

Cumulative Spatial Knowledge Distillation for Vision Transformers

Borui Zhao¹ Renjie Song¹ Jiajun Liang¹

¹MEGVII Technology

zhaoborui.gm@gmail.com, song.renjie@foxmail.com, liangjiajun@megvii.com

Abstract

*Distilling knowledge from convolutional neural networks (CNNs) is a **double-edged sword** for vision transformers (ViTs). It boosts the performance since the image-friendly local-inductive bias of CNN helps ViT learn faster and better, but leading to two problems: (1) Network designs of CNN and ViT are completely different, which leads to different semantic levels of intermediate features, making spatial-wise knowledge transfer methods (e.g., feature mimicking) inefficient. (2) Distilling knowledge from CNN limits the network convergence in the later training period since ViT’s capability of integrating global information is suppressed by CNN’s local-inductive-bias supervision.*

To this end, we present Cumulative Spatial Knowledge Distillation (CSKD). CSKD distills spatial-wise knowledge to all patch tokens of ViT from the corresponding spatial responses of CNN, without introducing intermediate features. Furthermore, CSKD exploits a Cumulative Knowledge Fusion (CKF) module, which introduces the global response of CNN and increasingly emphasizes its importance during the training. Applying CKF leverages CNN’s local inductive bias in the early training period and gives full play to ViT’s global capability in the later one. Extensive experiments and analysis on ImageNet-1k and downstream datasets demonstrate the superiority of our CSKD. Code: <https://github.com/Zzzzz1/CSKD>

1. Introduction

Recently, Vision Transformers (ViT) attract lots of attention in the computer vision community. Due to the powerful network capacity brought by the self-attention module, ViT achieves great performance on many vision tasks (e.g., image recognition[7, 22] and object detection[20, 3]). Knowledge distillation (KD) is a widely-used technique, which improves the performance of a lightweight model (the student) by transferring knowledge from a heavy one (the teacher). DeiT[34] first applies this technique to ViT, which proposes to distill knowledge from the convolutional neural networks (CNNs) for better training efficiency. The net-

work design of CNN is image-friendly since the convolution operation is translation invariant, *i.e.*, CNN has a local inductive bias which benefits the vision tasks. Thus, applying DeiT could help ViT converge faster and better for vision tasks. However, distilling knowledge from CNN is a double-edged sword. It introduces two problems, which are explained as follows.

Firstly, the network designs of CNN and ViT are different in three aspects: (1) The ways to extend the receptive field, (2) the ways to stack blocks, and (3) the types of normalization layers. It makes intermediate features hard to align in terms of semantics. Thus, feature-based distillation methods could be inefficient due to the misalignment. The inefficiency is worthy of attention because the features can provide spatial-wise supervision with abundant knowledge.

Secondly, CNN’s local inductive bias could hinder the full power of long-range dependencies and positional encoding of ViTs at convergence, although it helps ViT converge faster in the earlier stage (which is also observed in [4]). This phenomenon can be explained as follows: ViT’s network capability is stronger than CNN’s due to its powerful self-attention module. It could predict some hard training samples correctly that CNN can not handle. However, distilling knowledge from CNN’s predictions provides negative supervision, leading ViT to predict wrong labels. With the progress of training, this problem is going to get worse.

Previous works make efforts to solve the problems above. (1) To alleviate the first one, DearKD[4] presents the MHCA module which modifies ViT’s network architecture to align the intermediate features better. Despite the improved performance, extra module design is needed and ViT’s capability of modeling the global relation is weakened. (2) As for the second one, DearKD [4] presents a two-stage learning strategy, and the distillation loss is disabled in the second stage to bypass the problem. In CviT[27], multiple networks with various inductive biases are used as teachers to provide more abundant knowledge, so that ViT’s capability could be less suppressed. Both works boost the distillation performance. However, the former degrades into a baseline training scheme in the later period because it doesn’t apply knowledge distillation, and the latter makes

