Sample-level Adaptive Knowledge Distillation for Action Recognition

Ping Li, Chenhao Ping, Wenxiao Wang, and Mingli Song

Abstract—Knowledge Distillation (KD) compresses neural networks by learning a small network (student) via transferring knowledge from a pre-trained large network (teacher). Many endeavours have been devoted to the image domain, while few works focus on video analysis which desires training much larger model making it be hardly deployed in resource-limited devices. However, traditional methods neglect two important problems, i.e., 1) Since the capacity gap between the teacher and the student exists, some knowledge w.r.t. difficult-to-transfer samples cannot be correctly transferred, or even badly affects the final performance of student, and 2) As training progresses, difficult-to-transfer samples may become easier to learn, and vice versa. To alleviate the two problems, we propose a Sample-level Adaptive Knowledge Distillation (SAKD) framework for action recognition. In particular, it mainly consists of the sample distillation difficulty evaluation module and the sample adaptive distillation module. The former applies the temporal interruption to frames, i.e., randomly dropout or shuffle the frames during training, which increases the learning difficulty of samples during distillation, so as to better discriminate their distillation difficulty. The latter module adaptively adjusts distillation ratio at sample level, such that KD loss dominates the training with easy-to-transfer samples while vanilla loss dominates that with difficult-to-transfer samples. More importantly, we only select those samples with both low distillation difficulty and high diversity to train the student model for reducing computational cost. Experimental results on two video benchmarks and one image benchmark demonstrate the superiority of the proposed method by striking a good balance between performance and efficiency.

Index Terms—Knowledge distillation; action recognition; model compression; sample-level distillation.

I. Introduction

DEEP neural networks [1], [2] have exhibited its great power in image and video analysis. However, their model size grows up rapidly as the advanced technologies are updated continuously, and this problem becomes more serious in video analysis including action recognition as a fundamental task, leading to heavy computational cost and much difficulty in deploying them on resource-constrained devices. Hence, compressing neural networks is very necessary in highly-demanding environment, and Knowledge Distillation (KD) [3] has established itself as a useful model compression technique. Essentially, it transfers knowledge (e.g., outputs, gradients, intermediate features) from a powerful large network (i.e., teacher) to a lightweight small network (i.e., student) learned on training samples.

While KD methods [3]–[5] have been widely used in image analysis, there are only very few works exploring KD for
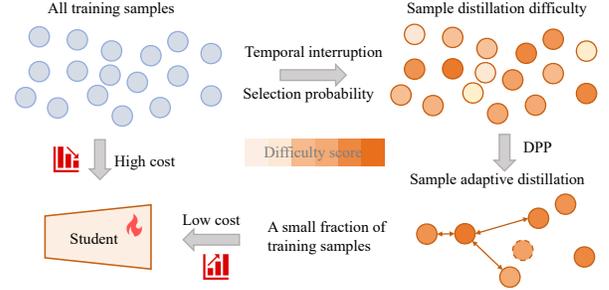
Fig. 1: Illustration of motivation. We first determine the difficulties of all samples by adopting the temporal interruption strategy and considering the sample selection probability. Then, only a small fraction of training samples are chosen in each epoch for a much lower cost and a better student performance.

action recognition. Some early works [6], [7] adopt cross-modal distillation by transferring knowledge from optical-flow network (teacher) to RGB network (student), and one recent work [8] adopts a pre-trained generative model (i.e., conditional variational autoencoder) for distilling attention-based feature knowledge. These methods all explore the early Convolutional Neural Network (CNN)-based action recognition models [9], [10] using both RGB and optical flow features, which is time-consuming, and they fail to consider the temporal dynamics in distillation. Importantly, they overlook **two** critical facts, i.e., **1)** Since the capacity gap between the teacher and the student exists, some knowledge *w.r.t.* difficult-to-transfer samples cannot be correctly transferred, or even badly affects the final performance of student; and **2)** As training progresses, difficult-to-transfer samples may become easier to learn, and vice versa. Meanwhile, when all samples are used during student training, it incurs large computational overheads. So it is necessary to select those samples more benefiting distillation to save cost, as illustrated in Fig. 1. For example, we choose only 10% samples to achieve comparable or even better performance than previous KD methods using all samples with SlowFast [11] model on UCF101 [12] dataset, while the training time is greatly reduced to only one-fourth of previous ones.

This work addresses the above issues from sample-level distillation perspective, and introduces two concepts, i.e., *sample distillation difficulty* and *sample distillation strength*. The former stems from a fact that different samples have varying difficulty in transferring the knowledge from teacher. The latter determines to what degree the KD loss should be more emphasized at sample level during student learning, and

it has negative correlation with sample distillation difficulty.

Regarding the first problem, we develop a sample distillation difficulty evaluation module which involves a temporal-dynamic interruption strategy. It adopts the dropout-shuffle technique that dynamically adjusts the interruption ratio of frames. Note that it can easily adapt to image analysis by conducting dropout-shuffle along the spatial dimension at pixel level. Essentially, it increases the learning difficulty (e.g., missing key frames, temporal misalignment among frames from different actions) of samples to enlarge the distillation loss between teacher and student. Usually, the larger the loss value, the lower the distillation difficulty is. This is because there exists large capacity difference to be diminished by distillation. Meanwhile, those samples selected multiples times are expected to have a low probability to be selected in future, which means these samples are difficult to transfer knowledge and should be less emphasized during distillation. The sample selection probability and the distillation loss value are coupled to determine the sample distillation difficulty.

Regarding the second problem, we design a sample adaptive distillation module, which dynamically calculates the sample distillation strength, according to the interruption ratio and the distillation difficulty. In particular, those highly difficult-to-transfer samples correspond to lower distillation strength and should learn more from ground-truth by vanilla loss. Meanwhile, difficult-to-transfer samples in early training may become easier in later training, since student learns more knowledge from teacher as training progresses. Hence, difficult-to-transfer samples may become less while easy-to-transfer ones become more during distillation. Moreover, we adopt the Determinantal Point Process (DPP) [13] sampling to select those diverse samples with low distillation difficulty to make the student more robust.

Overall, we propose a *Sample-level Adaptive Knowledge Distillation* (**SAKD**) framework for action recognition. It is an easy-to-use plug-in technique and can be easily embedded into existing KD methods. Extensive experiments on two video benchmarks, i.e., UCF101 [12] and Kinetics-400 [14], and one image benchmark CIFAR-100 [15], have validated both effectiveness and efficiency of our approach.

II. RELATED WORK

This section briefly reviews the related works in knowledge distillation and action recognition.

A. Knowledge Distillation

Knowledge Distillation (KD) reduces the inference model size without sacrificing the performance by transferring knowledge from teacher to student. Previous methods can be roughly divided into two groups, i.e., logit-based [3] [18] [19] ones and feature-based [21] [20] [4] ones. This allows student model to be deployed on resource-constrained devices efficiently.

Logit-based KD. Student imitates the output logits of teacher. Hinton *et al.* [3] developed the first KD method by transferring the output distributions via soft labels from teacher to student. There are several attempts to better handle the logits, e.g., Yang *et al.* [16] normalizes the non-target

logits from both teacher and student to use the soft labels for distillation; Wu *et al.* [17] sparsify the logits of teacher and store them in disk in advance to save the memory and computation cost, resulting in a tiny ViT (Vision Transformer) as student. To narrow down the gap of outputs, Li *et al.* [18] employ a dynamic and learnable temperature to control the task difficulty level during student learning but the model is difficult to converge; Sun *et al.* [19] set the temperature as the weighted standard deviation of logit and perform a plug-and-play Z-score preprocess of logit standardization, which facilitates only logit-based KD methods.

