, 2020. 2
- [80] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, pages 87.1–87.12, 2016. 1
- [81] Quanshi Zhang, Xu Cheng, Yilan Chen, and Zhefan Rao. Quantifying the knowledge in a dnn to explain knowledge distillation for classification. *TPAMI*, 45(4):5099–5113, 2022. 3
- [82] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*, pages 6848–6856, 2018. 1, 2
- [83] Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. Deep mutual learning. In *CVPR*, pages 4320–4328, 2018. 3
- [84] Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *CVPR*, pages 11953–11962, 2022. 2, 3, 5, 7, 8, 1, 4
- [85] Borui Zhao, Quan Cui, Renjie Song, and Jiajun Liang. Dot: A distillation-oriented trainer. In *ICCV*, pages 6189–6198, 2023. 3
- [86] Yian Zhao, Wenyu Lv, Shangliang Xu, Jinman Wei, Guanzhong Wang, Qingqing Dang, Yi Liu, and Jie Chen. Detsr beat yolos on real-time object detection. In *CVPR*, pages 16965–16974, 2024. 1, 2
- [87] Kaixiang Zheng and EN-HUI Yang. Knowledge distillation based on transformed teacher matching. In *ICLR*, 2024. 2, 3, 7, 1, 5
- [88] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, pages 1116–1124, 2015. 1
- [89] Zhaohui Zheng, Rongguang Ye, Qibin Hou, Dongwei Ren, Ping Wang, Wangmeng Zuo, and Ming-Ming Cheng. Localization distillation for object detection. *TPAMI*, 45(8):10070–10083, 2023. 1
- [90] Helong Zhou, Liangchen Song, Jiajie Chen, Ye Zhou, Guoli Wang, Junsong Yuan, and Qian Zhang. Rethinking soft labels for knowledge distillation: A bias-variance tradeoff perspective. In *ICLR*, 2021. 3
- [91] Zikai Zhou, Yunhang Shen, Shitong Shao, Huanran Chen, Linrui Gong, and Shaohui Lin. Rethinking centered kernel alignment in knowledge distillation. In *IJCAI*, pages 5680–5688, 2024. 7, 1
- [92] Martin Zong, Zengyu Qiu, Xinzhu Ma, Kunlin Yang, Chunya Liu, Jun Hou, Shuai Yi, and Wanli Ouyang. Better teacher better student: Dynamic prior knowledge for knowledge distillation. In *ICLR*, 2023. 2
- [93] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Wang, Lijuan Wang, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. *NeurIPS*, 36:19769–19782, 2024. 1