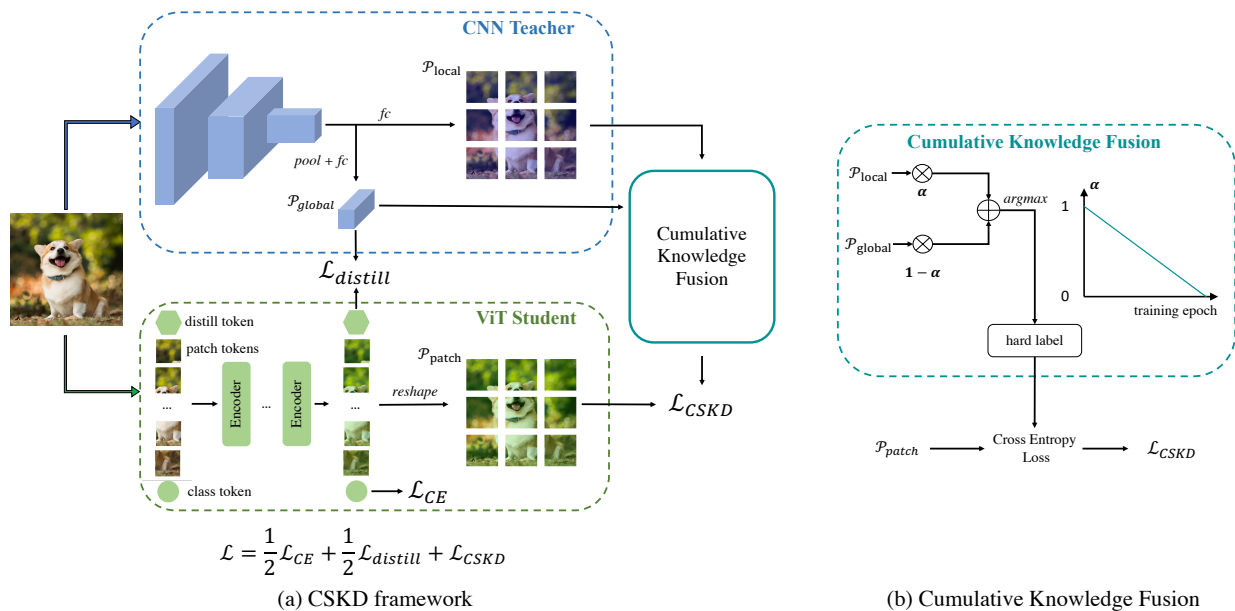


Figure 1. Illustration of our Cumulative Spatial Knowledge Distillation (CSKD).

the training process more computationally intensive.

In this paper, we present Cumulative Spatial Knowledge Distillation (CSKD) to alleviate the aforementioned problems in a simple but effective way. (1) CSKD generates spatial classification predictions from CNN’s last features and uses them as the supervision signals of the corresponding patch tokens of ViT (as shown in Figure 1 (a)). In this way, spatial-wise knowledge is transferred without introducing intermediate features. Thus, the first problem can be alleviated successfully. (2) CSKD exploits a Cumulative Knowledge Fusion (CKF) module which fuses the local and the global responses and increasingly emphasizes the importance of the global ones with the progress of training. (as shown in Figure 1 (b)). In this way, CKF leverages the local inductive bias for fast convergence in the early period and gives full play to ViT’s network capability in the later one. Experimental results on ImageNet-1k and downstream datasets demonstrate the effectiveness of our CSKD. We further provide visualizations to validate that CSKD leverages ViT’s network capability better.

Our contributions are summarized as follows.

- Two obstacles to distilling knowledge from CNN to ViT are discussed: (1) the spatial-wise knowledge transfer is inefficient and (2) the network convergence is limited in the later training period.
- We present Cumulative Spatial Knowledge Distillation (CSKD) to alleviate the two problems in a simple yet effective way. CSKD transfers spatial-wise knowledge without introducing intermediate features

and increasingly emphasizes the importance of global responses for better supervision.

- Extensive experiments validate the effectiveness of our CSKD. We further provide analysis and visualizations, demonstrating that CSKD gives full play to ViT’s powerful network capability.

2. Related Work

2.1. ConvNets and Vision Transformers

Convolutional neural networks (CNN) have been the mainstream on computer vision tasks since [17] firstly presents the convolution operation. Convolution is translation invariant (*i.e.*, with an image-friendly local inductive bias) since the convolution filters are replicated across the entire visual field and share the same parameters. Many works further investigate the network design of CNN. AlexNet[16] firstly propose a deep neural network that performs well on the large-scale ImageNet-1k dataset. GoogLeNet[31] and ResNet[8] further explore deeper and more effective architectures. Meanwhile, MobileNet[11, 30] and ShuffleNet[45, 23] series design lightweight networks to efficiently deploy CNNs on mobile devices.

Transformers are proposed by [35] for the machine translation task. The powerful self-attention module helps the network to model long-distance relations, making transformers the state of the art in many natural language processing (NLP) tasks. Vision transformers (ViTs) are firstly proposed by [7], which proves that the powerful global self-attention module can also boost performance for vision

tasks. Recently, methods introducing more image-friendly inductive bias are proposed for better performance. Swin Transformer[22] presents a shifted windowing scheme to limit self-attention computation to local windows. PVT[37] employs feature pyramid to reduce the computations and benefits dense-prediction vision tasks (*e.g.*, object detection). T2T-ViT[39] recursively aggregates neighboring tokens of an overlapping sliding window, improving the training efficiency. The above works demonstrate that image-friendly design is still important despite ViT’s powerful network capacity.

2.2. Knowledge Distillation

CNN. Knowledge distillation (KD) is proven efficient for transferring knowledge from heavy models (teacher) to lightweight ones (student), achieving great success on many vision tasks. KD is firstly proposed by [10], which introduces the teacher’s predictions as “soft” labels and uses them to supervise the student. Besides such logit-based methods[10, 24, 5, 46], feature-mimicking methods are proposed to leverage the spatial-wise knowledge, since the intermediate features have spatial dimensions. FitNets[28] firstly distill hints from intermediate features. Many works [42, 9, 25, 33] further delve into distilling feature-based knowledge from different perspectives. There are also many works[19, 36, 21, 43] applying KD to more vision tasks(*e.g.*, detection and segmentation).

