

# Curriculum Temperature for Knowledge Distillation

Zheng Li<sup>1</sup>, Xiang Li<sup>1\*</sup>, Lingfeng Yang<sup>2</sup>, Borui Zhao<sup>3</sup>,  
Renjie Song<sup>3</sup>, Lei Luo<sup>2</sup>, Jun Li<sup>2</sup>, Jian Yang<sup>1\*</sup>

<sup>1</sup> Nankai University

<sup>2</sup> Nanjing University of Science and Technology

<sup>3</sup> Megvii Technology

zhengli97@mail.nankai.edu.cn, {xiang.li.implus, csjyang}@nankai.edu.cn,

zhaoborui.gm@gmail.com, songrenjie@megvii.com, {yanglfnjust, cslluo, junli}@njust.edu.cn

## Abstract

Most existing distillation methods ignore the flexible role of the temperature in the loss function and fix it as a hyper-parameter that can be decided by an inefficient grid search. In general, the temperature controls the discrepancy between two distributions and can faithfully determine the difficulty level of the distillation task. Keeping a constant temperature, i.e., *a fixed level of task difficulty*, is usually sub-optimal for a growing student during its progressive learning stages. In this paper, we propose a simple curriculum-based technique, termed Curriculum Temperature for Knowledge Distillation (CTKD), which controls the task difficulty level during the student’s learning career through a dynamic and learnable temperature. Specifically, following an easy-to-hard curriculum, we gradually increase the distillation loss w.r.t. the temperature, leading to increased distillation difficulty in an adversarial manner. As an easy-to-use plug-in technique, CTKD can be seamlessly integrated into existing knowledge distillation frameworks and brings general improvements at a negligible additional computation cost. Extensive experiments on CIFAR-100, ImageNet-2012, and MS-COCO demonstrate the effectiveness of our method.

## Introduction

Knowledge distillation (Hinton, Vinyals, and Dean 2015) (KD) has received increasing attention from both academic and industrial researchers in recent years. It aims at learning a comparable and lightweight student by transferring the knowledge from a pretrained heavy teacher. The traditional process is implemented by minimizing the KL-divergence loss between two predictions obtained from the teacher/student model with a fixed temperature in the softmax layer. As depicted in (Hinton, Vinyals, and Dean 2015; Liu et al. 2022; Chandrasegaran et al. 2022), the temperature controls the smoothness of distribution and can faithfully determine the difficulty level of the loss minimization process. Most existing works (Tung and Mori 2019; Chen et al. 2020; Ji et al. 2021) ignore the flexible role of the temperature and empirically set it to a fixed value (e.g., 4). Differently, MKD (Liu et al. 2022) proposes to

learn the suitable temperature via meta-learning. However, it has certain limitations that require an additional validation set to train the temperature module, which complicates the training process. Besides, it mainly focuses on the strong data augmentation condition, neglecting that most existing KD methods work under normal augmentation. Directly combining MKD with existing distillation methods under strong augmentation may cause severe performance degradation (Das et al. 2020).

In human education, teachers always train students with simple curricula, which start from easier knowledge and gradually present more abstract and complex concepts when students grow up. This curriculum learning paradigm has inspired various machine learning algorithms (Caubrière et al. 2019; Duan et al. 2020). In knowledge distillation, LFME (Xiang, Ding, and Han 2020) adopt the classic curriculum strategy and propose to train the student gradually using samples ordered in an easy-to-hard sequence. RCO (Jin et al. 2019) propose to utilize the sequence of the teacher’s intermediate states as the curriculum to gradually guide the learning of a smaller student. The progressive curricula based on data samples and models can help students learn better representations during distillation, but it requires a careful curriculum design and complex computational process, making it hard to deploy into existing methods.

