

A Good Student is Cooperative and Reliable: CNN-Transformer Collaborative Learning for Semantic Segmentation

Jinjing Zhu¹ Yunhao Luo³ Xu Zheng¹ Hao Wang⁴ Lin Wang^{1,2*}

¹ AI Thrust, HKUST(GZ) ² Dept. of CSE, HKUST ³ Brown University ⁴ Alibaba Cloud, Alibaba Group

zhujinjing.hkust@gmail.com, devinluo27@gmail.com, zhengxu128@gmail.com, cashenry@126.com, linwang@ust.hk

Project Page: <https://vlislab22.github.io/CTCL/>




Figure 1: (a) Our CNN-ViT collaborative learning framework can learn the compact ViT-models and CNN-based models simultaneously while achieving the SoTA segmentation performance than prior methods. (b) We propose the first online KD framework to collaboratively learn compact CNN-based and ViT-based models by selecting and exchanging reliable knowledge between them.

Abstract

In this paper, we strive to answer the question ‘how to collaboratively learn convolutional neural network (CNN)-based and vision transformer (ViT)-based models by selecting and exchanging the reliable knowledge between them for semantic segmentation?’ Accordingly, we propose an online knowledge distillation (KD) framework that can simultaneously learn compact yet effective CNN-based and ViT-based models with two key technical breakthroughs to take full advantage of CNNs and ViT while compensating their limitations. Firstly, we propose heterogeneous feature distillation (HFD) to improve students’ consistency in low-layer feature space by mimicking heterogeneous features between CNNs and ViT. Secondly, to facilitate the two students to learn reliable knowledge from each other, we propose bidirectional selective distillation (BSD) that can dynamically transfer selective knowledge. This is achieved by 1) region-wise BSD determining the directions of knowledge transferred between the corresponding regions in the feature space and 2) pixel-wise BSD discerning which of the prediction knowledge to be transferred in the logit space.

Extensive experiments on three benchmark datasets demonstrate that our proposed framework outperforms the state-of-the-art online distillation methods by a large margin, and shows its efficacy in learning collaboratively between ViT-based and CNN-based models.

1. Introduction

Semantic segmentation [6, 25, 39] is a crucial and challenging vision task, which aims to predict a category label for each pixel in the input image. Although the state-of-the-art (SoTA) segmentation methods have achieved remarkable performance, they often require prohibitive computational costs. This limits their applications to resource-limited scenarios, e.g., autonomous driving [12]. Consequently, growing attention has been paid to model compression aiming at obtaining more compact networks. It can be roughly divided into quantization [10, 13, 38], pruning [4, 26, 32], and knowledge distillation (KD) [27, 30, 33]. The standard KD paradigm aims to learn a compact yet effective student model under the guidance of a high-capacity teacher model. For instance, CD [29] proposes a channel-wise KD approach by normalizing the activation map of

*Corresponding author.

each channel. IFVD [36] characterizes the intra-class feature variation (IFV) and makes the student model mimic the IFV of the teacher model.

Recently, vision transformer (ViT) achieves comparable or even better performance than that of CNNs thanks to the computing paradigm, *e.g.*, multi-head self-attention (MHSA). For instance, PVT [34, 35] and Swin Transformer [9, 24] extract the pyramid features from the high-resolution images and achieve SoTA performance on various benchmarks. To minimize the model complexity, SegFormer [39] proposes a hierarchically structured transformer encoder to learn a simple yet efficient ViT-based model.

In this paper, we strive to collaboratively learn *compact yet effective* CNN-based and ViT-based models for semantic segmentation. Intuitively, we explore an online KD paradigm for this goal. Existing online KD methods for classification employ a ‘Dual-Student’ framework (without the pre-trained model) by enabling the students to learn from each other in a one-stage learning manner [1, 14, 31, 42]. For example, Deep Mutual Learning (DML) [42] proposes to make the CNN-based students teach each other in the training process. KDCL [14] enables students with different capacities to learn collaboratively to generate reliable soft supervision and boost their classification performance. However, naively applying these CNN-based KD methods is less effective and even leads to performance drops (see Fig. 1 (a)). The reasons are that: 1) The discrepancies in the feature and prediction space between CNNs and ViT caused by the distinct computing paradigms make it challenging to perform online KD. 2) These methods only transfer knowledge in the logit space while more reliable and informative knowledge does exist in the *feature space*. 3) There are considerable model size gap and learning capacity gap between CNNs and ViT. Intuitively, we ask a question: ‘*how to collaboratively learn CNN-based and ViT-based models by selecting and exchanging the reliable knowledge between them for semantic segmentation?*’

In light of this, we propose, to the best of our knowledge, the **first** online KD strategy to further push the limit of CNNs and ViT for semantic segmentation (See Fig. 1 (b)). Our method enjoys two key technical breakthroughs. Firstly, we propose heterogeneous feature distillation (**HFD**) to make the students learn the heterogeneous features from each other for complementary knowledge in the low-layer feature space. Concretely, the ViT-based student takes the low-level features from the CNN-based student as guidance and vice versa. Then, consistency between the low-layer features of CNN-based and ViT-based students is imposed to encourage them to compensate for their limitations. Secondly, to transfer reliable knowledge between CNNs and ViT, we propose a bidirectional selective distillation (**BSD**) module that selectively distills the reliable region-wise and pixel-wise knowledge. Specifically,

the region-wise distillation dynamically transfers reliable knowledge of regions in the feature space by determining the directions of transferring knowledge. Similarly, pixel-wise distillation discerns which of the prediction knowledge to be transferred in the logit space. Note that these bidirectional distillation approaches are both guided by the cross entropy between predictions and ground-truth (GT) labels.

In summary, our main contributions are four-fold: **(I)** We introduce the *first* online collaborative learning strategy to collaboratively learn compact ViT-based and CNN-based models for semantic segmentation. **(II)** We propose HFD to facilitate CNNs and ViT learning global and local feature representations correspondingly. **(III)** We propose BSD to distill knowledge between ViT and CNNs in the feature and logit spaces. **(IV)** Our proposed method consistently achieves new state-of-the-art performance on three benchmark datasets for semantic segmentation.