Feature-based KD. It makes the intermediate feature of student approach that of teacher [20]. For example, Han *et al.* [21] employ one layer feature of teacher as the target of student and compute the Euclidean distance of their features. Afterwards, some works attempt to transfer knowledge by activation map [22], feature distribution [23], and pairwise similarity [24], but these early works fail to consider the dynamics of the relation between teacher and student. To simplify the knowledge of teacher model, Wang *et al.* [25] propose a dual KD framework to reduce the side effects of teacher and utilize the optimal transport distance to measure the difference of feature maps between teacher and student. In addition, Zhu *et al.* [26] verify the existence of undistillable classes by illustrating their correlation with capacity gap between teacher and student, the similarity of which is used to reflect the distillability level of student model.

B. Action Recognition

Action recognition model assigns class label to the video involving one action, e.g., *running*. Early works adopt CNNs as the backbone, either using two-stream network [10], [27] to extract both RGB features along the spatial dimension and optical flow features along the temporal dimension, or using 3D CNNs [9], [28]–[30] to capture the spatiotemporal features. Taking advantage of two-stream network and 3D CNNs, Feichtenhofer *et al.* [11] combine them together and design the typical SlowFast network, where a slow pathway captures spatial semantics and a fast pathway captures motion at fine temporal resolution. Recent works [31]–[34] apply Transformer [2] to the action recognition task by using the model pre-trained on large-scale dataset. More recently, Wang *et al.* [8] present the generative feature KD framework to train student, which employs the conditional variational auto-encoder to extract attention-based features, but it brings about additional cost incurred by auto-encoder.

III. METHOD

This section introduces the framework of Sample-level Adaptive KD in action recognition task, and describes the principles of the sample distillation difficulty evaluation and adaptive distillation strategy.

A. Problem Definition and Preliminary

KD learns a small network (i.e., student $\mathcal{S}$) by transferring the knowledge from a large network (i.e., teacher $\mathcal{T}$), where
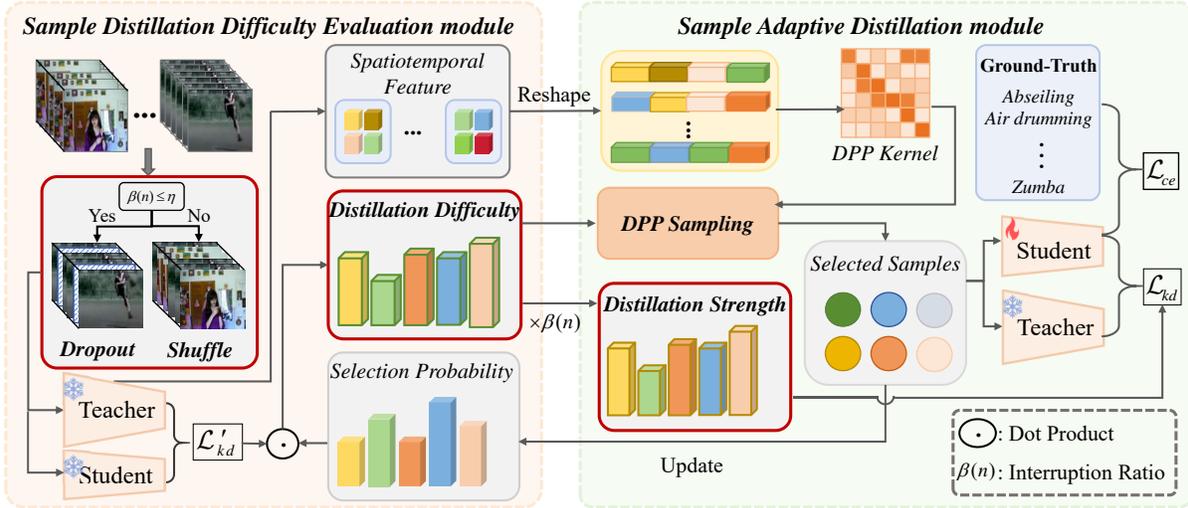
Fig. 2: Overall framework of the Sample-level Adaptive Distillation (SAKD) approach.

the latter is usually a pre-trained model with rich knowledge. Action recognition models are mostly very large and hardly to be deployed in resource-limited devices. Given a source video dataset, we divide it into training set $\mathcal{D}_{train} = \{\mathcal{V}_i, y_i\}_{i=1}^N$ and test set $\mathcal{D}_{test} = \{\mathcal{V}'_{i'}, y'_{i'}\}_{i'=1}^{N'}$, where $\mathcal{V}$ is video sample and $y \in \mathbb{R}$ is its action label from K categories. The random sampling is operated on each video to obtain T frames. For the i-th video clip, its RGB frames with height H and width W are stacked in a tensor $\mathbf{X}_i \in \mathbb{R}^{3 \times T \times H \times W}$. Thus, it yields a training set $\mathcal{X} = \{\mathbf{X}_i, y_i\}_{i=1}^N$ for student learning and a test set $\mathcal{X}' = \{\mathbf{X}'_{i'}, y'_{i'}\}_{i'=1}^{N'}$ for inference.

During training, the video clips in $\mathcal{X}$ are input to both teacher $\mathcal{T}$ and student $\mathcal{S}$, which adopt the commonly used backbone (e.g., ResNet [35]) for action recognition models such as Slowfast [11] and Temporal Pyramid Network [30] to derive the spatiotemporal feature maps, i.e., $\mathbf{F} \in \mathbb{R}^{N \times c \times h \times w}$, where $\{c, h, w\}$ denote channel, height, and width. The dimensions of feature maps $\mathbf{F}^{\mathcal{T}}, \mathbf{F}^{\mathcal{S}}$ of both teacher and student are forced to be the same by a mapping function $f(\cdot)$, e.g., 1×1 convolution or auto-encoder.

For feature-based KD, the distillation loss minimizes the gap of features between teacher and student to make student imitate teacher, i.e.,

$$\mathcal{L}_{kd}^{fea} = \frac{1}{N} \sum_{i=1}^N \|f(\mathbf{F}_i^{\mathcal{T}}) - f(\mathbf{F}_i^{\mathcal{S}})\|_2^2, \qquad (1)$$

where $\|\cdot\|_2$ denotes the ℓ_2-norm.

For logit-based KD, the distillation loss minimizes the discrepancy between the output probabilities of teacher and student, i.e.,

$$\mathcal{L}_{kd}^{logit} = \frac{1}{N} \sum_{i=1}^N \tau^2 KL(\sigma(q_i^{\mathcal{T}}/\tau), \sigma(q_i^{\mathcal{S}}/\tau)), \qquad (2)$$

where $KL(\cdot)$ is the Kullback-Leibler divergence, $\{q^{\mathcal{T}}, q^{\mathcal{S}}\}$ are the logits of teacher and student respectively, $\sigma(\cdot)$ is the softmax function, and τ (set to 4 [3]) is the temperature to scale the smoothness of two probability distributions.

The vanilla loss of action recognition model adopts the Cross-Entropy (CE) loss, i.e., $\mathcal{L}_{ce} = -\frac{1}{N} \sum_{i=1}^N y_i \log q_i^{\mathcal{S}}$.