ViT. To benefit the ViT community with KD, DeiT[34] proposes to distill knowledge from CNNs to vision transformers, using CNN’s local inductive bias for better training efficiency. After that, DearKD[4] proposes a two-stage learning framework, distilling knowledge from CNN’s intermediate features only at the early stage. DearKD learns stronger global representations yet needs extra designs for feature alignment and training schemes. CivT[27] focuses on distilling knowledge from more models with different inductive biases. It introduces both convolution and involution[18] networks as teachers, achieving better performance. Previous works prove the potential of distilling knowledge from different-inductive-bias networks, which motivates us to explore further.

3. Distilling Knowledge from CNN is a Double-edged Sword

3.1. Recap of Vision Transformer Distillation

Vision Transformers Inspired by NLP successes, vision transformers (ViTs) are firstly proposed by [7]. ViT takes an image as several “ 16×16 words” and utilizes transformers for the image recognition task. Firstly, an input image (*e.g.*, with 224×224 resolution and K channels) is split into N discrete patches with patch size P . Each patch is a vector with KP^2 channels. Thus we have $N = 224^2/P^2$.

Then, every patch is mapped to D dimensions as patch embeddings (with position embeddings to retain positional information). An extra learnable embedding named “class token” is introduced for the final classification prediction. ViT feeds these tokens to stacked transformer encoders to model global relations. Every transformer is composed of multi-head self-attention modules (MHSA), MLP blocks, and layer normalization (LN) blocks. At last, ViT takes the output class token from the last encoder as the features for classification and feeds them to a fully-connected layer for the final predictions.

DeiT To further improve the training efficiency, DeiT first proposes to train ViT via transferring knowledge from CNN. It utilizes CNN’s image-friendly inductive bias for better supervision. Concretely, DeiT introduces a new “distillation token” to ViT. The distillation token takes the output logits of the CNN teacher network as the supervision signal. The corresponding loss function $\mathcal{L}_{\text{distill}}$ could be either KL-Divergence (like the traditional KD[10]) or cross-entropy loss (the target label is the hard decision of the teacher). The training loss of DeiT composes of the default cross-entropy loss supervised by the target label (denoted as \mathcal{L}_{CE}) and the aforementioned distillation loss, Eqn. (1).

$$\mathcal{L} = \frac{1}{2}\mathcal{L}_{\text{CE}} + \frac{1}{2}\mathcal{L}_{\text{distill}} \quad (1)$$

As for the inference phase, DeiT takes the average of class token output and distillation token output as the final output. DeiT [34] reveals two phenomena. (1) CNNs are better teachers than ViTs because the inductive bias of CNNs provides more image-friendly guidance. (2) Hard distillation is better than soft distillation¹.

3.2. A Double-edged Sword

DeiT achieves much better training efficiency by distilling knowledge from CNN. However, distilling knowledge from a teacher with such a different inductive bias is a *double-edged sword*. On the bright side, CNN’s local inductive bias helps ViT converge faster and better. On the dark side, the difference in network architectures and network capabilities limits the distillation potential.

Different network architectures. CNN’s network design is completely different from ViT’s. The main differences are as follows: (1) The ways to *extend the receptive field* are different between CNN and ViT. CNN mainly extends its receptive field by downsampling the features step by step (*e.g.*, the feature map of the first stage is 16×16 and the second is 8×8 for spatial dimensions). Stacking

¹Thus, we simply use hard distillation by default in the rest of this paper.

spatial-wise convolution (e.g., 3×3) also provides assistance. ViT extends the network receptive field by introducing self-attention. Self-attention makes ViT have a global receptive field from the first stage. (2) The ways to *stack blocks* are also different. All blocks share the same architectures in ViT, yet CNN’s stage designs are different from each other (e.g., the channel dimensions and the number of bottleneck modules are different among different stages in a ResNet50 [8] network). (3) The *types of normalization layers* are different between CNN and ViT. CNN always uses batch normalization [13] and group normalization [38] while ViT utilizes layer normalization [1].

All the differences above make the intermediate features of CNNs and ViTs have different semantic levels. It makes the intermediate features hard to align. Thus, feature-based distillation methods (i.e., feature-mimicking) could be inefficient due to the misalignment. The inefficiency should be attached much importance to because the features can provide spatial-wise supervision with abundant knowledge. To transfer spatial-wise knowledge, DearKD[4] presents the MHCA module cast in the ViT network architecture, achieving better performance than DeiT. However, it modifies the network structure and limits the capability of global relation modeling.

Different network capability. As is well-known, the excellent performance of ViT is owing to the global receptive field provided by the self-attention module. Relatively speaking, the network capability of CNN is not *that powerful* since there is no module having the global-relation modeling ability in CNN’s architecture. However, stronger teachers are favored in the knowledge distillation community since a better-performance network could transfer more valuable knowledge². Distilling knowledge from CNN to ViT *violates the traditional practice* since the CNN teacher is less powerful. Thus, when the student network is converged, its performance could surpass the teacher’s. At this point, the supervision signal of the teacher could be harmful. Previous works attempt to alleviate this problem from different perspectives. DearKD[4] presents a two-stage learning strategy, deprecating the distillation term in the second stage. It achieves better performance yet wastes the training time in the second stage (i.e., without applying knowledge distillation). CviT[27] uses multiple networks as teachers whose network designs are different (e.g., CNN and INN [18]) to provide knowledge of various inductive biases. It enriches the knowledge yet makes the training more computationally intensive because of the introduction of multiple teachers.