2. Related work

KD for Segmentation. The mainstream methods [1, 2, 16, 18, 21, 23, 29, 36, 40] for segmentation mostly focus on learning a compact CNN-based student model by distilling the knowledge from a cumbersome CNN-based teacher model with the same network architecture. CD [29] proposes a channel-wise KD approach by normalizing the activation map of each channel. IFVD [36] characterizes the intra-class feature variation (IFV) and makes the student mimic the IFV of the teacher. SSTKD [19] exploits the structural and statistical knowledge to enrich low-level information of the student model. *Differently, we propose a first online KD approach, which collaboratively learns compact yet effective CNN-based and ViT-based models for segmentation.*

Vision Transformer has demonstrated its effectiveness on several vision tasks but is less applicable in case of limited computational resources. Recently, several attempts have been made to obtain compact ViT models via network pruning [41] or KD [20]. Moreover, some works [7] combine the advantages of CNNs and ViT, and design hybrid models for classification. *By contrast, we explore simultaneously learning compact yet effective CNNs and ViT models by bidirectionally learning the feature and prediction information from both models for semantic segmentation.*

Online KD. Some works [1, 5, 21, 42] focus on the online KD without a pre-trained teacher model. DML [42] proposes a mutual learning strategy, where an ensemble of students’ logits is deployed, for classification task. Co-distillation [1] further extends this idea and explores the potential in distributed learning. ONE [22] constructs a multi-branch network and assembles the on-the-fly logit information from the branches to enhance the performance on the target network. CLNN [31] proposes multiple generated classifier heads to obtain supplementary information for improving the generalization ability of the target network.

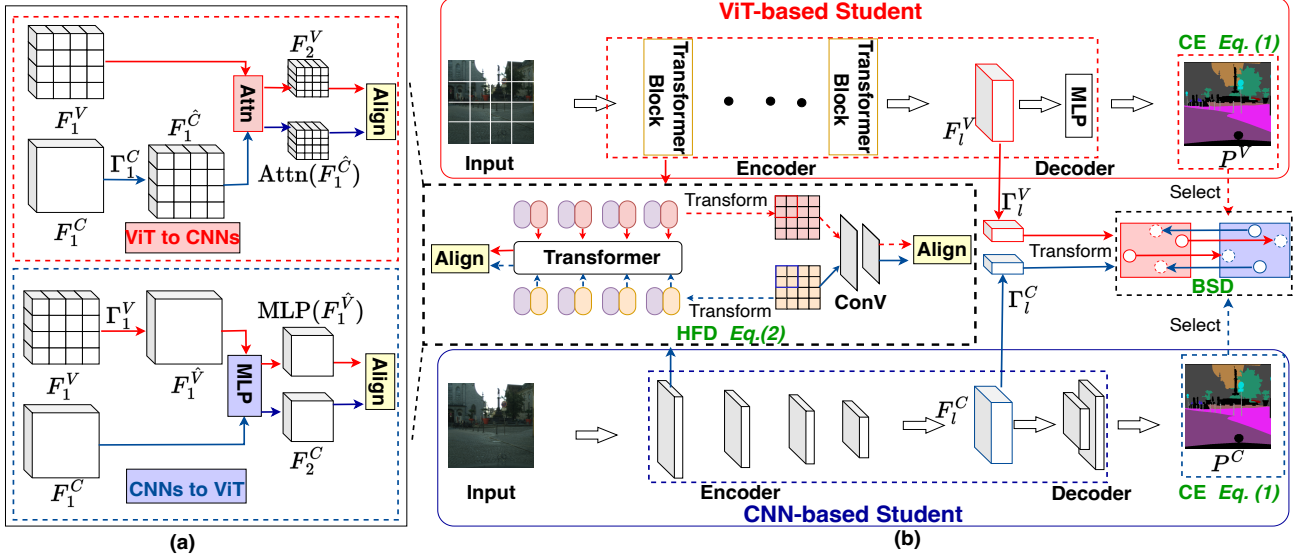


Figure 2: Illustration of the proposed framework, containing three parts: ViT-based student, CNN-based student, and KD modules. (a) Heterogeneous feature distillation (HFD) enables CNN-based and ViT-based students to learn from each other in the low-layer feature space. (b) Our proposed framework. The online KD strategy is optimized via three loss functions: (i) a cross entropy loss, (ii) heterogeneous feature distillation loss, and (iii) a bidirectional selective distillation (BSD) loss.

KDCL [14] aggregates the outputs of numerous students with different learning capacities to generate high-quality labels for supervision. PCL [37] integrates online ensemble and collaborative learning into a unified framework. *Unlike these works generating a soft target in the logit space and transferring knowledge between the isomorphic CNN-based models, we introduce a collaborative learning strategy between the heterogeneous CNN-based and ViT-based models. We propose to bidirectionally exchange the reliable knowledge in feature and logit spaces for semantic segmentation.*

3. The Proposed Approach

3.1. Overview

An overview of the proposed framework is depicted in Fig. 2(b), which consists of three components: a CNN-based student $f(\theta^C)$, a ViT-based student $f(\theta^V)$, and the proposed KD modules. Given an input image set X , our objective is to enable $f(x; \theta^V)$ and $f(x; \theta^C)$ to learn collaboratively that can assign a pixel-wise label $l \in 1, \dots, K$ to each pixel $p_{i,j}$ in image $x \in X (x \in \mathbb{R}^{H \times W \times 3})$ more accurately than the student itself, where H and W are the height and width of x , K is the number of categories. To achieve this goal, given specific input x , we attain the segmentation prediction maps (P^C and P^V) and feature representations (F^C and F^V) from the two students $f(x; \theta^C)$ and $f(x; \theta^V)$, respectively, which can be formulated as:

$$(P^C, F^C) = f(x; \theta^C), \quad (P^V, F^V) = f(x; \theta^V).$$

The pixel-wise segmentation loss $CE(\cdot)$ is based on the cross-entropy (CE) loss with the ground-truth (GT) label :

$$\begin{aligned} \mathcal{L}_{CE}^C &= \frac{1}{H \times W} \sum_{h=1}^H \sum_{w=1}^W CE(\sigma(P^C_{(h,w)}), y_{(h,w)}), \\ \mathcal{L}_{CE}^V &= \frac{1}{H \times W} \sum_{h=1}^H \sum_{w=1}^W CE(\sigma(P^V_{(h,w)}), y_{(h,w)}). \end{aligned} \quad (1)$$

Here, σ is the softmax function, and $y_{(h,w)}$ denotes the GT of the (h, w) -th pixel of image x . Our key ideas are two folds. To compensate for the limitations of CNNs and ViT, we first propose HFD to align the features in the low-layer feature space. Secondly, we propose a BSD module to selectively enable both two students to mimic the region-wise and pixel-wise information from each other. We now describe the technical details.