In this paper, we propose a simple and elegant curriculum-based approach, called Curriculum Temperature for Knowledge Distillation (CTKD), which enhances the distillation performance by progressively increasing the learning difficulty level of the student through a dynamic and learnable temperature. The temperature is learned during the student’s training process with a *reversed* gradient that aims to maximize the distillation loss (i.e., increasing the learning difficulty) between teacher and student in an adversarial manner. Specifically, the student is trained under a designed curriculum via the learnable temperature: following the easy-to-hard principle, we gradually increase the distillation loss w.r.t. the temperature, resulting in increased learning difficulty through simply adjusting the temperature dynamically. This operation can be easily implemented by a non-parametric gradient reversal layer (Ganin and Lempitsky 2015) to reverse the gradients of the temperature, which hardly introduces extra computation budgets. Furthermore,

\*Corresponding Author. This work is partially done when Zheng Li is an intern at Megvii.  
Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

based on the curriculum principle, we explore two (global and instance-wise) versions of the learnable temperature, namely Global-T and Instance-T respectively. As an easy-to-use plug-in technique, CTKD can be seamlessly integrated into most existing state-of-the-art KD frameworks and achieves comprehensive improvement at a negligible additional computation cost.

In summary, our contributions are as follows:

- We propose to adversarially learn a dynamic temperature hyperparameter during the student’s training process with a reversed gradient that aims to maximize the distillation loss between teacher and student.
- We introduce simple and effective curricula which organize the distillation task from easy to hard through a dynamic and learnable temperature parameter.
- Extensive experiment results demonstrate that CTKD is a simple yet effective plug-in technique, which consistently improves existing state-of-the-art distillation approaches with a substantial margin on CIFAR-100, ImageNet and MS-COCO.

## Related Work

**Curriculum Learning.** Originally proposed by (Bengio et al. 2009), curriculum learning (Wang, Chen, and Zhu 2021) is a way to train networks by organizing the order in which tasks are learned and incrementally increasing the learning difficulty (Morero et al. 2017; Caubrière et al. 2019). This training strategy has been widely applied in various domains, such as computer vision (Wu et al. 2018; Sinha, Garg, and Larochelle 2020) and natural language processing (Platanios et al. 2019; Tay et al. 2019). Curriculum Dropout (Morero et al. 2017) dynamically increases the dropout ratios in order to improve the generalization ability of the model. PG-GANs (Karras et al. 2017) learn to sequentially generate images from low-resolution to high-resolution, and also grew both generator and discriminator simultaneously. In knowledge distillation, various works (Xiang, Ding, and Han 2020; Zhao et al. 2021) adopt the curriculum learning strategy to train the student model. LFME (Xiang, Ding, and Han 2020) proposes to use the teacher as a difficulty measure and organize the training samples from easy to hard so that the model can receive a less challenging schedule. RCO (Jin et al. 2019) proposes to utilize the sequence of teachers’ intermediate states as a curriculum to supervise the student at different learning stages.

**Knowledge Distillation.** KD (Hinton, Vinyals, and Dean 2015) aims at effectively transferring the knowledge from a pretrained teacher model to a compact and comparable student model. Traditional methods propose to match the output distributions of two models by minimizing the Kullback-Leibler divergence loss with a fixed temperature hyperparameter. To improve distillation performance, existing methods have designed various forms of knowledge transfer. It can be roughly divided into three types, logit-based (Chen et al. 2020; Li et al. 2020b; Zhao et al. 2022), representation-based (Yim et al. 2017; Chen et al. 2021) and relationship-based (Park et al. 2019; Peng et al. 2019) methods. The temperature controls the smoothness of probability distributions

and can faithfully determine the difficulty level of the distillation process. As discussed in (Chandrasegaran et al. 2022; Liu et al. 2022), a lower temperature will make the distillation pays more attention to the maximal logits of teacher output. On the contrary, a higher value will flatten the distribution, making the distillation focus on the logits. Most works ignore the effectiveness of the temperature on distillation and fix it as a hyperparameter that can be decided by an inefficient grid search. However, keeping a constant value, i.e., a fixed level of distillation difficulty, is sub-optimal for a growing student during its progressive learning stages.