In summary, distilling knowledge from CNN to ViT is a

²There are also some works[40] investigating how to distill knowledge from less-powerful networks. Yet distilling knowledge from better ones is still the mainstream.

double-edged sword. The local inductive bias of CNN helps ViT learn image-friendly knowledge. However, the difference in architectural design makes the intermediate features hard to align, limiting the efficiency of spatial-wise knowledge transfer. Meanwhile, the gap in network capability makes CNN’s supervision harmful to some degree in the later training period.

4. Cumulative Spatial Knowledge Distillation

Two problems of distilling knowledge from CNN to ViT are revealed in the previous section. Driven by the analysis, we present a simple but effective method named Cumulative Spatial Knowledge Distillation (CSKD). CSKD transfers spatial-wise knowledge without introducing intermediate features and alleviates the capability mismatch issue by cumulative learning global-wise knowledge.

Notations We denote the input image as \mathcal{X} and the target label as \mathcal{Y} . The CNN teacher and the ViT student are denoted as \mathcal{T} and \mathcal{S} , respectively. C denotes the number of classes for the image recognition task. For simplification, we omit the batch size of training samples.

4.1. Spatial-wise Knowledge Transfer

Firstly, we generate dense classification predictions to provide spatial-wise knowledge. Dense predictions could transfer spatial-wise information and avoid complex feature aligners like [4]. Ways to generate dense predictions are different between CNN and ViT. As for CNN, the last feature map (with $H \times W$ resolution) is fed to a global average pooling layer and a final fully-connected layer (classifier) to generate the global output logits (denoted as $\mathcal{P}_{global}^{\mathcal{T}} \in \mathbb{R}^{1 \times C}$). We remove the pooling operation and take the features of each position as an *independent sample*. Then, we feed them to the classifier, generating dense predictions of every corresponding position. We denote the dense predictions from CNNs as $\mathcal{P}^{\mathcal{T}} \in \mathbb{R}^{1 \times H \times W \times C}$. As for ViTs, we take every patch token from the last transformer encoder as the final feature of the corresponding position. We further feed the final features to the fully-connected classifier to generate dense predictions. Dense predictions³ from ViTs are denoted as $\mathcal{P}^{\mathcal{S}} \in \mathbb{R}^{1 \times H \times W \times C}$ ⁴. Since every patch token is embedded with the corresponding sub-image and position information, regarding them as spatial dense predictions is reasonable. At last, we take the hard decisions (i.e., *argmax*) of dense predictions from CNN as target labels (denoted as $\mathcal{Y}^{\mathcal{T}}$) and use cross-entropy (CE) as the loss function for spatial knowledge transfer.

³The shape of ViT’s dense predictions is (1, N, C) by default. Details about aligning them to the same shape of CNN’s predictions are attached to the supplement.

⁴In Figure 1, $\mathcal{P}^{\mathcal{T}}$ and $\mathcal{P}^{\mathcal{S}}$ are rewritten as \mathcal{P}_{local} and \mathcal{P}_{patch} respectively for easy understanding.

4.2. Cumulative Knowledge Fusion

To eliminate the effect of capacity mismatch, we introduce a Cumulative Knowledge Fusion (CKF) module. According to [4], CNN helps ViT converge better in the early stage of training due to the different inductive biases, but limits the network convergence in the later period. As aforementioned, DearKD presents a two-stage distillation framework and disables the knowledge distillation to mitigate this problem. In contrast, we believe that CNN could also help ViT learn better representations in the later period. The key factor to achieve this goal is *using better supervision signals*. In the later training period, each patch token of ViT has already modeled global relations among other tokens. Thus, supervising the tokens with global information instead of local information could help them leverage their capability and learn better representations. Thus, we present the Cumulative Knowledge Fusion module (CKF). As illustrated in Figure 1 (b), CKF fuses CNN’s global (\mathcal{P}_{global}^T) and local (\mathcal{P}^T) responses during the entire training process. It prefers local ones in the early stage and global ones in the later stage. CKF generates the dense target labels according to :

$$\mathcal{Y}^T = \underset{c}{\operatorname{argmax}} \left(\alpha \mathcal{P}^T + (1 - \alpha) \mathcal{P}_{global}^T \right), \quad (2)$$

where α is a coefficient that negatively correlates with the training duration (e.g., $\alpha = 1 - t/t_{max}$, where t denotes the current epoch and t_{max} denotes the number of total epochs). Thus, the loss function of our CSKD can be written as :

$$\mathcal{L}_{CSKD} = \operatorname{CE}(\mathcal{P}^S, \mathcal{Y}^T) \quad (3)$$

As Eqn. (2) suggests, the more training time increases the more global knowledge is transferred. Even if there is a capacity mismatch between CNNs and ViTs, supervising local patches with global knowledge could alleviate the problem and further improve the distillation performance. We explore visualizations in Section 5.5 to verify it.

In summary, CSKD transfers spatial-wise knowledge through the patch tokens of ViT without introducing feature mimicking. It also benefits ViT in the later training period by increasingly emphasizing the importance of global responses. We keep the cross-entropy loss and the distillation loss in DeiT [34]. The total loss can be written as:⁵

$$\mathcal{L} = \frac{1}{2} \mathcal{L}_{CE} + \frac{1}{2} \mathcal{L}_{distill} + \mathcal{L}_{CSKD} \quad (4)$$

Algorithm 1 provides CSKD’s PyTorch-style pseudocode.