3.2. Heterogeneous Feature Distillation (HFD)

Inspired by the observations that CNNs are hard-coded to attend only locally while ViT does not learn to attend locally in earlier layers [28], we propose a novel HFD module to make the students learn the heterogeneous features from each other for complementary knowledge in the low-layer feature space. Specifically, it is efficiently achieved by aligning the transformed features between CNNs and ViT (see Fig.2 (a)). At the top of Fig.2 (a), we transfer knowledge from the ViT-based student to the CNN-based student. To match the shapes and channels between the first-

layer features F_1^C of CNN-based student and the first-stage features F_1^V , we utilize a linear transformation Γ_1^C which consists of 1×1 convolution (conv) and pooling layers. And F_1^C is transformed to be $F_1^{\hat{C}} = \Gamma_1^C(F_1^C)$. Then, the second-stage ViT block ‘Attn’ has two inputs: (a) the feature F_1^V and (b) the transformed feature $F_1^{\hat{C}}$ and outputs second-stage feature $F_2^V = \text{Attn}(F_1^V)$ and $\text{Attn}(F_1^{\hat{C}})$. To enable the low-layer features of CNNs to mimic low-layer features of ViT, we align F_2^V and $\text{Attn}(F_1^{\hat{C}})$ by using cosine distance and use the discrepancy to optimize CNNs. Similarly, as shown in the right of figure of Fig. 2(a), we exploit the linear transformation Γ_1^V to match the spatial size of CNNs and ViT. We also utilize a linear transformation Γ_1^V comprising of 1×1 convolution (conv) and pooling layers to transform F_1^V as $F_1^{\hat{V}} = \Gamma_1^V(F_1^V)$. The second-layer of CNNs MLP takes the transformed features $F_1^{\hat{V}}$ and $F_1^{\hat{C}}$ as inputs, and outputs $\text{MLP}(F_1^{\hat{V}})$ and $F_2^C = \text{MLP}(F_2^C)$, correspondingly. Then, aligning these outputs with the cosine distance facilitates ViT-based student to learn from CNN-based student and thus improves the performance of ViT. Finally, we make ViT-based student learn the local feature representations and CNN-based student learn global feature representations by HFD, which is defined as:

$$\begin{aligned} \mathcal{L}_{\text{HFD}}^C &= \cos(\text{Attn}(F_1^{\hat{C}}), F_2^V), \\ \mathcal{L}_{\text{HFD}}^V &= \cos(\text{MLP}(F_1^{\hat{V}}), F_2^C), \end{aligned} \quad (2)$$

here \cos is the cosine distance measuring the consistency between CNNs and ViT. $\text{Attn}(F_1^{\hat{C}})$ is defined as

$$\text{Attn}(F_1^{\hat{C}}) = \text{softmax} \left(\frac{F_1^{\hat{C}} W^Q (F_1^{\hat{C}} W^K)^T}{\sqrt{d}} \right) (F_1^{\hat{C}} W^V),$$

where W^Q , W^K , and W^V are the projections of Attn and d is the number of multi-head in ViT. Similarly, MLP block is the second layer of CNNs.

3.3. Bidirectional Selective Distillation (BSD)

Due to the different performance in different regions between the ViT and CNN students, we intend to dynamically select useful knowledge between the two students in the feature space, so as to benefit each other. However, there is a challenging problem: ‘*how to decide the directions of transferring knowledge for different regions during training?*’. To this end, we propose to manage the directions of KD via combining the predictions and GT labels, where we regard the directions of KD for different regions as a sequential decision-making problem. Consequently, we propose a directional selective distillation (BSD) for *enabling students to learn collaboratively*, as shown in Fig. 3. Our BSD module transfers knowledge in two aspects. Firstly, the region-wise distillation determines the distillation direction of each region for supervising students in each region. Secondly,

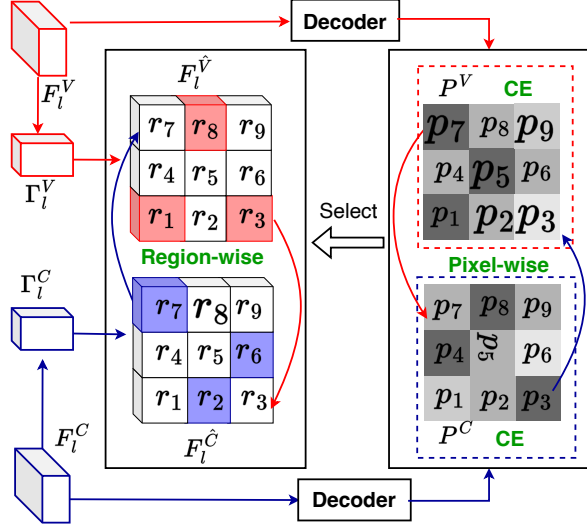


Figure 3: Illustration of the proposed BSD. For region-wise BSD, the colorful cubes mean these regions have more reliable knowledge than the same regions from other students. Then the reliable knowledge is transferred from colorful cubes to white cubes and from more reliable regions to less reliable regions. For pixel-wise BSD, the darker the color of the square, the more accurate the predictions of pixels. Like region-wise, pixel-wise teaches pixels with less accurate predictions from the reliable knowledge of pixels with more accurate predictions.

the pixel-wise distillation decides which of the prediction knowledge to be transferred in the logit space.

3.3.1 Region-wise distillation

Given the last-layer feature $F_l^C \in \mathbb{R}^{\hat{H} \times \hat{W} \times \hat{D}}$ and the last-stage feature F_l^V , we exploit 1×1 conv and pooling layers to transform F_l^C and F_l^V for matching the channels and shapes of them. The transformation functions are denoted as Γ_l^C and Γ_l^V , respectively. Then $F_l^{\hat{C}} = \Gamma_l^C(F_l^C)$ matches the dimensions of $F_l^{\hat{V}} = \Gamma_l^V(F_l^V)$ and $F_l^{\hat{C}} \in \mathbb{R}^{\hat{H} \times \hat{W} \times \hat{D}}$. To transfer knowledge from regions between students, we calculate the cross-student region-wise similarity matrix $S_{(\hat{h}, \hat{w})} = \cos(F_l^{\hat{C}}_{(\hat{h}, \hat{w})}, (F_l^{\hat{V}})_{(\hat{h}, \hat{w})})$. Then, we exploit the cross entropy between the predictions and GT labels to quantify the most reliable knowledge for each region between the students. As shown in Fig. 3, the red grids $\{r_1, r_3, r_8\}$ indicate that the knowledge of these regions are more reliable than the white cubes $\{r_1, r_3, r_8\}$, which means the CE losses of these red regions are smaller than blue regions. Then, the direction of KD for these three regions is from red regions to white regions. Specifically, we utilize the matrix $\hat{m}_{(\hat{h}, \hat{w})} \in \mathbb{R}^{\hat{H} \times \hat{W}}$ which is 0 or 1 to