Recently, MKD (Liu et al. 2022) proposes to learn the temperature by performing meta-learning on the extra validation set. It mainly works on the ViT (Dosovitskiy et al. 2020) backbone with strong data augmentation while most existing KD methods work under normal augmentation. Directly applying MKD to other distillation methods may weaken the effect of distillation (Das et al. 2020). Our proposed CTKD is more efficient than MKD since we don’t need to pay the effort to split and preserve an extra validation set. Besides, CTKD works under normal augmentation, so it can be seamlessly integrated into existing KD frameworks. The detailed comparison and discussion are attached in the supplement.

## Method

In this section, we first review the concept of knowledge distillation and then introduce our proposed curriculum temperature knowledge distillation technique.

### Background

Knowledge distillation (Hinton, Vinyals, and Dean 2015), as one of the main network compression techniques, has been widely used in many vision tasks (Liu et al. 2019; Ye et al. 2019; Li et al. 2021b, 2022). The traditional two-stage distillation process usually starts with a pre-trained cumbersome teacher network. Then a compact student network will be trained under the supervision of the teacher network in the form of soft predictions or intermediate representations (Romero et al. 2014; Yim et al. 2017). After the distillation, the student can master the expertise of the teacher and use it for final deployment. Given the labeled classification dataset  $D = \{(x_i, y_i)\}_{i=1}^I$ , the Kullback-Leibler (KL) divergence loss is used to minimize the discrepancy between the soft output probabilities of the student and teacher model:

$$L_{kd}(q^t, q^s, \tau) = \sum_{i=1}^I \tau^2 KL(\sigma(q_i^t/\tau), \sigma(q_i^s/\tau)), \quad (1)$$

where  $q^t$  and  $q^s$  denote the logits produced by teacher and student,  $\sigma(\cdot)$  is the softmax function, and  $\tau$  is the temperature to scale the smoothness of two distributions. As discussed in previous works (Hinton, Vinyals, and Dean 2015; Liu et al. 2022), a lower  $\tau$  will sharpen the distribution, enlarge the difference between two distributions and make distillation focus on the maximal logits of teacher prediction. While a higher  $\tau$  will flatten the distribution, narrow the gap between two models and make the distillation focus on whole logits. Therefore, the temperature value  $\tau$  can