⁵There is a default hyperparameter to weight \mathcal{L}_{CSKD} . Since we conduct ablation studies on this hyperparameter and validate the robustness among different values in Table 9, it is omitted for simplification.

Algorithm 1 Pseudocode of CSKD in PyTorch.

```
# f_cnn: CNN features before global pooling
# f_vit: ViT features of all patch tokens
# cls_cnn: classifier of the teacher
# cls_vit: classifier of the student
# T: the temperature for KD

# get CNN's global & dense predictions
b, c1, h, w = f_cnn.shape
p_global = cls_cnn(f_cnn.mean(dim=[2,3]))
f_cnn = f_cnn.permute(0, 2, 3, 1)
f_cnn = f_cnn.reshape(b, -1) # (b*h*w, c)
p_cnn = cls_cnn(f_cnn)

# get ViT's dense predictions
b, c2, hw = f_vit.shape
f_vit = f_vit.permute(0, 2, 1)
f_vit = f_vit.reshape(b, -1)
p_vit = cls_vit(f_vit)

# calculate loss_cskd
alpha = 1 - epoch / max_epoch
p_fusion = alpha*p_cnn + (1-alpha)*p_global
target = p_fusion.argmax(dim=1)

loss_cskd = CrossEntropyLoss(p_vit, target)
```

5. Experiments

In this section, we mainly conduct experiments on ImageNet-1k[29] to evaluate CSKD. Besides, experiments on transfer learning to downstream tasks validate the generalization ability of CSKD. We also provide ablation studies and visualizations for better understanding.

5.1. Implementation Details

Our implementation *exactly* follows the practice in DeiT[34]. Concretely, we use a RegNetY-16GF CNN model as the teacher, which achieves 82.9% top-1 accuracy on the ImageNet-1k validation set. All models are trained for 300 epochs with 1024 batch size. We use AdamW as the optimizer with a 0.001 initial learning rate, and cosine decay as the learning scheduler (with a warmup for the first 5 epochs). The weight decay is set as 0.05. Input images are at resolution 224×224 . Mixup[44], Cutmix[41], and Auto-Augment[6] are utilized for data augmentation, as the same as DeiT[34]. Stochastic depth[12] is not employed since we already apply knowledge distillation which is stronger regularization.

As the same as DeiT[34]⁶, we also train tiny, small, and base models named CSKD-Ti, CSKD-S, and CSKD-B, respectively. Table 1 summarizes the model information.

⁶Since we only consider the knowledge distillation setting, all “DeiT” represent the distilled ones for simplification in the rest of the paper.

model	#heads	#layers	#params	throughput
CSKD-Ti	3	12	5M	2536
CSKD-S	6	12	22M	940
CSKD-B	12	12	86M	292

Table 1. Variants of the model architecture. The larger model DeiT-B and CSKD-B have the same architecture as the ViT-B. Embedding dimensions and the number of heads are different among different variants. Dimension per head is always 64. The throughput (img/sec) is measured at resolution 224×224 .

model	#params	image size	val top-1	val-real top-1
<i>CNNs</i>				
ResNet18[8]	12M	224^2	69.8	77.3
ResNet50[8]	25M	224^2	76.2	82.5
ResNet101[8]	45M	224^2	77.4	83.7
ResNet152[8]	60M	224^2	78.3	84.1
RegNetY-4GF*[26]	21M	224^2	80.0	86.4
RegNetY-8GF*[26]	39M	224^2	81.7	87.4
RegNetY-16GF*[26]	84M	224^2	82.9	88.1
EffiNet-B0[32]	5M	224^2	77.1	84.1
EffiNet-B3[32]	12M	224^2	81.6	86.8
EffiNet-B6[32]	43M	224^2	84.0	88.8
<i>ViTs</i>				
ViT-B/16[7]	86M	384^2	77.9	83.6
ViT-L/16[7]	307M	384^2	76.5	82.2
DeiT-Ti[34]	6M	224^2	74.5	82.1
DeiT-S[34]	22M	224^2	81.2	86.8
DeiT-B[34]	87M	224^2	83.4	88.3
DearKD-Ti[4]	5M	224^2	74.8	-
DearKD-S[4]	22M	224^2	81.5	-
DearKD-B[4]	86M	224^2	83.6	-
CSKD-Ti	6M	224^2	76.3	83.6
CSKD-S	22M	224^2	82.3	87.9
CSKD-B	87M	224^2	83.8	88.6

Table 2. Results on ImageNet-1k. * means that RegNets are optimized with similar optimization procedures as DeiT, serving as teachers for DeiT and our CSKD. “val-real” represents the results of the ImageNet Real validation set.

5.2. Results on ImageNet-1k

We mainly evaluate our CSKD on the ImageNet-1k[29] dataset. ImageNet-1k is a well-known large-scale dataset that consists of 1000 classes. The training set contains 1.28 million images and the validation set contains 50k images. We also use ImageNet Real[2] for further evaluation, which corrects some wrong labels of the validation set.