B. Overall Framework

The overall framework of Sample-level Adaptive Knowledge Distillation (SAKD) for action recognition is illustrated in Fig. 2, which is composed of the Sample Distillation Difficulty (SDD) evaluation module and the Sample Adaptive Distillation (SAD) module. Unlike previous works using all to learn student model, SAKD selects only a small fraction of training samples for optimizing the model to save the cost. Essentially, we first evaluate the distillation difficulty of all training samples, and then select a small fraction of easy-to-transfer samples with diversity to train student model. Those selected samples form a selection subset $\mathcal{D}_{sel}$ dynamically in every epoch. The working mechanism of SAKD is described below. For example, we empirically found that the performance (e.g., accuracy) of our method using only 10% samples is comparable or even surpasses that of previous KD methods using all samples, and the training time is greatly reduced to only one-fourth of the previous ones.

First, the video samples in $\mathcal{D}_{train}$ are fed into the sample distillation difficulty evaluation module to compute the distillation difficulty score ζ for each sample. Note that the dropout-shuffle strategy is applied to video frames as a temporal interruption skill to enhance samples for increasing their learning difficulty, such that the distillation loss is enlarged for better evaluating distillation difficulty. These samples are fed into both teacher and student to derive the spatiotemporal features $\{\mathbf{F}^{\mathcal{T}}, \mathbf{F}^{\mathcal{S}}\}$ or output logits $\{q^{\mathcal{T}}, q^{\mathcal{S}}\}$ by forward propagation. They are used to calculate the distillation loss $\mathcal{L}_{kd}^{fea}$ or $\mathcal{L}_{kd}^{logit}$ using all samples. The loss value and the sample selection rate (i.e., historical selection times of samples) are used to evaluate the distillation difficulty at sample level. Moreover, we adopt the DPP [13] to evaluate the diversity among the selected samples, and choose those diverse samples with low distillation difficulty to participate in training student model. In addition, we estimate the sample distillation strength according

to the distillation difficulty and the interruption ratio per epoch, for better transferring knowledge from teacher to student by regarding it as an adaptive weighting ratio of the distillation loss and the vanilla loss.

C. Sample Distillation Difficulty Evaluation

Previous works treat all samples equally during distillation, while neglecting a fact that it may be trapped into an adverse situation, where both difficult-to-transfer and easy-to-transfer samples are emphasized equally which prevents more knowledge being transferred from teacher. We call this phenomenon "*distillation bottleneck*". The reason behind this is different samples contribute differently during distillation due to the large gap between teacher and student. It is more difficult for some samples to perform knowledge transfer from teacher while it is easier for other samples. If all samples are equally treated, the student performance may saturate after certain training epoch, since the majority of samples are likely to be regarded as easy-to-transfer ones by student during later training. This inspires us to evaluate the sample-level distillation difficulty by less emphasizing the difficult-to-transfer samples and more emphasizing the easy-to-transfer samples per epoch. Here, we discriminate difficult-to-transfer samples in two aspects: 1) if the sample selection probability in the next epoch is low, student may hardly learn more knowledge from teacher with more epochs; 2) if the distillation loss difference between neighbouring epochs is small but the loss value is large due to the teacher-student gap, or the distillation loss itself in current epoch is small, this suggests the knowledge is difficult to transfer for this sample.

Mathematically, we define the sample distillation difficulty ζ_i as:

$$\zeta_i = 1/(\mathbf{p}_i^{sel} \cdot \tilde{\mathcal{L}}_i^{kd}). \tag{3}$$

where $\mathbf{p}_i^{sel}$ denotes the sample selection probability, and $\tilde{\mathcal{L}}_i^{kd}$ is the distillation loss on the enhanced samples by the dynamic-temporal interruption strategy, which dynamically does the dropout-shuffle operations on video samples. This adds the learning difficulty and diversifies samples during distillation. Details are shown below.

Given the i-th video sample $\mathbf{X}_i \in \mathbb{R}^{3 \times T \times H \times W}$, we randomly dropout some frames or shuffle frames at some interruption rate $\beta(\cdot) \in (0, 1]$, which is a function of the current epoch number n, i.e.,

$$\beta(n) = 1 - (1 - n/N_{epoch})^\theta, \tag{4}$$

where N_{epoch} is the maximum epoch, and the power θ is set to 0.9 as indicated by poly learning rate [36]. Here we introduce a threshold $\eta \in (0, 1]$ to control whether adopting the dropout or shuffle operation. When the interruption rate is less than η, we randomly dropout a percentage $\beta(n)$ of the frames in a video along the temporal dimension; when it is larger than or equal to η, we randomly shuffle a percentage η of the frames in a batch, where there exist different actions. Motivated by the curriculum learning [18], the interruption rate function adopts the polynomial learning rate policy. In another word, the interruption on frames is weak at early training

and is gradually enhanced to increase the training difficulty of samples.

By the dynamic interruption operation, we derive the enhanced video sample $\tilde{\mathbf{X}}_i$, which is fed into teacher model and student model to compute the distillation loss $\tilde{\mathcal{L}}_i^{kd}$. This loss value acts as an important factor that influences the evaluation of sample distillation difficulty. Usually, the larger the loss value, the lower the distillation difficulty is. The sample with lower distillation difficulty has a high probability to be selected for learning the student model. Meanwhile, those samples ever selected multiple times in different epochs are expected to have a low probability to be selected again, such that more different samples have the opportunity of participating in the student learning. This will also improve the generalization ability of student model during inference. To achieve this, we define the sample selection probability $\mathbf{p}_i^{sel}$ as:

$$\mathbf{p}_i^{sel} = 1/(\omega_i + \epsilon), \tag{5}$$

where ω_i is the selection times of the i-th sample during training, and the constant $\epsilon > 0$ (set to 1 empirically) is a smoothing factor to avoid the zero value of denominator. The normalization is applied to the vector ω to ensure the probability will not be too small due to a large epoch. This selection probability acts as a correction factor that influences the evaluation of sample distillation difficulty.

At the early training stage, the difficult-to-transfer samples with too large loss values might be wrongly assigned with a low distillation difficulty value according to Eq. (3). Hence, we weight the loss value by a correction factor, i.e., the sample selection probability $\mathbf{p}_i^{sel}$. It makes sense because the correction factor will become small when the sample is selected multiple times, thus weakening the impact of difficult-to-transfer samples with large loss. As the training goes on, the number of easy-to-transfer samples will increase since the learning ability of student model is boosted gradually with more epochs. In this situation, both teacher and student may yield similar logits or intermediate features, leading to small distillation loss value. So, the samples with too similar outputs from teacher and student should not be selected in future, because it is hard for student to learn richer knowledge from teacher with these samples any more. These samples are also regarded as difficult-to-transfer ones.

D. Sample Adaptive Distillation Module

Traditional methods always fix the ratio of the vanilla loss and the KD loss. Worse still, this ratio is identified by grid searching from some range of hyper-parameters, which is inefficient at expensive computational cost. To overcome this drawback, we propose the sample adaptive distillation strategy, which dynamically estimates the distillation strength of each sample according to the sample distillation difficulty, i.e., the above ratio is adaptively adjusted to the current student model.