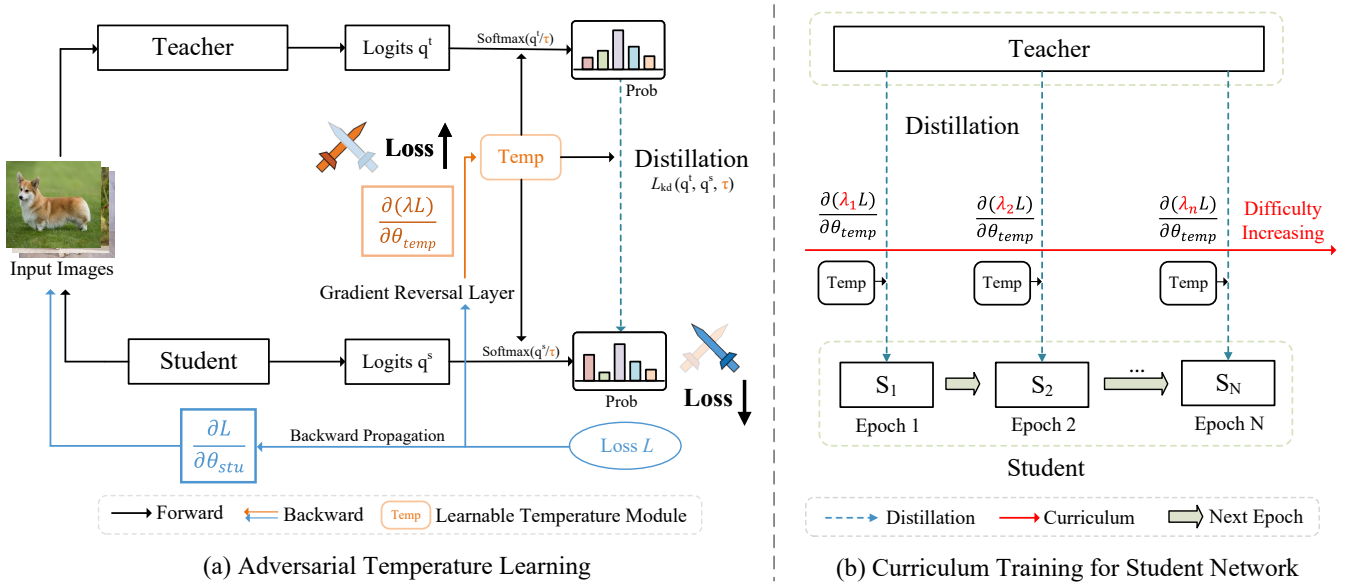


Figure 1: An overview of our proposed Curriculum Temperature for Knowledge Distillation (CTKD). (a) We introduce a learnable temperature module that predicts a suitable temperature  $\tau$  for distillation. The gradient reversal layer is proposed to reverse the gradient of the temperature module during the backpropagation. (b) Following the easy-to-hard curriculum, we gradually increase the parameter  $\lambda$ , leading to increased learning difficulty w.r.t. temperature for the student.

faithfully determine the difficulty level of the KD loss minimization process by affecting the probability distribution.

### Adversarial Distillation

For a vanilla distillation task, the student  $\theta_{stu}$  is optimized to minimize the task-specific loss and distillation loss. The objective of the distillation process can be formulated as follows:

$$\min_{\theta_{stu}} L(\theta_{stu}) = \min_{\theta_{stu}} \sum_{x \in D} \alpha_1 L_{task}(f^s(x; \theta_{stu}), y) + \alpha_2 L_{kd}(f^t(x; \theta_{tea}), f^s(x; \theta_{stu}), \tau). \quad (2)$$

where  $L_{task}$  is the regular cross-entropy loss for the image classification task,  $f^t(\cdot)$  and  $f^s(\cdot)$  denotes the function of teacher and student.  $\alpha_1$  and  $\alpha_2$  are balancing weights.

In order to control the learning difficulty of the student via dynamic temperature, inspired by GANs (Goodfellow et al. 2014), we propose to adversarially learn a dynamic temperature module  $\theta_{temp}$  that predicts a suitable temperature value  $\tau$  for the current training. This module is optimized in the opposite direction of the student, intending to maximize the distillation loss between the student and teacher. Different from vanilla distillation, the student  $\theta_{stu}$  and temperature module  $\theta_{temp}$  play the two-player mini-max game with the following value function  $L(\theta_{stu}, \theta_{temp})$ :

$$\begin{aligned} & \min_{\theta_{stu}} \max_{\theta_{temp}} L(\theta_{stu}, \theta_{temp}) \\ & = \min_{\theta_{stu}} \max_{\theta_{temp}} \sum_{x \in D} \alpha_1 L_{task}(f^s(x; \theta_{stu}), y) \\ & + \alpha_2 L_{kd}(f^t(x; \theta_{tea}), f^s(x; \theta_{stu}), \theta_{temp}). \end{aligned} \quad (3)$$

We apply the alternating algorithm to solve the problem in Eqn. (3), fixing one set of variables and solving for the other set. Formally, we can alternate between solving these two subproblems:

$$\hat{\theta}_{stu} = \arg \min_{\theta_{stu}} L(\theta_{stu}, \hat{\theta}_{temp}), \quad (4)$$

$$\hat{\theta}_{temp} = \arg \max_{\theta_{temp}} L(\hat{\theta}_{stu}, \theta_{temp}). \quad (5)$$