decides the direction of KD for each region. The value of $\hat{m}_{(\hat{h}, \hat{w})}$ is 1 when the CE of this region in CNN-based prediction map is smaller than that of ViT-based prediction map, enabling the knowledge to be transferred from $F_l^{\hat{C}}(\hat{h}, \hat{w})$ to $F_l^{\hat{V}}(\hat{h}, \hat{w})$, and vice versa for $\hat{m}_{(\hat{h}, \hat{w})}=0$. Note that for matching the size of $F_l^{\hat{C}}$ and P^C , we divide the prediction map into $\hat{H} \times \hat{W}$ size. A $\frac{H}{\hat{H}} \times \frac{W}{\hat{W}}$ sized prediction map P^C at the same location corresponds to one region in $F_l^{\hat{C}}$. After determining the direction of KD for each region, the two students can exchange reliable region-wise knowledge. Therefore, we weight the similarity matrix $S_{(\hat{h}, \hat{w})}$ based on the matrix $\hat{m}_{(\hat{h}, \hat{w})}$, which denotes the process of KD from the more reliable regions to these less reliable regions. To achieve this region-wise KD, we propose to minimize the loss function as follows:

$$\begin{aligned} \mathcal{L}_R^C &= \frac{1}{\hat{H} \times \hat{W} - \hat{M}} \sum_{\hat{h}=1}^{\hat{H}} \sum_{\hat{w}=1}^{\hat{W}} (1 - \hat{m}_{(\hat{h}, \hat{w})}) S_{(\hat{h}, \hat{w})}, \\ \mathcal{L}_R^V &= \frac{1}{\hat{M}} \sum_{\hat{h}=1}^{\hat{H}} \sum_{\hat{w}=1}^{\hat{W}} m_{(\hat{h}, \hat{w})} S_{(\hat{h}, \hat{w})}, \end{aligned} \quad (3)$$

where $\hat{M} = \sum_{\hat{h}=1}^{\hat{H}} \sum_{\hat{w}=1}^{\hat{W}} \hat{m}_{(\hat{h}, \hat{w})}$.

3.3.2 Pixel-wise distillation

Previous KD approaches [29, 36] for semantic segmentation apply the fundamental response-based distillation loss \mathcal{L}_{KL} for the stable gradient descent optimization: $\mathcal{L}_{\text{KL}} = \frac{1}{\hat{H} \times \hat{W}} \sum_{\hat{h}=1}^{\hat{H}} \sum_{\hat{w}=1}^{\hat{W}} \text{KL}(P^C_{(h,w)} \| P^V_{(h,w)})$, where $\text{KL}(\cdot)$ is the Kullback-Leibler divergence (KL divergence) between two probabilities. However, due to the performance gap, the heterogeneous students have their own strengths in predicting different segmentation categories. Therefore, the pixel-wise distillation *aims to transfer the knowledge of more reliable pixel-wise predictions to less reliable pixel-wise predictions in the logit space*. As shown in Fig. 3, the black squares $\{p_1, p_5, p_7\}$ of P^V means that the CE losses of these pixels are smaller than the gray squares $\{p_1, p_5, p_7\}$ of P^C . Therefore, we transfer the reliable knowledge from these more black squares to light black squares. Specifically, we utilize the matrix $m_{(h,w)} \in R^{H \times W}$ which is 0 or 1 to decide the direction of KD for each pixel. The value of $m_{(h,w)}$ is 1 when the CE of pixel from CNN-based student is smaller than that of ViT, enabling the knowledge to be transferred from $p_{(h,w)}^C$ to $p_{(h,w)}^V$, and vice versa for $m_{(h,w)}=0$. Moreover, we use the KL divergence to the effectiveness of transferring knowledge from ViT-based student to CNN-based student: $\text{KL}(P^C_{(h,w)} \| P^V_{(h,w)})$. After determining the direction of KD for each pixel, the two students can exchange useful pixel-wise knowledge. Therefore, we weight the KL divergence $\text{KL}(P^C_{(h,w)} \| P^V_{(h,w)})$ based on the matrix $m_{(h,w)}$, which denotes the process of

KD from the more reliable pixels to these less reliable pixels. To achieve pixel-wise distillation, we propose to minimize the loss function as follows:

$$\begin{aligned} \mathcal{L}_P^C &= \frac{1}{H \times W - M} \sum_{i=1}^H \sum_{j=1}^W (1 - m_{(h,w)}) \text{KL}(P^C_{(h,w)} \| P^V_{(h,w)}), \\ \mathcal{L}_P^V &= \frac{1}{M} \sum_{i=1}^H \sum_{j=1}^W m_{(h,w)} \text{KL}(P^V_{(h,w)} \| P^C_{(h,w)}), \end{aligned} \quad (4)$$

where $M = \sum_{h=1}^H \sum_{w=1}^W m_{(h,w)}$. Finally, combining the region-wise and pixel-wise KD losses, the BSD loss is defined as:

$$\begin{aligned} \mathcal{L}_{\text{BSD}}^C &= \mathcal{L}_R^C + \alpha \mathcal{L}_P^C, \\ \mathcal{L}_{\text{BSD}}^V &= \mathcal{L}_R^V + \alpha \mathcal{L}_P^V, \end{aligned} \quad (5)$$

where α is the trade-parameter to balance the region-wise and pixel-wise losses, and α is set to 1.

3.4. Optimization

Overall, the objectives of the proposed method for CNN-based and ViT-based students are given as

$$\begin{aligned} \mathcal{L}^C &= \mathcal{L}_{\text{CE}}^C + \beta \mathcal{L}_{\text{HFD}}^C + \gamma \mathcal{L}_{\text{BSD}}^C, \\ \mathcal{L}^V &= \mathcal{L}_{\text{CE}}^V + \beta \mathcal{L}_{\text{HFD}}^V + \gamma \mathcal{L}_{\text{BSD}}^V, \end{aligned} \quad (6)$$

where β and γ are hyperparameters and set to 0.1 and 1, respectively.

4. Experiments and Evaluation

4.1. Setup

Datasets. In this work, we conduct extensive experiments to demonstrate the effectiveness of the proposed method on three public datasets: **PASCAL VOC 2012** [11], **Cityscapes** [8], and **CamVid** [3]. Following previous works [29, 36], we adopt the augmented **PASCAL VOC 2012** set [15] consisting of 10,582 training and 1,449 validation images with 21 pixel-wise annotated classes. **Cityscapes** is a dataset for urban scene understanding and consists of 5,000 fine-annotated 1024×2048 images with 19 categories for segmentation. **CamVid** is another widely used urban scene dataset with 11 classes, such as building, tree, sky, car, road, etc., and the 12th class indicates unlabeled data. It contains 367 training, 101 validation, and 233 testing images of 720×960 , where we resize them to 360×480 following previous work.