Table 2 reports the comparison results on ImageNet-1k and demonstrates the effectiveness of CSKD. CSKD-Ti, CSKD-S, and CSKD-B improve the performance by +1.8%, +1.1%, and +0.4% respectively. And the results on ImageNet Real also verify the performance gain. Besides, we train CSKD-Ti for a longer time (*i.e.*, 1000 epochs). As shown in Table 3, CSKD-Ti-1k achieves 78.1% top-1 accuracy, surpassing the baseline by a large margin,

further proving the effectiveness.

model	DeiT-Ti-1k	DearKD-Ti-1k	CSKD-Ti-1k
val top-1	76.6	77.0	78.1

Table 3. We train our CSKD-Ti for 1000 epochs (CSKD-Ti-1k). CSKD-Ti achieves 78.1% top-1 accuracy, surpassing DeiT and DearKD by a large margin.

5.3. Transfer Learning to Downstream Datasets

We conduct experiments on transfer learning to downstream datasets. CIFAR10[15], CIFAR100[15], Cars[14] and iNat19 are considered. Table 4 reports the dataset information. As shown in Table 5, our CSKD achieves significantly better transfer performances than DeiT and DearKD on all the downstream datasets. Especially for the challenging iNat19 fine-grained classification task, CSKD-Ti, CSKD-S, and CSKD-B surpass the corresponding baselines by +1.1%, +0.8%, and +0.4% respectively.

dataset	training size	val size	#classes
CIFAR10	50,000	10,000	10
CIFAR100	50,000	10,000	100
Stanford Cars	8,144	8,041	196
iNaturalist 2019	265,240	3,030	1,010

Table 4. Downstream dataset information for transfer learning.

	CIFAR10	CIFAR100	Cars	iNat19
DeiT-Ti*	98.1	86.1	92.1	76.6
DearKD-Ti	97.5	85.7	89.0	-
CSKD-Ti	98.5	87.0	93.1	77.7
DeiT-S*	98.7	89.2	91.7	80.9
DearKD-S	98.4	89.3	91.3	-
CSKD-S	99.1	90.3	93.7	81.7
DeiT-B*	99.1	91.3	92.9	82.1
DearKD-B	99.2	91.1	92.7	-
CSKD-B	99.3	91.4	94.0	82.5

Table 5. Transfer performance in downstream tasks. * represents that the results are based on our implementation.

5.4. Ablation Study

We ablate key elements of our design in CSKD. All models are trained based on CSKD-B following the settings in Section 5.1. First, we investigate how distilling spatial-wise knowledge and CKF module work. Table 6 reports the results of (1) the DeiT baseline and our CSKD that distills knowledge from (2) the local responses, (3) the global ones, (4) combining both by a simple summation, and (5) combining the both with CKF. The results indicate that: distilling spatial-wise knowledge can improve the performance (81.2% vs. 81.6%). Moreover, introducing CKF significantly improves the performance from 81.5% (simply combining global and local responses) to 82.3%.

local	global	CKF	val top-1
			81.2
✓			81.6(+0.4)
	✓		81.4(+0.2)
✓	✓		81.5(+0.3)
✓	✓	✓	82.3(+1.1)

Table 6. Ablation study on distilling knowledge from CNN’s local and global responses. The great effect of our CKF is also verified (81.5% vs. 82.3).

We further explore different loss types. In traditional knowledge distillation methods, KL-Divergence is the commonly used loss function for logit-based distillation (*i.e.*, soft distillation) instead of the cross-entropy loss (*i.e.*, hard distillation). Results in Table 7 demonstrate that hard distillation is more suitable for distilling knowledge from CNN to ViT, while the difference is marginal between them (82.3% vs. 82.1%).

distillation type	hard	soft	soft+hard
val top-1	82.3	82.1	82.2

Table 7. Ablations on different distillation loss types.

We study different strategies for decaying α in the CKF module. We use linear decay, cosine decay, and parabolic decay for comparison. Results in Table 8 show that our CKF is insensitive to the decaying strategy since the results are pretty much the same (82.3% vs. 82.2%).

strategy	val top-1
$1 - t/t_{max}$	82.3
$\cos(\pi/2 \cdot t/t_{max})$	82.2
$(1 - t/t_{max})^2$	82.2

Table 8. Different strategies for decaying α .

loss weight	0.5	1.0	2.0
val top-1	82.1	82.3	82.1

Table 9. Applying CSKD with different loss weight.

Lastly, we conduct experiments on tuning loss weights for CSKD. The results in Table 9 indicate that weighing CSKD in a range of $\{0.5, 2.0\}$ achieves similar performance and 1.0 is the best one. In summary, the ablation studies above prove the effectiveness of CSKD and CKF, and also reveal that CSKD is hyperparameter insensitive.

5.5. Visualizations and Analysis

In this part, we provide more analysis for a better understanding of CSKD’s working mechanism. We mainly study attention distance, attention heatmap, and accuracy dynamics. Visualizations demonstrate the effectiveness of our CSKD and the CKF module.




Figure 2. Mean attention distance of DeiT-B and CSKD-B. The final stages of CSKD have significantly larger attention distances than DeiT, which suggests that CSKD leverages the global capacity of ViTs in a better way.