Formally, the distillation strength $\alpha_{i,n}$ of the i-the sample in the n-th epoch is defined as:

$$\alpha_{i,n} = \lambda \alpha_{i,n-1} + (1 - \lambda)\beta(n)/\zeta_i, \tag{6}$$

where the first term is the distillation strength of the previous epoch, the second term directly affects the distillation strength,

and the hyper-parameter $\lambda \in (0, 1)$ that balances the contribution of the two terms. The first term consider the previous distillation strength. The second term reveals two facts: 1) the distillation strength α gradually rises up as the sample interruption rate $\beta(n)$ in Eq. (4) grows monotonically with increasing epochs, and 2) the distillation strength α is inversely proportional to the distillation difficulty ζ, which makes sense since difficult-to-transfer samples should be less emphasized during distillation.

Another problem to be solved is to adaptively select those easy-to-transfer samples and diverse samples per epoch. For the former, we resort to the sample distillation difficulty, i.e., the samples with low distillation difficulty are treated as easy-to-transfer samples. For the latter, we adopt the Determinantal Point Process (DPP) [37] to evaluate the diversity of samples. Specifically, we define a DPP kernel matrix $\mathbf{\Lambda} \in \mathbb{R}^{N \times N}$ that measures the similarity of feature matrices. When the column vectors are more dissimilar, their angles are greater and thus the determinant of the kernel matrix, $\ln \det(\mathbf{\Lambda})$, is larger.

To obtain the above feature matrices, it goes through the following steps. First, we feed video sample $\tilde{\mathbf{X}}_i$ into teacher model to derive the feature map $\mathbf{F}_i \in \mathbb{R}^{c \times h \times w}$. Second, we fuse the feature maps of two neighbouring epochs to yield the fusion feature map $\tilde{\mathbf{F}}_i \in \mathbb{R}^{c \times h \times w}$, i.e.,

$$\tilde{\mathbf{F}}_{i,n} = \lambda \tilde{\mathbf{F}}_{i,n-1} + (1 - \lambda)\mathbf{F}_i, \qquad (7)$$

where $\lambda \in (0, 1)$ is the same as in Eq. (6) to balance the fusion ratio of the current epoch and the previous epoch. Here, we update the feature by considering that of previous one to avoid large fluctuations of feature during training, which helps to stabilize the subsequent DPP sampling process. The fusion feature map $\tilde{\mathbf{F}}_i$ is reshaped to $\hat{\mathbf{F}}_i \in \mathbb{R}^d$ by flattening the feature map to a vector at sample level, where $d = c \cdot h \cdot w$. Now we can compute the above DPP kernel matrix, i.e., $\mathbf{\Lambda} = \hat{\mathbf{F}} \cdot \hat{\mathbf{F}}^\top$, where $\hat{\mathbf{F}} \in \mathbb{R}^{N \times d}$ is the feature matrix of all samples.

To achieve the goal of selecting those samples with both low distillation difficulty and high diversity, we optimize the following objective function, i.e.,

$$\max_{\mathcal{A}} \gamma \sum_{j=1}^{N_{sel}} \frac{1}{\zeta_{proj(j)}} + (1 - \gamma) \ln \frac{\det(\mathbf{\Lambda}_{\mathcal{A}})}{\det(\mathbf{\Lambda} + \mathbf{I})}, \qquad (8)$$

where $\mathbf{I} \in \mathbb{R}^{N \times N}$ is an identity matrix, $\mathcal{A}$ is an index set of those selected samples, $N_{sel} = r \cdot N$ is the sample selection number, r is a selection ratio, and hyper-parameter $\gamma \in [0, 1]$ is used to tradeoff the former distillation difficulty term and the latter diversity term. Here, there exists one-to-one projection $proj(\cdot)$ between the new index $j \in \{1, 2, \ldots, N_{sel}\}$ in the selected sample set $\mathcal{A}$ and the index i of the source sample, and vice versa, i.e., $i = proj(j)$. The matrix $\mathbf{\Lambda}_{\mathcal{A}}$ denotes the kernel matrix derived from the features of those selected samples, which forms the selection subset $\mathcal{D}_{sel} \subset \mathcal{D}_{train}$. We adopt the greedy algorithm [13] to optimize the above objective function by Cholesky decomposition to reduce the computational complexity.

TABLE I: Statistics of data. "K"/"N" is class/sample number.

Dataset	Training		Validation		Test	
	K	N	K	N	K	N
UCF101 [12]	101	9,537	-	-	101	3,783
Kinetics-400 [14]	400	234,619	400	19,761	-	-
CIFAR-100 [15]	100	50,000	-	-	100	10,000

E. Loss Function

For action recognition task, the total loss consists of the vanilla classification loss $\mathcal{L}_{ce}$ and the KD loss $\mathcal{L}_{kd}$, and it is calculated on those selected video clips per epoch, i.e.,

$$\mathcal{L}_{total} = \frac{1}{N_{sel}} \sum_{j=1}^{N_{sel}} (1 - \alpha_{proj(j)})\mathcal{L}_j^{ce} + \alpha_{proj(j)}\mathcal{L}_j^{kd}. \qquad (9)$$

where $\alpha_{proj(j)} \in (0, 1)$ is a dynamic distillation strength of the j-th selected video clip per epoch, and $\mathcal{L}_j^{kd}$ takes the form of Eq. (1) or (2). During the early training, difficult-to-transfer samples are expected to be less emphasized during distillation and student should learn knowledge more from ground-truth labels than from teacher, e.g., sample distillation strength α is less than 0.5. As the training progresses, the learning ability of student model is boosted gradually. Therefore, those originally difficult-to-transfer samples may become easier to transfer during the later training, and simultaneously the sample distillation strength gets larger by dominating the distillation process.

IV. EXPERIMENT

All experiments were performed on a server equipped with four 11G GeForce 2080Ti graphics cards. The codes are compiled with PyTorch 1.7, Python 3.8, and CUDA 10.1.

A. Datasets

We conduct experiments on two video benchmarks including UCF101 [12][1] and Kinetics-400 [14][2], and one image benchmark CIFAR-100 [15][3]. Statistics are in Table I.

UCF101[4] [12] consists of daily-life action videos collected from YouTube, covering 101 different categories with a total of 13,320 video clips, whose total duration is approximately 27 hours. The videos are organized into 25 groups by category, with each group containing 4 to 7 videos. The video resolution is 320×240, and the frame rate is 25 fps. We use the official splits with all images uniformly cropped to a size of 224×224.

Kinetics-400[5] [14] was originally released by DeepMind and contains video clips collected from YouTube, covering 400 different action categories. Each category has at least 400 videos, with each video clip lasting approximately 10 seconds. It includes 234,619 samples for the training set and 19,761 videos for the validation set, with all input images uniformly cropped to a size of 256×256.