Implementation and Evaluation We implement our method on PyTorch framework. For CNN-based students, we adopt the widely applied segmentation architecture DeepLabV3+ with encoders of MobileNetV2 and ResNet-50; for ViT-based students, we utilize the efficient SegFormer with encoders of MiT-B1 and MiT-B2, which have comparable or smaller parameters with their CNN counterparts, respectively. Due to the page limit, we put the details in the supplementary.

Dataset	Method	MobileNet	MiT-B1	Δ	ResNet-50	MiT-B2	Δ
PASCAL VOC 2012	Vanilla	67.54	78.48	0.00	76.05	82.03	0.00
	Offline KD	67.40 _{-0.14}	78.87 _{+0.39}	+0.25	76.77 _{+0.68}	82.19 _{+0.16}	+0.88
	DML	67.43 _{-0.11}	78.76 _{+0.28}	+0.17	76.51 _{+0.46}	82.10 _{+0.07}	+0.53
	KDCL	67.41 _{-0.13}	78.76 _{+0.28}	+0.15	76.46 _{+0.41}	82.01 _{-0.02}	+0.39
	IFVD	67.70 _{+0.16}	77.61 _{-0.87}	-0.71	76.52 _{+0.47}	81.52 _{-0.51}	-0.03
	Ours	69.57 _{+2.03}	79.17 _{+0.69}	+2.72	76.99 _{+0.94}	82.67 _{+0.64}	+1.58

Table 1: Comparison with the SoTA KD methods on the **PASCAL VOC 2012** dataset for our CNN-based (MobileNetV2 and ResNet-50) and ViT-based (MiT-B1 and MiT-B2) students.

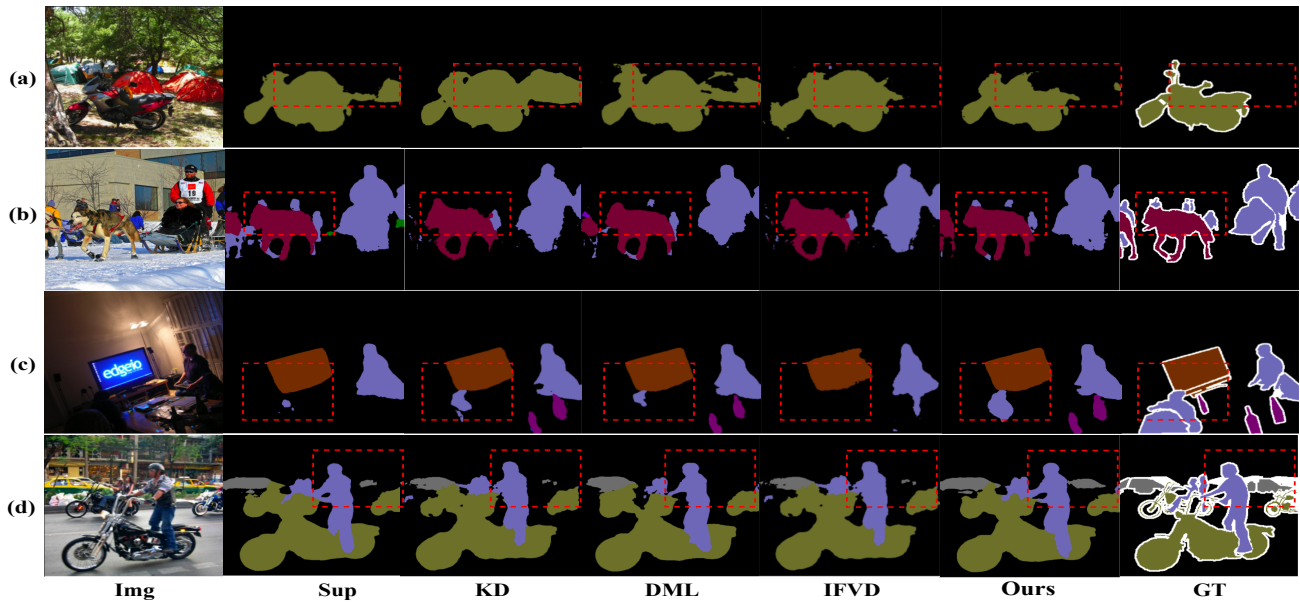


Figure 4: **Visual results on PASCAL VOC 2012.** (a) MobileNetV2, (b) ResNet-50, (c) MiT-B1 and (d) MiT-B2.

In each dataset, CNN-based students are trained by mini-batch stochastic gradient descent (SGD) where the momentum is 0.9, and weight decay is 0.0005; and ViT-based students are trained by AdamW optimizer with a learning rate 0.00006 and weight decay of 0.01. We train Pascal VOC 2012 for 60 epochs with image size 512×512 , where the learning rate for CNN-based models is set to 0.0025 and ViT-based to 0.00006. We evaluate the performance by the mean Intersection over Union (mIoU) score and report our results on the validation sets. We use center-crop evaluation for Pascal VOC 2012 and sliding windows evaluation for Cityscapes. We randomly crop images as 512×512 inputs trained with a batch size of 4. It is worth mentioning that for Offline KD and IFVD approaches, the CNN-based and ViT-based students guide each other as teachers.

4.2. Experiments results

We conduct experiments with two pairs: MobileNetV2 and MiT-B1; ResNet-50 and MiT-B2. It is worth noting that we build two tasks on online KD, which means that students

can transfer knowledge from each other. And we compare our proposed method with some SoTA KD methods Offline KD [17], DML [42], KDCL [14], and IFVD [36]. Furthermore, the Δ is the sum of the performance improvement of each pair compared with Vanilla.

Results on PASCAL VOC 2012. We first evaluate the proposed method on the PASCAL VOC 2012 dataset and report the quantitative results in Tab. 1. Our findings show that existing SoTA KD methods designed for isomorphic models have inferior generalization abilities in online KD between CNN and ViT, compared to the vanilla KD method. In contrast, our proposed approach exhibits significantly better performance. Specifically, our method improves the mIoU of MobileNetV2 and MiT-B1 by **+2.72%**. In contrast, the prior SoTA KD methods DML and KDCL only achieved a mIoU increment of **+0.25%** and **+0.15%**, respectively.