Figure 3. Validation accuracy dynamics. It indicates that CKF mainly works in the later training period.

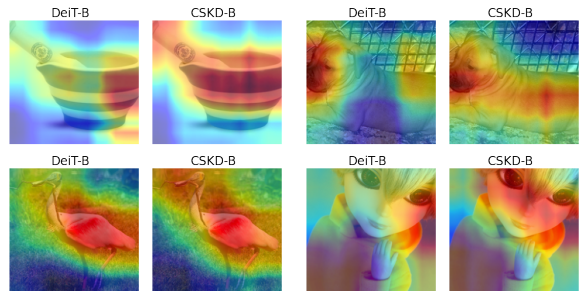


Figure 4. Attention Heatmaps from DeiT-B and our CSKD-B. CSKD-B focuses more attention on the salient object.

Attention Distance. Firstly, we visualize the attention distance of each transformer encoder. Following ViT[7], the attention distance is computed for 128 example images by averaging the distance between the query pixel and all other pixels weighted by the attention weight. A larger attention distance means that the global capability is better leveraged. As shown in Figure 2, the attention distance of CSKD-B is higher than the baseline, especially for the latter stages. It suggests that CSKD makes ViT reach its full potential. We further explore the attention distance for the downstream tasks to validate that the learned global capability is generalizable and transferable in Figure 5.

Accuracy Dynamics. We investigate how CKF influences

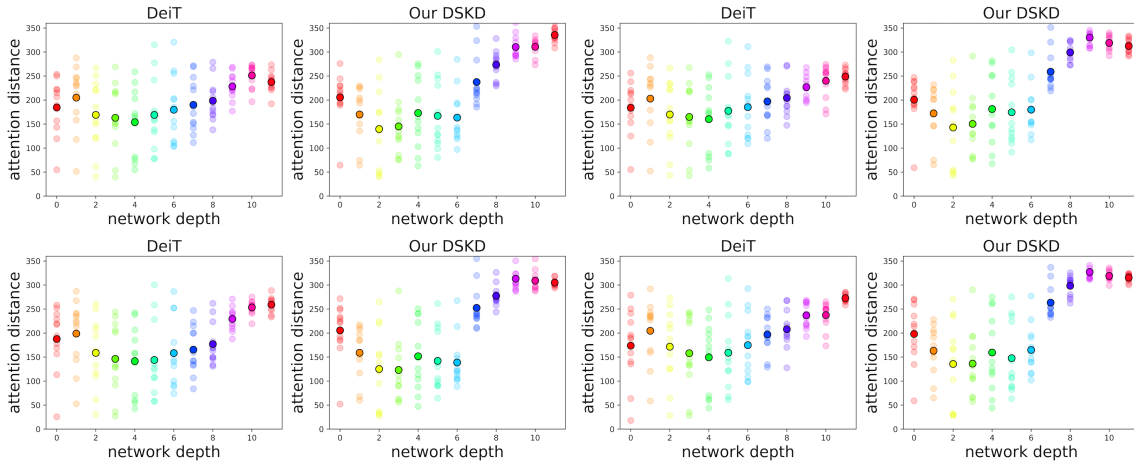


Figure 5. **Mean attention distance** of DeiT-B and CSKD-B for downstream tasks. Results indicate that the global capability of CSKD is generalizable and transferable since the attention distance of DeiT-B is always larger.

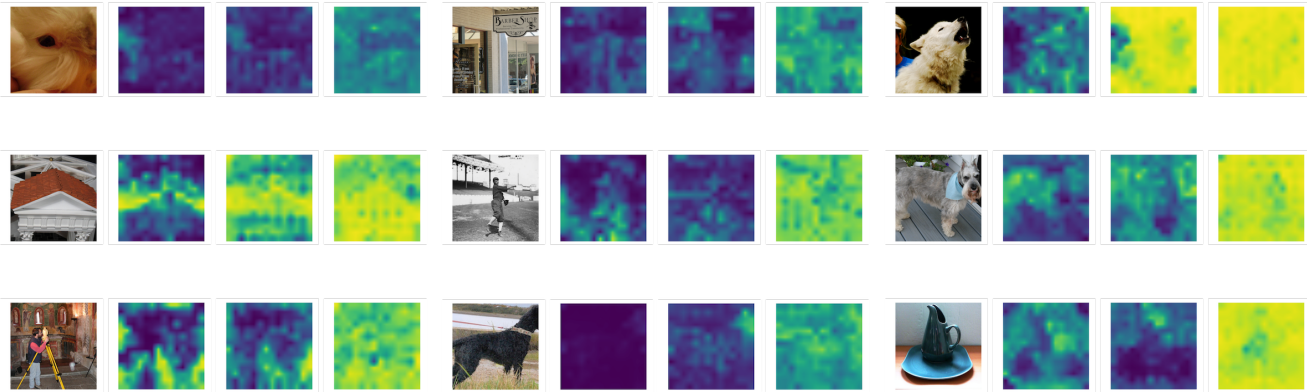


Figure 6. Responses of patch tokens of different epochs. The first column is the input image. The second, third, and fourth ones are response heatmaps of the 100th, 200th, and 300th epochs, respectively.

the training process by visualizing the validation accuracy dynamics. As shown in Figure 3, our CKF module mainly works in the later training period *as designed*. The performance gap becomes larger in the latter epochs. It suggests that our CKF gives full play to ViT’s network capability when the network is almost converged.