---

[1] https://www.crcv.ucf.edu/data/UCF101.php
[2] https://deepmind.com/research/open-source/kinetics
[3] http://www.cs.toronto.edu/ kriz/cifar.html
[4] https://www.crcv.ucf.edu/data/UCF101.php
[5] https://deepmind.com/research/open-source/kinetics

TABLE II: Batch size and training epoch.

| Model | Batch Size | | | | N_{epoch} | |
| | Our method | | Student training | | | |
	UCF101	Kinetics	UCF101	Kinetics	UCF101	Kinetics
SlowFast [11]	96	128	32	48	50	200
TPN [30]	24	-	4	-	50	-
Video ST [33]	24	48	4	8	100	200

TABLE III: Performance on UCF101 [12] in terms of Top-1/5 Accuracy (%) / training time (min). "r" is sample selection ratio.

| Method | Type | r | SlowFast [11] | | | VideoST [33] | | | TPN [30] | | |
			Top1↑	Top5↑	min↓	Top1↑	Top5↑	min↓	Top1↑	Top5↑	min↓
-	Teac.	100%	93.61	98.64	10.2	92.18	98.01	22.2	90.21	96.51	32.8
	Stud.	100%	85.15	95.80	6.2	86.17	97.03	18.1	82.23	93.24	20.7
KD [3] arXiv'15	Vani.	100%	87.92	97.56	6.4	89.38	98.25	20.0	86.02	94.60	23.4
	Ours	10%	**90.73**	**99.20**	4.6	**89.43**	**99.04**	11.1	**87.15**	96.61	11.3
	Ours*	10%	89.88	98.80	**1.4**	87.70	98.72	**3.9**	85.55	**97.01**	**4.9**
DKD [39] CVPR'22	Vani.	100%	90.50	98.46	6.8	91.60	98.43	19.9	85.40	94.05	23.5
	Ours	10%	**91.07**	**99.31**	4.6	**91.78**	**99.06**	12.2	**87.92**	**97.59**	13.0
	Ours*	10%	89.93	98.83	**1.5**	91.44	98.72	**4.0**	86.20	96.88	**5.0**
CTKD [18] AAAI'23	Vani.	100%	87.12	96.45	6.9	89.61	97.93	20.1	86.12	94.57	24.2
	Ours	10%	**90.67**	98.88	4.8	**89.81**	**99.08**	12.1	**86.65**	96.84	11.5
	Ours*	10%	89.61	**98.91**	**1.4**	87.83	97.82	**4.0**	85.96	**97.10**	**4.9**
GKD [8] AAAI'24	Vani.	100%	86.83	96.03	7.0	88.24	97.97	20.8	86.37	94.63	24.9
	Ours	10%	**89.92**	**98.51**	5.3	**88.42**	**98.16**	13.8	**88.72**	**97.28**	13.9
	Ours*	10%	87.92	98.34	**1.8**	87.23	98.03	**4.6**	87.64	97.04	**4.8**
CrossKD [40] CVPR'24	Vani.	100%	90.79	99.10	6.5	87.63	98.06	19.8	88.05	97.51	23.2
	Ours	10%	**90.82**	98.91	4.9	**87.98**	**98.38**	11.6	**88.55**	**97.96**	12.0
	Ours*	10%	90.04	**99.23**	**1.5**	86.94	97.16	**3.8**	87.69	97.29	**4.9**
DualKD [25] TIP'24	Vani.	100%	86.65	96.82	7.0	89.27	98.06	20.5	86.17	94.75	26.2
	Ours	10%	**90.32**	**98.79**	5.0	**89.58**	**99.87**	13.0	**86.67**	**97.21**	11.4
	Ours*	10%	89.64	98.72	**1.6**	87.95	97.72	**4.3**	85.86	96.76	**5.0**

CIFAR-100[6] [15] contains 100 categories, each of which has 600 RGB images of size 32×32. The training set has 50,000 images and the test set has 10,000 images. All images are uniformly rescaled to 32×32.

Following previous works [5], [18], [38], we adopt the commonly used Top-1 accuracy and Top-5 accuracy as the evaluation metrics. Also we report the elapsed time of each epoch to show the training efficiency. Top-1 or Top-5 accuracy evaluates those samples whose ground-truth class takes up the top or one of the five leading positions of the candidate class set.

B. Experimental Setup

For UCF101 [12], the action recognition model parameters are initialized with the model pre-trained on Kinetics-400 [14]. For UCF101 and Kinetics-400, the number of frames T in a clip after sampling is 32 and 16, respectively. The Stochastic Gradient Descent optimizer is used with a momentum of 0.9. The initial learning rate for UCF101 and Kinetics-400 is set to 1e-2, with a weight decay factor of 1e-4 after each epoch. For Kinetics-400, following [33], both the tiny version (SwinT) and the small version (SwinS) of Video Swin Transformer (Video ST) [33] are pre-trained on ImageNet, while models are randomly initialized. For CIFAR-100, following [4], it has an initial learning rate of 0.1 and a decay factor of 0.1 applied at epochs 150, 180, and 210. During student training, the parameters of teacher are frozen and do not participate in back-propagation. For CIFAR-100 [15], the batch size of our method using frozen model is 1024, the batch size of student training is 256, and the training epoch is set to 240. The batch size and training epoch settings of two video datasets are summarized in Table II.

To reduce the training cost, we select a small fraction of samples with low distillation difficulty and high diversity every epoch or every five epochs. For UCF101 and Kinetics-400, the sample selection ratio r is 0.1, i.e., 10% of all training samples are selected every time, and this rate is set to 0.5 for CIFAR-100 since images have much less redundancy than video. The hyper-parameters are set as follows: λ of the distillation strength in Eq. (6) and the feature fusion in Eq. (7) is set to 0.1, γ of the sample selection using DPP in Eq. (8) is set to 0.5, and the interruption (dropout-shuffle) ratio threshold η is set to 0.5.

Inference. Given a test video or image, we do the normalization before feeding them into the student model to output the estimated action class or image label.

[6]http://www.cs.toronto.edu/ kriz/cifar.html

C. Compared Methods

We compare our SAKD method with two groups of State-Of-The-Art (SOTA) KD methods: 1) *logit-based* ones include KD (vanilla Knowledge Distillation) [3], DKD (Decoupled KD) [39], and CTKD (Curriculum Temperature KD) [18]; 2) *feature-based* ones include CrossKD (Cross-head KD) [40], DualKD [25], and GKD (Generative model based KD) [8]. We use the source codes publicly available from the original papers, and we try the best to implement DualKD [25] by ourself since its code is unavailable. Our code is available in the attached file.

To verify the generalization ability of our SAKD method on both video and image samples, we examine the compared KD methods on three typical action recognition models, including SlowFast [11], TPN (Temporal Pyramid Network) [30], and VideoST (Video SwinTransformer) [33], as well as two typical image classification models (settings follows [4]), including ResNet (Residual Neural Network) [35] and WRN (Wide ResNet) [41]. For SlowFast, we sample 16 or 8 or 4 frames in a clip when the step is set to 8 or 16, termed SF16x8 or SF8x8 or SF4x16. For TPN teacher or student, we sample 32 or 8 frames along the temporal dimension and scale up frames along the spatial dimension by a factor of 2 or 8, termed TPN-f32s2 or TPN-f8s8. For VideoST, SwinS is teacher whose total Transformer layer at Stage 3 is 18, while SwinT is student whose that layer number is 6. For SlowFast and TPN, teacher and student adopt ResNet101 and ResNet50 as the backbone respectively.

D. Quantitative Results

We apply all compared KD methods on action recognition models, whose results are reported in Table III (UCF101 [12]) and Table IV (Kinetics-400 [14]). The results on image models are reported in Table V (CIFAR-100 [15]). The best records are highlighted in boldface, and the second best ones

TABLE IV: Performance on Kinetics-400 [14] in terms of Top-1/5 Accuracy (%) / training time (hour). "r" is sample selection ratio.