The optimization process for Eqn. (4) and Eqn. (5) can be conducted via stochastic gradient descent (SGD). The student  $\theta_{stu}$  and temperature module  $\theta_{temp}$  parameters are updated as follows:

$$\theta_{stu} \leftarrow \theta_{stu} - \mu \frac{\partial L}{\partial \theta_{stu}}, \quad (6)$$

$$\theta_{temp} \leftarrow \theta_{temp} + \mu \frac{\partial L}{\partial \theta_{temp}}. \quad (7)$$

where  $\mu$  is the learning rate.

In practice, we implement the above adversarial process (i.e., Eqn. (7)) by a non-parametric Gradient Reversal Layer (GRL) (Ganin and Lempitsky 2015). The GRL is inserted between the softmax layer and the learnable temperature module, as shown in Fig. 1(a).

### Curriculum Temperature

Keeping a constant learning difficulty is sub-optimal for a growing student during its progressive learning stages. In school, human teachers always teach students with curricula, which start with basic (easy) concepts, and then gradually present more advanced (difficult) concepts when students




Figure 2: The illustrations of global and instance-wise temperature modules.  $B$  denotes the batch size,  $C$  denotes the number of classes.  $\tau$  is the final temperature.

grow up. Humans will learn much better when the tasks are organized in a meaningful order.

Inspired by curriculum learning (Bengio et al. 2009), we further introduce a simple and effective curriculum which organizes the distillation task from easy to hard via directly scaling the loss  $L$  by magnitude  $\lambda$  w.r.t. the temperature, i.e.,  $L \rightarrow \lambda L$ . Consequently, the  $\theta_{temp}$  would be updated by:

$$\theta_{temp} \leftarrow \theta_{temp} + \mu \frac{\partial(\lambda L)}{\partial \theta_{temp}}. \quad (8)$$

At the beginning of training, the junior student has limited representation ability and requires to learn basic knowledge. We set the initial  $\lambda$  value to 0 so that the junior student can focus on the learning task without any constraints. By gradually increasing  $\lambda$ , the student learns more advanced knowledge as the distillation difficulty increases. Specifically, following the basic concept of curriculum learning, our proposed curriculum satisfies the following two conditions:

(1) Given the unique variable  $\tau$ , the distillation loss w.r.t. the temperature module (simplified as  $L_{kd}(\tau)$ ) gradually increases, i.e.,

$$L_{kd}(\tau_{n+1}) \geq L_{kd}(\tau_n), \quad (9)$$

(2) The value of  $\lambda$  increases, i.e.,

$$\lambda_{n+1} \geq \lambda_n. \quad (10)$$

where  $n$  represents the  $n$ -th step of training.

In our method, when training at  $E_n$  epoch, we gradually increase  $\lambda$  with a cosine schedule as follows:

$$\lambda_n = \lambda_{min} + \frac{1}{2}(\lambda_{max} - \lambda_{min}) \left(1 + \cos\left(\left(1 + \frac{\min(E_n, E_{loops})}{E_{loops}}\right)\pi\right)\right). \quad (11)$$

where  $\lambda_{max}$  and  $\lambda_{min}$  are ranges for  $\lambda$ .  $E_{loops}$  is the hyper-parameter that gradually varies the difficulty scale  $\lambda$ . In our method, we default to set  $\lambda_{max}$ ,  $\lambda_{min}$  and  $E_{loops}$  to 1, 0 and 10, respectively. This curriculum indicates that the parameter  $\lambda$  increases from 0 to 1 during 10 epochs of training and keeps 1 until the end. Detailed ablation studies are conducted in Table 6 and Table 8.

## Learnable Temperature Module

In this section, we introduce two versions of the learnable temperature module, namely Global-T and Instance-T.