Attention Heatmaps. We visualize the attention heatmaps to validate the superiority of our CSKD. Visualizations in Figure 4 demonstrate that the attention heatmaps of CSKD are better than the baseline, *i.e.*, CSKD focuses more attention on the salient object of the image.

Responses Dynamics of Patch Tokens. We visualize the response dynamics of patch tokens during the training process. As shown in Figure 6, CKF helps patch tokens predict more accurate labels with the progress of training, *i.e.*, the global relation is modeled better and better.

6. Conclusion

Distilling knowledge from CNN to ViT achieves great success. In this paper, we reveal two problems limiting the distillation performance. On one hand, spatial-wise knowledge transfer is inefficient since the features are hard to align between CNN and ViT. On the other hand, ViT’s network capability is limited in the later training period. To address these problems, we present Cumulative Spatial Knowledge Distillation (CSKD). CSKD distills knowledge to patch tokens from the corresponding local responses of CNN, transferring spatial-wise knowledge without introducing intermediate features. Moreover, CSKD introduces a Cumulative Knowledge Fusion (CKF) module, which benefits ViT from the local-inductive-bias knowledge in the early period and gives full play to ViT’s capability later. Experiments and analysis on ImageNet-1k and downstream tasks demonstrate the superiority of CSKD.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 4
- [2] Lucas Beyer, Olivier J Hénaff, Alexander Kolesnikov, Xiao-hua Zhai, and Aäron van den Oord. Are we done with imagenet? *arXiv preprint arXiv:2006.07159*, 2020. 6
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCVur*. Springer, 2020. 1
- [4] Xianing Chen, Qiong Cao, Yujie Zhong, Jing Zhang, Shenghua Gao, and Dacheng Tao. Dearth: Data-efficient early knowledge distillation for vision transformers. In *CVPR*, 2022. 1, 3, 4, 5, 6
- [5] Jang Hyun Cho and Bharath Hariharan. On the efficacy of knowledge distillation. In *ICCV*, 2019. 3
- [6] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *CVPR*, 2019. 5
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. 2020. 1, 2, 3, 6, 7
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2, 4, 6
- [9] Byeongho Heo, Jeessoo Kim, Sangdoon Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *ICCV*, 2019. 3
- [10] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *arXiv:1503.02531*, 2015. 3
- [11] 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. *arXiv preprint arXiv:1704.04861*, 2017. 2
- [12] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *ECCV*, 2016. 5
- [13] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*. PMLR, 2015. 4
- [14] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 554–561, 2013. 6
- [15] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 2017. 2
- [17] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998. 2
- [18] Duo Li, Jie Hu, Changhu Wang, Xiangtai Li, Qi She, Lei Zhu, Tong Zhang, and Qifeng Chen. Involution: Inverting the inheritance of convolution for visual recognition. In *CVPR*, 2021. 3, 4
- [19] Quanquan Li, Shengying Jin, and Junjie Yan. Mimicking very efficient network for object detection. In *CVPR*, 2017. 3
- [20] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. *arXiv preprint arXiv:2203.16527*, 2022. 1
- [21] Yifan Liu, Ke Chen, Chris Liu, Zengchang Qin, Zhenbo Luo, and Jingdong Wang. Structured knowledge distillation for semantic segmentation. In *CVPR*, 2019. 3
- [22] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 1, 3
- [23] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet V2: Practical guidelines for efficient cnn architecture design. In *ECCV*, 2018. 2
- [24] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. Improved knowledge distillation via teacher assistant. In *AAAI*, volume 34, 2020. 3
- [25] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *CVPR*, 2019. 3
- [26] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *CVPR*, 2020. 6
- [27] Sucheng Ren, Zhengqi Gao, Tianyu Hua, Zihui Xue, Yonglong Tian, Shengfeng He, and Hang Zhao. Co-advise: Cross inductive bias distillation. In *CVPR*, 2022. 1, 3, 4
- [28] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *ICLR*, 2015. 3
- [29] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. ImageNet large scale visual recognition challenge. *IJCV*, 2015. 5, 6
- [30] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobilenetV2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018. 2
- [31] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. 2
- [32] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019. 6
- [33] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *ICLR*, 2020. 3
- [34] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, 2021. 1, 3, 5, 6

- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 2
- [36] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. In *CVPR*, 2019. 3
- [37] 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 *ICCV*, 2021. 3
- [38] Yuxin Wu and Kaiming He. Group normalization. In *ECCV*, 2018. 4
- [39] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zi-Hang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *ICCV*, 2021. 3
- [40] Li Yuan, Francis EH Tay, Guilin Li, Tao Wang, and Jiashi Feng. Revisiting knowledge distillation via label smoothing regularization. In *CVPR*, 2020. 4
- [41] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *CVPR*, 2019. 5
- [42] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2017. 3
- [43] Feng Zhang, Xiatian Zhu, and Mao Ye. Fast human pose estimation. In *CVPR*, 2019. 3
- [44] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 5
- [45] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*, 2018. 2
- [46] Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *CVPR*, 2022. 3