Method	Type	r	SlowFast [11] Top1↑	Top5↑	hr↓	VideoST [33] Top1↑	Top5↑	hr↓
-	Teacher	100%	62.77	84.55	3.3	74.07	91.05	3.7
	Student	100%	51.08	76.05	2.6	70.81	89.61	2.9
KD [3] arXiv'15	Vanilla	100%	54.81	78.95	2.8	71.98	90.29	3.1
	Ours	10%	55.26	79.85	2.1	73.01	91.00	2.4
	Ours*	10%	54.37	78.23	0.8	72.69	90.64	1.1
DKD [39] CVPR'22	Vanilla	100%	56.17	80.20	2.9	73.95	92.10	3.1
	Ours	10%	56.74	81.72	2.2	74.07	92.13	2.3
	Ours*	10%	56.35	80.88	0.8	73.92	91.47	1.0
CTKD [18] AAAI'23	Vanilla	100%	55.21	79.95	3.0	72.23	90.72	3.2
	Ours	10%	55.99	81.02	2.3	73.12	91.13	2.3
	Ours*	10%	54.82	80.92	0.8	72.83	90.87	1.1
CrossKD [40] CVPR'24	Vanilla	100%	58.07	79.26	2.8	74.71	91.86	3.0
	Ours	10%	62.18	81.02	2.2	75.07	92.06	2.2
	Ours*	10%	60.28	80.54	0.8	74.85	91.72	1.0
DualKD [25] TIP'24	Vanilla	100%	56.25	79.85	3.1	74.09	91.29	3.3
	Ours	10%	57.02	81.95	2.4	74.87	91.82	2.5
	Ours*	10%	56.37	81.83	0.9	74.28	91.39	1.2

TABLE V: Performance on CIFAR-100 [15] in terms of Top-1/5 Accuracy (%) / training time (sec). "r" is sample selection ratio.

Method	Type	r	WRN [41] WRN40-2 WRN16-2	WRN40-2 WRN40-1	ResNet [35] ResNet56 ResNet20	ResNet110 ResNet32	ResNet32×4 ResNet8×4
-	Teac.	100%	75.61 (45.2s)	75.61 (54.2s)	72.34 (50.3s)	74.31 (68.2s)	79.42 (39.8s)
	Stud.	100%	73.26 (38.2s)	71.09 (46.2s)	69.06 (43.2s)	70.87 (60.3s)	71.50 (33.8s)
KD [3] arXiv'15	Vani.	100%	74.34 (40.2s)	73.02 (48.8s)	70.81 (46.2s)	73.18 (64.6s)	72.23 (36.1s)
	Ours	50%	74.44 (33.6s)	72.86 (39.4s)	70.62 (40.6s)	73.36 (50.8s)	72.85 (26.3s)
	Ours*	50%	74.23 (22.8s)	72.69 (27.4s)	70.55 (26.6s)	73.12 (36.0s)	72.05 (21.6s)
DKD [39] CVPR'22	Vani.	100%	76.17 (40.2s)	74.91 (50.4s)	71.86 (45.3s)	73.41 (64.2s)	75.92 (36.2s)
	Ours	50%	75.95 (33.6s)	74.88 (38.2s)	71.35 (35.4s)	73.87 (48.6s)	76.02 (27.0s)
	Ours*	50%	75.62 (22.4s)	74.54 (27.8s)	70.92 (25.9s)	72.92 (35.4s)	75.57 (21.5s)
CTKD [18] AAAI'23	Vani.	100%	74.58 (41.2s)	72.83 (50.2s)	71.08 (46.9s)	72.83 (65.2s)	72.52 (36.2s)
	Ours	50%	74.98 (35.1s)	72.93 (33.8s)	70.92 (41.2s)	72.91 (53.6s)	72.63 (26.8s)
	Ours*	50%	74.83 (23.5s)	72.85 (28.0s)	70.64 (27.0s)	72.62 (36.8s)	72.20 (21.4s)
CrossKD [40] CVPR'24	Vani.	100%	73.41 (40.9s)	71.33 (49.5s)	69.95 (46.2s)	71.58 (64.1s)	71.90 (36.7s)
	Ours	50%	74.01 (36.2s)	71.42 (41.0s)	69.73 (41.6s)	71.42 (48.5s)	71.98 (27.4s)
	Ours*	50%	73.87 (23.6s)	71.11 (28.0s)	69.21 (26.8s)	71.05 (35.3s)	71.48 (21.3s)
DualKD [25] TIP'24	Vani.	100%	74.24 (41.3s)	73.86 (50.3s)	70.77 (46.8s)	72.87 (65.3s)	72.50 (36.9s)
	Ours	50%	74.01 (34.4s)	73.53 (40.9s)	70.82 (40.4s)	72.98 (50.6s)	73.52 (26.3s)
	Ours*	50%	73.78 (23.4s)	73.29 (28.3s)	70.12 (26.8s)	72.55 (36.2s)	72.18 (22.3s)

TABLE VI: Ablation studies on components in terms of Top-1/5 Accuracy (%). "Diff" denotes distillation difficulty evaluation module, "Ada." denotes adaptive distillation module.

Method	Dif.	Ada.	UCF101 SlowFast Top1↑	Top5↑	VideoST Top1↑	Top5↑	Kinetics-400 SlowFast Top1↑	Top5↑	VideoST Top1↑	Top5↑	CIFAR100 ResNet Top1↑	Top5↑	WRN Top1↑	Top5↑
KD [3]			78.91	94.33	84.77	96.07	50.23	74.23	66.24	87.23	68.01	90.91	72.01	92.24
	✓		81.74	96.10	86.50	97.72	53.01	78.27	71.24	90.23	69.25	91.80	72.95	92.79
		✓	88.95	98.33	86.07	96.64	51.98	76.44	70.92	90.05	68.75	91.63	73.18	92.84
	✓	✓	89.88	98.80	87.70	98.72	54.37	78.23	72.69	90.64	70.55	92.85	74.23	92.87
CrossKD [40]			86.39	97.67	83.65	96.88	56.21	76.28	68.34	88.02	67.08	89.92	70.23	91.72
	✓		88.44	99.18	85.19	97.09	57.72	77.85	72.83	91.23	67.97	90.85	71.75	92.36
		✓	89.06	98.59	84.73	96.96	57.41	79.81	70.88	90.03	68.47	91.19	72.09	92.73
	✓	✓	90.04	99.23	86.94	97.16	60.28	80.54	74.85	92.06	69.21	92.93	73.87	93.87

are underlined. Here, "Vanilla" (Vani.) denotes existing KD method, "Ours" or "Ours*" denotes our method selecting samples every epoch or every five epochs.

From Table III and Table IV, we observe that our SAKD method as a plug-in technique consistently improves the performance of three typical action recognition models including SlowFast [11], VideoST [33], and TPN [30] across six SOTA KD methods on UCF101 and Kinetics-400. Compared to the baseline (row 1) without knowledge distillation, both vanilla KD methods and ours perform better, which supports the claim that student indeed learns knowledge from teacher by KD. Meanwhile, vanilla KD methods consumes more time than the baseline due to the additional overheads incurred by the distillation module. On the contrary, our method greatly reduces training time, e.g., from 4.8 min to 1.5 min with SlowFast on CTKD method at a faster training speed by over three times. Moreover, our SAKD method outperforms the vanilla KD ones by using only 10% of training samples. This is because ours takes into account the different distillation difficulty of samples and adaptively adjusts the distillation strength for each sample as training progresses, which allows the student to learn more knowledge from teacher. Especially, those difficult-to-transfer samples receive more emphasis during student training, further improving the performance.