Note that the offline KD method IFVD, designed to transfer knowledge between isomorphic models, impedes the performance of online KD between CNN and ViT. This leads to a drop in performance of **-0.71%** in mIoU. The first

Dataset	Method	MobileNet	MiT-B1	Δ	ResNet-50	MiT-B2	Δ
Cityscapes	Vanilla	73.23	74.95	0.00	76.83	78.77	0.00
	Offline KD	74.11 ^{+0.84}	75.50 ^{+0.55}	+1.43	77.68 ^{+0.85}	78.84 ^{+0.07}	+0.92
	DML	73.68 ^{+0.45}	75.13 ^{+0.18}	+0.63	77.22 ^{+0.39}	78.91 ^{+0.14}	+0.53
	KDCL	73.41 ^{+0.28}	75.51 ^{+0.56}	+0.74	77.94 ^{+1.11}	78.81 ^{+0.04}	+1.15
	IFVD	73.13 ^{-0.10}	75.25 ^{+0.30}	+0.20	77.57 ^{+0.74}	78.90 ^{+0.13}	+0.83
	Ours	74.42^{+1.19}	75.62^{+0.37}	+1.86	78.03^{+1.20}	79.71^{+0.94}	+2.14

Table 2: Comparison with the SoTA KD methods on the **Cityscapes** dataset for our CNN-based (MobileNetV2 and ResNet-50) and ViT-based (MiT-B1 and MiT-B2) students.

Dataset	Method	MobileNet	MiT-B1	Δ	ResNet-50	MiT-B2	Δ
CamVid	Vanilla	71.28	76.10	0.00	73.97	77.04	0.00
	Offline KD	69.48 ^{-1.80}	75.76 ^{-0.34}	-2.14	71.90 ^{-2.07}	77.29 ^{+0.21}	-1.28
	DML	70.73 ^{-0.55}	75.85 ^{-0.25}	-0.80	73.75 ^{-0.22}	77.15 ^{+0.11}	-0.41
	KDCL	71.96 ^{+0.68}	76.40 ^{+0.30}	+0.98	73.19 ^{-0.78}	77.56 ^{+0.51}	-0.26
	IFVD	71.08 ^{-0.20}	75.38 ^{-0.72}	-0.92	74.22 ^{+0.25}	77.25 ^{+0.21}	+0.46
	Ours	73.09^{+1.81}	77.04^{+0.94}	+2.75	75.26^{+1.29}	78.52^{+1.48}	+2.77

Table 3: Comparison with the SoTA KD methods on the **CamVid** dataset for our CNN-based (MobileNetV2 and ResNet-50) and ViT-based (MiT-B1 and MiT-B2) students.

reason mentioned in the introduction causes this outcome. Our proposed method consistently outperforms SoTA KD methods with larger backbone models, ResNet50 and MiT-B2, by achieving a **+1.58%** mIoU increment. This result indicates the superiority of our method, which enables students to learn heterogeneous features from each other to acquire complementary knowledge in the feature space.

Fig. 4 shows the qualitative results and a comparison of the SoTA KD methods on the PASCAL VOC 2012 dataset. Intuitively, the vanilla KD method (3rd column) and DML (4th column) produce unsatisfactory predictions and even erroneous segmentation. In contrast, the results of our method are closer to the ground truth with better segmentation. *These outcomes demonstrate the effectiveness and superiority of the proposed BSD module, which selectively distills the reliable region-wise and pixel-wise knowledge.*

Results on Cityscapes. Tab. 2 presents the quantitative results on Cityscapes validation set. Our proposed method consistently outperforms the SoTA KD methods. In comparison with other online KD methods, our method demonstrates a remarkable increase of mIoU by **+1.86%**, which is much higher than the online DML’s improvement of **+0.63%**, while KDCL and IFVD show only **+0.74%** and **+0.20%** improvements in mIoU, respectively. Our method also outperforms the offline KD methods, Offline KD and IFVD, in terms of segmentation results. Specifically, our method achieves an improvement of **+2.14%** in mIoU with the larger backbone ResNet-50 and MiT-B2, while IFVD and KDCL only show **+0.83%** and **+1.15%** improvements, respectively.

Results on CamVid. Tab. 3 presents a comparison of our proposed method with SoTA KD methods on the CamVid

\mathcal{L}_s	\mathcal{L}_{HFD}	\mathcal{L}_{BSD}	MobileNetV2	MiT-B1	Δ
✓	✗	✗	67.54	78.48	0
✓	✓	✗	67.89 ^{+0.35}	78.78 ^{+0.30}	+0.65
✓	✗	✓	69.19 ^{+1.65}	78.91 ^{+0.43}	+2.08
✓	✓	✓	69.57 ^{+2.03}	79.17 ^{+0.69}	+2.72

Table 4: Ablation of two components of the proposed method evaluated on **PASCAL VOC 2012**.

\mathcal{L}_R	\mathcal{L}_P	MobileNetV2	MiT-B1	Δ
✗	✗	67.54	78.48	0
✓	✗	67.92 ^{+0.38}	78.82 ^{+0.34}	+0.72
✗	✓	68.87 ^{+1.33}	78.81 ^{+0.33}	+1.66
✓	✓	69.19 ^{+1.65}	78.91 ^{+0.43}	+2.08

Table 5: Ablation of two components of BSD evaluated on **PASCAL VOC 2012**.

dataset. The results demonstrate the significant performance improvement of ViT-based and CNN-based student models achieved by our method. In contrast to the students without distillation, our method produces a remarkable improvement of **1.81%**, **0.94%**, **1.29%**, and **1.48%** in MobileNetV2, MiT-B1, ResNet-50, and MiT-B2, respectively. While most of the previous KD methods show decreased performance on this dataset, our method exhibits better generalization and enables the students to learn collaboratively. Additionally, our method outperforms the compared KD methods, regardless of the choice of different architectures and backbones for the student networks.

Method	MobileNetV2	MiT-B2	Δ	ResNet-50	MiT-B1	Δ
Vanilla	67.54	82.03	0.00	76.05	78.48	0.00
Ours	69.21 _{+1.67}	82.27 _{+0.24}	+1.91	77.59 _{+1.54}	79.56 _{+1.08}	+2.62

Table 6: Comparison with the Vanilla methods on the **PASCAL VOC 2012** dataset for our CNN-based (MobileNetV2 and ResNet-50) and ViT-based (MiT-B1 and MiT-B2) students.

α	0.1	0.5	1.0	2.0	5.0	γ	0.1	0.5	1.0	2.0	5.0
MobileNetV2	68.14 _{+1.41}	68.76 _{+1.22}	69.19 _{+1.65}	69.47 _{+1.93}	70.69 _{+3.15}	MobileNetV2	68.57 _{+1.03}	68.92 _{+1.38}	69.03 _{+1.49}	69.08 _{+1.54}	70.12 _{+2.58}
MiT-B1	78.94 _{+0.46}	78.83 _{+0.35}	78.91 _{+0.48}	78.55 _{+0.07}	78.25 _{-0.23}	MiT-B1	79.10 _{+0.62}	78.95 _{+0.47}	78.94 _{+0.46}	78.60 _{+0.12}	78.30 _{-0.18}
Δ	1.06	+1.57	+2.08	+2.00	+2.92	Δ	+1.65	+1.85	+1.95	+1.66	+2.40

Table 7: Sensitivity of α evaluated on **PASCAL VOC 2012**.