In addition, similar performance improvements can be found on CIFAR-100 as shown in Table V, which reports the results of six different backbone pairs in two types including WRN [41] and ResNet [35]. While KD methods bring about improvements on the baseline, the gain is smaller using 50% of training samples compared with that on video. This might because that CIFAR-100 is an image database with few redundancy and it may lose some representative images when applying the dropout strategy. This is a bit different from that of video with much redundancy.

E. Ablation Studies

To probe into the details of our method, some ablation studies on the individual components, the interruption threshold η, the distillation strength λ, and the sample selection ratio r were conducted on UCF101 and CIFAR-100. Here, we use ResNet56/ResNet20 and WRN40-2/WRN16-2 as teacher/student for CIFAR-100. We do ablations using KD [3] by selecting samples every five epochs, and hyper-parameters keep still as during training unless specified. Please refer to supplementary for more results.

Individual component. We examine the sample distillation difficulty (SDD) evaluation module and the sample adaptive distillation (SAD) module, whose results are shown in Table VI, where the baseline without "✓" is the vanilla KD method with DPP sampling. As shown in the top, using SDD module (row 2) improves the performance by about 2.8% in terms of Top-1 accuracy with SlowFast model on UCF101, and that improvement with VideoST is larger on Kinetics-400, i.e., 5.0%; the performance gains are significant (i.e., 10%) by adding SAD module (row 3) with SlowFast on UCF101; coupling the two module (row 4) achieves the best performance. Similar observations are also found by using CrossKD in the bottom group. This demonstrates considering the sample distillation difficulty and adaptively updating the sample distillation strength are beneficial for transferring knowledge from teacher, regardless of logit-based or feature-

TABLE VII: Ablation studies on the interruption threshold η.

η	UCF101 [12]				CIFAR100 [15]			
	Top1↑		Top5↑		Top1↑		Top5↑	
	SlowFast	VideoST	SlowFast	VideoST	ResNet	WRN	ResNet	WRN
0.0	88.34	86.82	98.43	96.59	69.01	73.19	91.52	93.27
0.1	88.42	87.01	98.67	97.33	69.55	73.08	91.82	93.34
0.3	89.27	87.52	98.64	98.05	70.02	73.64	92.37	93.76
0.5	89.88	87.70	98.80	98.72	70.55	73.87	92.85	93.79
0.7	89.53	87.30	98.07	98.84	70.46	73.82	92.93	93.86
0.9	89.14	87.03	97.99	97.93	69.32	73.65	92.36	93.54
1.0	88.78	86.87	96.87	97.70	69.43	73.76	92.18	93.20

TABLE VIII: Ablation on distillation strength hyper-parameter λ.

λ	UCF101 [12]				CIFAR100 [15]			
	Top1↑		Top5↑		Top1↑		Top5↑	
	SlowFast	VideoST	SlowFast	VideoST	ResNet	WRN	ResNet	WRN
0.0	89.69	87.24	98.64	98.03	69.81	73.20	92.01	93.39
0.1	89.96	87.79	98.80	98.72	70.55	73.87	92.85	93.97
0.3	89.91	87.70	98.67	98.35	70.23	73.56	92.63	93.41
0.5	89.87	87.65	98.52	98.49	70.02	73.69	92.45	93.32
0.7	89.77	87.28	98.53	98.02	69.87	73.48	92.32	93.67
0.9	89.43	86.72	98.42	97.98	68.85	72.87	92.18	93.25
1.0	89.38	86.01	98.36	97.43	69.16	73.47	92.02	93.59

based methods.

Interruption threshold η. We vary η from 0 to 1 with seven grids, and show the results in Table VII. As the threshold increases from 0, the performance is progressively enhanced and mostly achieves the best around 0.5, but it tends to degrade after the peak point. This indicates that the threshold should be neither too large nor too small, since it decides whether doing random dropout ($< \eta$) or random shuffle on video frames ($\geq \eta$).

Distillation strength λ. We vary λ from 0 to 1 with seven grids, and show the results in Table VIII. From the table, the performance achieves the best when λ takes 0.1, i.e., the previous distillation strength contributes 10%, while the current one contributes 90% to the model. This suggests the history distillation strength cannot be neglected and the current one dominates the adaptive distillation process.

Sample selection ratio r. We vary r from 1% to 100% and depict the Top-1 accuracy and the training time in Fig. 3. As drawn by the curves, the accuracy rapidly approaches (VideoST) or even surpasses (SlowFast) that of baseline (vanilla KD) by only 10% of training samples on two video datasets during student learning. The elapsed training time is very low compared to that of baseline using all training samples. Note that CIFAR-100 is an image dataset which is much easier for classification, which makes knowledge distillation more difficult, so it desires almost 50% (ResNet) or 30% (WRN) of training samples to achieve that of baseline.

F. Qualitative Results

To intuitively show the advantage of our method, we randomly choose several videos from UCF101 [12] and Kinetics-400 [14] and visualize the feature maps in Fig. 4. These features are from the 2nd block in the slow path of SlowFast [11]

(a) UCF101 [12]-SlowFast [11]   (b) UCF101 [12]-VideoST [33]

(c) CIFAR-100 [15]-ResNet [35]   (d) CIFAR-100 [15]-WRN [41]
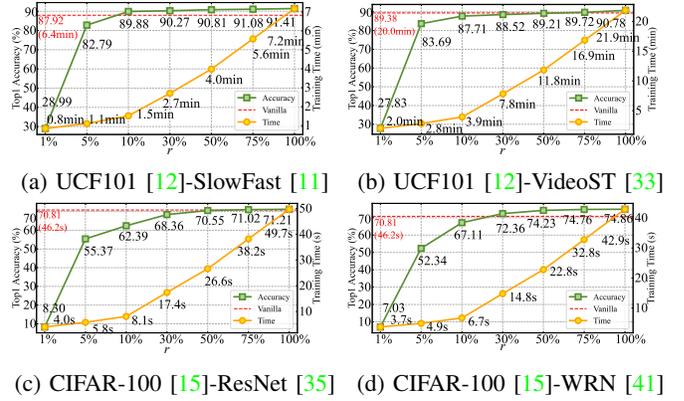
Fig. 3: Performance vs. sample selection ratio r.
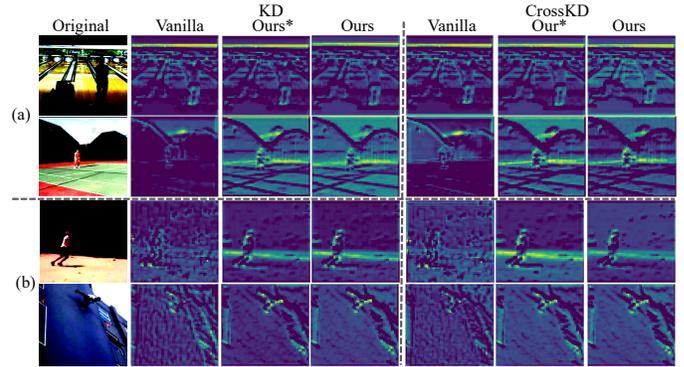
Fig. 4: Feature map visualization of our method on SlowFast [11]. (a) UCF101 [12]; (b) Kinetics-400 [14].

model by applying our method to vanilla KD [3] and vanilla CrossKD [40]. As shown in Fig. 4(a), vanilla KD or CrossKD fails to clearly show the contour in dark area (bottom in row 1 and top in row 2), which may cause a wrong recognition (e.g., "bowling" → "standing"). As shown in Fig. 4(b), vanilla KD or CrossKD generates the feature maps with many noisy pixels which adds the discriminating difficulty of actions. Our method not only highlights the contours in dark area but also greatly reduces the number of noisy pixels, regardless of selecting samples every epoch (Ours) or every five epochs (Ours*). This validates the robustness of our method in adverse condition.

V. Conclusion

This work explores the knowledge distillation problem at sample level for action recognition, by developing an efficient adaptive distillation framework. It is inspired by a fact that the transfer difficulty varies across different samples during distillation. In particular, we first evaluate sample distillation difficulty by considering the sample selection probability and the distillation loss of the interrupted samples, after which we calculate the sample distillation strength based on the interruption rate and the distillation difficulty. Meanwhile, we select only a small fraction of samples with both low distillation difficulty and high diversity to train student model. Empirical studies on several benchmarks validate that our

method achieves comparable performance or even surpasses prevailing SOTA methods at much lower computational cost.

REFERENCES

[1] A. Krizhevsky, I. Sutskever, and G. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1106–1114, 2012.

[2] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017, pp. 5998–6008.

[3] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint*, vol. arXiv:1503.02531, 2015.

[4] C. Pham, Van-Anh, N. Trung, L. Dinh, and P. G. Carneiro, "Frequency attention for knowledge distillation," *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 2277–2286, 2024.

[5] M. Zong, Z. Qiu, X. Ma, K. Yang, C. Liu, J. Hou, S. Yi, and W. Ouyang, "Better teacher better student: dynamic prior knowledge for knowledge distillation," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.

[6] N. Crasto, P. Weinzaepfel, K. Alahari, and C. Schmid, "MARS: motion-augmented rgb stream for action recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 7882–7891.

[7] J. C. Stroud, D. A. Ross, C. Sun, J. Deng, and R. Sukthankar, "D3D: distilled 3d networks for video action recognition," in *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2020, pp. 614–623.

[8] G. Wang, P. Zhao, Y. Shi, C. Zhao, and S. Yang, "Generative model-based feature knowledge distillation for action recognition," in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2024, pp. 15 474–15 482.

[9] S. Ji, W. Xu, M. Yang, and K. Yu, "3d convolutional neural networks for human action recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 35, pp. 221–231, 2012.

[10] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 27, 2014.

[11] C. Feichtenhofer, H. Fan, J. Malik, and K. He, "Slowfast networks for video recognition," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019, pp. 6202–6211.

[12] K. Soomro, A. R. Zamir, and M. Shah, "Ucf101: A dataset of 101 human actions classes from videos in the wild," *arXiv preprint arXiv:1212.0402*, 2012.

[13] L. Chen, G. Zhang, and E. Zhou, "Fast greedy map inference for determinantal point process to improve recommendation diversity," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018, pp. 5627–5638.

[14] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijaya-narasimhan, F. Viola, T. Green, T. Back, P. Natsev, M. Suleyman, and A. Zisserman, "The kinetics human action video dataset," *arXiv preprint arXiv:1705.06950*, 2017.

[15] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," Canadian Institute for Advanced Research, Tech. Rep., 2009.

[16] Z. Yang, A. Zeng, Z. Li, T. Zhang, C. Yuan, and Y. Li, "From knowledge distillation to self-knowledge distillation: A unified approach with normalized loss and customized soft labels," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2023, pp. 17 185–17 194.

[17] K. Wu, J. Zhang, H. Peng, M. Liu, B. Xiao, J. Fu, and L. Yuan, "Tinyvit: Fast pretraining distillation for small vision transformers," in *Proceedings of the European Conference on Computer Vision (ECCV)*, vol. 13681, 2022, pp. 68–85.

[18] Z. Li, X. Li, L. Yang, B. Zhao, R. Song, L. Luo, J. Li, and J. Yang, "Curriculum temperature for knowledge distillation," in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2023, pp. 1504–1512.

[19] S. Sun, W. Ren, J. Li, R. Wang, and X. Cao, "Logit standardization in knowledge distillation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 15 731–15 740.

[20] X. Liu, L. Li, C. Li, and A. Yao, "NORM: knowledge distillation via n-to-one representation matching," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.

[21] S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2015, pp. 1135–1143.

[22] S. Zagoruyko and N. Komodakis, "Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.

[23] N. Passalis and A. Tefas, "Learning deep representations with probabilistic knowledge transfer," in *Proceedings of the European Conference on Computer Vision (ECCV)*, vol. 11215, 2018, pp. 283–299.

[24] F. Tung and G. Mori, "Similarity-preserving knowledge distillation," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019, pp. 1365–1374.

[25] R. Wang, Y. Hao, L. Hu, X. Li, M. Chen, Y. Miao, and I. Humar, "Efficient crowd counting via dual knowledge distillation," *IEEE Transactions on Image Processing (TIP)*, vol. 33, pp. 569–583, 2024.

[26] Y. Zhu, N. Liu, Z. Xu, X. Liu, W. Meng, L. Wang, Z. Ou, and J. Tang, "Teach less, learn more: On the undistillable classes in knowledge distillation," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

[27] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. V. Gool, "Temporal segment networks: Towards good practices for deep action recognition," in *Proceedings of the European Conference on Computer Vision (ECCV)*, vol. 9912, 2016, pp. 20–36.

[28] J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 4724–4733.

[29] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3d convolutional networks," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 4489–4497.

[30] C. Yang, Y. Xu, J. Shi, B. Dai, and B. Zhou, "Temporal pyramid network for action recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 591–600.

[31] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.

[32] K. Li, Y. Wang, Y. Li, Y. Wang, Y. He, L. Wang, and Y. Qiao, "Unmasked teacher: Towards training-efficient video foundation models," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2023, pp. 19 948–19 960.

[33] Z. Liu, J. Ning, Y. Cao, Y. Wei, Z. Zhang, S. Lin, and H. Hu, "Video swin transformer," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 3202–3211.

[34] W. Wu, Y. Song, Z. Sun, J. Wang, C. Xu, and W. Ouyang, "What can simple arithmetic operations do for temporal modeling?" in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2023, pp. 13 666–13 676.

[35] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.

[36] Y. Wu and L. Liu, "Selecting and composing learning rate policies for deep neural networks," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 14, pp. Article No. 22, 1–25, 2023.

[37] A. Kulesza and B. Taskar, "Determinantal point processes for machine learning," *Foundations and Trends in Machine Learning*, vol. 5, pp. 123–286, 2012.

[38] Y. Xie, H. Zhang, X. Xu, J. Zhu, and S. He, "Towards a smaller student: capacity dynamic distillation for efficient image retrieval," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 16 006–16 015.

[39] B. Zhao, Q. Cui, R. Song, Y. Qiu, and J. Liang, "Decoupled knowledge distillation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 11 953–11 962.

[40] J. Wang, Y. Chen, Z. Zheng, X. Li, M.-M. Cheng, and Q. Hou, "Crosskd: Cross-head knowledge distillation for object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 16 520–16 530.

[41] S. Zagoruyko and N. Komodakis, "Wide residual networks," in *Proceedings of the British Machine Vision Conference (BMVC)*, 2016, pp. 19–22.