Table 9: Sensitivity of γ evaluated on **PASCAL VOC 2012**.

β	0.05	0.1	0.5	1.0
MobileNetV2	67.60 _{+0.06}	67.89 _{+0.35}	67.76 _{+0.22}	67.44 _{-0.10}
MiT-B1	78.70 _{+0.22}	78.78 _{+0.30}	78.69 _{+0.21}	78.43 _{-0.05}
Δ	+0.28	+0.65	+0.43	-0.15

Table 8: Sensitivity of β evaluated on **PASCAL VOC 2012**.

4.3. Ablation study and analysis

Effectiveness of the two proposed modules. We investigate the impact of enabling and disabling the two components of our proposed method on the PASCAL VOC 2012 dataset using MobileNetV2 and MiT-B1. Tab. 4 reports the results of the different student settings. The table shows that both proposed components can enhance the performance of both students, and the selection of reliable knowledge aids better collaborative learning. Specifically, the BSD module improves performance by **2.08%**, demonstrating the effectiveness of selecting reliable knowledge to transfer between heterogeneous models.

Effectiveness of the two components of BSD. Tab. 5 demonstrates the effectiveness of the different components in the BSD module. The CNN-based and ViT-based students with the region-wise distillation module achieve results of **67.92%** and **78.42%**, respectively. The pixel-wise distillation module boosts the improvement to **2.08%**, a substantial enhancement to the students’ performance. We observe that pixel-wise distillation can enhance CNNs with a small capacity guided by ViT with a larger capacity. However, distilling knowledge from CNNs improves ViT slightly, indicating a further potential for exploration.

Sensitivity of α , β , and γ . Tab. 7, 8, and 9 report the mIoU(%) of the students with different ratios of α , β , and γ on the PASCAL VOC 2012 dataset. The students’ encoders are MobileNetV2 and MiT-B1.

As shown in Tab. 7, increasing the importance of α significantly improves the performance of CNNs due to the proposed BSD module’s transfer of reliable knowledge from ViT to CNN. However, the performance of ViT degrades slightly when $\alpha=5.0$. As our goal is to facilitate the

two students’ learning from each other, we choose $\alpha=1.0$ as it presents the best trade-off for both students.

From Tab. 8, it can be seen that aligning the heterogeneous features in the low-layer feature space directly may degrade the students’ performance due to the considerable learning capacity gap between CNNs and ViT. Therefore, we set β as 0.1 to facilitate the students to learn from each other, which leads to an absolute improvement of **+0.65%**, demonstrating that the HFD module enables the students to learn the heterogeneous features from each other for complementary knowledge in the low-layer feature space.

In Tab. 9, it shows that as γ increases, the performance of CNNs is continually improved, but the performance of ViT degrades slightly. Therefore, we set γ to 1.0 as it shows the best trade-off for improving the performance of both CNNs and ViT simultaneously. The results prove that our approach is suitable for situations when there is a significant performance gap between homogeneous students.

Students with different performance ability. To demonstrate the effectiveness of our proposed method, we conduct experiments with two pairs: MobileNetV2 and MiT-B2, and ResNet-50 and MiT-B1. Tab. 6 shows the quantitative results on the PASCAL VOC 2012 dataset. Compared to the vanilla methods, our proposed method achieves a dramatic increase in mIoU by **+1.91%** and **+2.62%**, respectively. The results demonstrate the effectiveness of our proposed method in improving the performance of different heterogeneous students with varying performance abilities. *More details of the method can be found in suppl. material.*

5. Conclusion

This paper presented the **first** online KD framework to collaboratively learn compact yet effective CNN-based and ViT-based models for semantic segmentation. Specifically, we proposed the heterogeneous feature distillation (HFD) module to improve students’ consistency in the low-layer feature space by mimicking heterogeneous features between CNNs and ViT. Then, we also proposed bidirectional selective distillation (BSD) to select reliable region-

wise and pixel-wise knowledge to transfer and enable students to learn from each other better. Comparison with the SoTA KD methods for semantic segmentation shows that our proposed method significantly outperforms these SoTA methods by a large margin and demonstrates our proposed method’s effectiveness for semantic segmentation.

Limitation and future work: Our method has one limitation. The cross-model distillation may lead to an unbalanced performance gain in the online KD training process. That is, if one student’s knowledge is less instructive, the performance of the other student may be marginally improved. Therefore, we will explore the online KD between the heterogeneous ViT-based and CNN-based models with more distinct learning capacities. Moreover, it is promising to extend the proposed collaborative learning paradigm to learn other tasks than semantic segmentation or the cross tasks between depth estimation and semantic segmentation.

Acknowledgement: This joint paper is supported by Alibaba Cloud, Alibaba Group through Alibaba Innovative Research Program, and the National Natural Science Foundation of China (NSF) under Grant No. NSFC22FYT45.

References

- [1] Rohan Anil, Gabriel Pereyra, Alexandre Passos, Róbert Ormándi, George E. Dahl, and Geoffrey E. Hinton. Large scale distributed neural network training through online distillation. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [2] Lucas Beyer, Xiaoohua Zhai, Amélie Royer, Larisa Markeeva, Rohan Anil, and Alexander Kolesnikov. Knowledge distillation: A good teacher is patient and consistent. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 10915–10924. IEEE, 2022.
- [3] Gabriel J. Brostow, Jamie Shotton, Julien Fauqueur, and Roberto Cipolla. Segmentation and recognition using structure from motion point clouds. In *Computer Vision - ECCV 2008, 10th European Conference on Computer Vision, Marseille, France, October 12-18, 2008, Proceedings, Part I*, volume 5302 of *Lecture Notes in Computer Science*, pages 44–57. Springer, 2008.
- [4] Linhang Cai, Zhulin An, Chuanguang Yang, Yangchun Yan, and Yongjun Xu. Prior gradient mask guided pruning-aware fine-tuning. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022*, pages 140–148. AAAI Press, 2022.
- [5] Defang Chen, Jian-Ping Mei, Can Wang, Yan Feng, and Chun Chen. Online knowledge distillation with diverse peers. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 3430–3437. AAAI Press, 2020.
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):834–848, 2018.
- [7] Yinpeng Chen, Xiyang Dai, Dongdong Chen, Mengchen Liu, Xiaoyi Dong, Lu Yuan, and Zicheng Liu. Mobileformer: Bridging mobilenet and transformer. *CoRR*, abs/2108.05895, 2021.
- [8] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 3213–3223. IEEE Computer Society, 2016.
- [9] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 12114–12124. IEEE, 2022.
- [10] Steven K. Esser, Paul A. Merolla, John V. Arthur, Andrew S. Cassidy, Rathinakumar Appuswamy, Alexander Andreopoulos, David J. Berg, Jeffrey L. McKinstry, Timothy Melano, Davis R. Barch, Carmelo di Nolfo, Pallab Datta, Arnon Amir, Brian Taba, Myron D. Flickner, and Dharmendra S. Modha. Convolutional networks for fast, energy-efficient neuromorphic computing. *Proc. Natl. Acad. Sci. USA*, 113(41):11441–11446, 2016.
- [11] Mark Everingham, S. M. Ali Eslami, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *Int. J. Comput. Vis.*, 111(1):98–136, 2015.
- [12] Di Feng, Christian Haase-Schütz, Lars Rosenbaum, Heinz Hertlein, Claudius Gläser, Fabian Timm, Werner Wiesbeck, and Klaus Dietmayer. Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. *IEEE Trans. Intell. Transp. Syst.*, 22(3):1341–1360, 2021.
- [13] Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W. Mahoney, and Kurt Keutzer. A survey of quantization methods for efficient neural network inference. *CoRR*, abs/2103.13630, 2021.
- [14] Qiushan Guo, Xinjiang Wang, Yichao Wu, Zhipeng Yu, Ding Liang, Xiaolin Hu, and Ping Luo. Online knowledge distillation via collaborative learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 11017–11026. Computer Vision Foundation / IEEE, 2020.
- [15] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *2011 international conference on computer vision*, pages 991–998. IEEE, 2011.
- [16] Tong He, Chunhua Shen, Zhi Tian, Dong Gong, Changming Sun, and Youliang Yan. Knowledge adaptation for efficient

- semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 578–587. Computer Vision Foundation / IEEE, 2019.
- [17] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. *CoRR*, abs/1503.02531, 2015.
- [18] Deyi Ji, Haoran Wang, Mingyuan Tao, Jianqiang Huang, Xian-Sheng Hua, and Hongtao Lu. Structural and statistical texture knowledge distillation for semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 16855–16864. IEEE, 2022.
- [19] Deyi Ji, Haoran Wang, Mingyuan Tao, Jianqiang Huang, Xian-Sheng Hua, and Hongtao Lu. Structural and statistical texture knowledge distillation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16876–16885, 2022.
- [20] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. Tinybert: Distilling BERT for natural language understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 4163–4174. Association for Computational Linguistics, 2020.
- [21] Jangho Kim, Minsung Hyun, Inseop Chung, and Nojun Kwak. Feature fusion for online mutual knowledge distillation. In *25th International Conference on Pattern Recognition, ICPR 2020, Virtual Event / Milan, Italy, January 10-15, 2021*, pages 4619–4625. IEEE, 2020.
- [22] Xu Lan, Xiatian Zhu, and Shaogang Gong. Knowledge distillation by on-the-fly native ensemble. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 7528–7538, 2018.
- [23] Yifan Liu, Ke Chen, Chris Liu, Zengchang Qin, Zhenbo Luo, and Jingdong Wang. Structured knowledge distillation for semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 2604–2613. Computer Vision Foundation / IEEE, 2019.
- [24] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, Furu Wei, and Baining Guo. Swin transformer V2: scaling up capacity and resolution. *CoRR*, abs/2111.09883, 2021.
- [25] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 3431–3440. IEEE Computer Society, 2015.
- [26] Christos Louizos, Max Welling, and Diederik P. Kingma. Learning sparse neural networks through l₀ regularization. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [27] Dae Young Park, Moon-Hyun Cha, Changwook Jeong, Daesin Kim, and Bohyung Han. Learning student-friendly teacher networks for knowledge distillation. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 13292–13303, 2021.
- [28] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 12116–12128, 2021.
- [29] Changyong Shu, Yifan Liu, Jianfei Gao, Zheng Yan, and Chunhua Shen. Channel-wise knowledge distillation for dense prediction*. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 5291–5300. IEEE, 2021.
- [30] Wonchul Son, Jaemin Na, Junyong Choi, and Wonjun Hwang. Densely guided knowledge distillation using multiple teacher assistants. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 9375–9384. IEEE, 2021.
- [31] Guocong Song and Wei Chai. Collaborative learning for deep neural networks. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 1837–1846, 2018.
- [32] Dong Wang, Lei Zhou, Xueni Zhang, Xiao Bai, and Jun Zhou. Exploring linear relationship in feature map subspace for convnets compression. *CoRR*, abs/1803.05729, 2018.
- [33] Lin Wang and Kuk-Jin Yoon. Knowledge distillation and student-teacher learning for visual intelligence: A review and new outlooks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(6):3048–3068, 2022.
- [34] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 548–558. IEEE, 2021.
- [35] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. PVT v2: Improved baselines with pyramid vision transformer. *Comput. Vis. Media*, 8(3):415–424, 2022.
- [36] Yukang Wang, Wei Zhou, Tao Jiang, Xiang Bai, and Yongchao Xu. Intra-class feature variation distillation for semantic segmentation. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part VII*, volume 12352 of *Lecture Notes in Computer Science*, pages 346–362. Springer, 2020.
- [37] Guile Wu and Shaogang Gong. Peer collaborative learning for online knowledge distillation. In *Thirty-Fifth AAAI*

Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021, pages 10302–10310. AAAI Press, 2021.

- [38] Jiaxiang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, and Jian Cheng. Quantized convolutional neural networks for mobile devices. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 4820–4828. IEEE Computer Society, 2016.
- [39] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M. Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 12077–12090, 2021.
- [40] Anbang Yao and Dawei Sun. Knowledge transfer via dense cross-layer mutual-distillation. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XV*, volume 12360 of *Lecture Notes in Computer Science*, pages 294–311. Springer, 2020.
- [41] Jinnian Zhang, Houwen Peng, Kan Wu, Mengchen Liu, Bin Xiao, Jianlong Fu, and Lu Yuan. Minivit: Compressing vision transformers with weight multiplexing. *CoRR*, abs/2204.07154, 2022.
- [42] Ying Zhang, Tao Xiang, Timothy M. Hospedales, and Huchuan Lu. Deep mutual learning. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 4320–4328. Computer Vision Foundation / IEEE Computer Society, 2018.