**Global-T.** The global version consists of only one learnable parameter, predicting one value  $T_{pred}$  for all instances,

---

## Algorithm 1: Curriculum Temperature Distillation

---

**Input:** Training dataset  $D = \{(x_i, y_i)\}_{i=1}^I$ ; Total training Epoch  $N$ ; Pre-trained Teacher  $\theta_{tea}$ ; Learnable Temperature Module  $\theta_{temp} \in \{\theta_{Global}, \theta_{Instance}\}$ ;

**Output:** Well-trained Student  $\theta_{stu}$ ;

**Initialize:** Epoch  $n=1$ ; Randomly initialize  $\theta_{stu}, \theta_{temp}$ ;

- 1: **while**  $n \leq N$  **do**
  - 2:     **for** data batch  $x$  in  $D$  **do**
  - 3:         Forward propagation through  $\theta_{tea}$  and  $\theta_{stu}$  to obtain predictions  $f^t(x; \theta_{tea}), f^s(x; \theta_{stu})$ ;
  - 4:         Obtain temperature  $\tau$  by  $\theta_{temp}$  in Eqn. (12) and parameter  $\lambda_n$  in Eqn. (11);
  - 5:         Calculate the loss  $L$  and update  $\theta_{stu}$  and  $\theta_{temp}$  by backward propagation as Eqn. (6) and Eqn. (8);
  - 6:     **end for**
  - 7:      $n=n+1$ ;
  - 8: **end while**
- 

as shown in Fig. 2(a). This efficient version does not bring additional computational costs to the distillation process since it only involves a single learnable parameter.

**Instance-T.** To achieve a better distillation performance, one global temperature is not accurate enough for all instances. We further explore the instance-wise variant, termed Instance-T, which predicts a temperature for all instances *individually*, e.g., for a batch of 128 samples, we predict 128 corresponding temperature values. Inspired by GFLv2 (Li et al. 2020a, 2021a), we propose to utilize the statistical information of probability distribution to control the smoothness of itself. Specifically, a 2-layer MLP is introduced in our work, which takes two predictions as input and outputs predicted value  $T_{pred}$ , as shown in Fig. 2(b). During training, the module will automatically learn the implicit relationship between original and smoothed distribution.

To ensure the non-negativity of the temperature parameter and keep its value within a proper range, we scale the predicted  $T_{pred}$  with the following equation:

$$\tau = \tau_{init} + \tau_{range}(\delta(T_{pred})). \quad (12)$$

where  $\tau_{init}$  denotes the initial value,  $\tau_{range}$  denotes the range for  $\tau$ ,  $\delta(\cdot)$  is the sigmoid function,  $T_{pred}$  is the predicted value. We default to set  $\tau_{init}$  and  $\tau_{range}$  to 1 and 20, so that all normal values can be included.

Compared to Global-T, Instance-T can achieve better distillation performance due to its better representation ability.

Teacher	RN-56	RN-110	RN-110	WRN-40-2	WRN-40-2	VGG-13	WRN-40-2	VGG-13	RN-50	RN-32x4	RN-32x4
Acc	72.34	74.31	74.31	75.61	75.61	74.64	75.61	74.64	79.34	79.42	79.42
Student	RN-20	RN-32	RN-20	WRN-16-2	WRN-40-1	VGG-8	SN-V1	MN-V2	MN-V2	SN-V1	SN-V2
Acc	69.06	71.14	69.06	73.26	71.98	70.36	70.50	64.60	64.60	70.50	71.82
Vanilla KD	70.66	73.08	70.66	74.92	73.54	72.98	74.83	67.37	67.35	74.07	74.45
CTKD	71.19	73.52	70.99	75.45	73.93	73.52	75.78	68.46	68.47	74.48	75.31
	(+0.53)	(+0.44)	(+0.33)	(+0.53)	(+0.39)	(+0.54)	(+0.95)	(+1.09)	(+1.12)	(+0.41)	(+0.86)

Table 1: Top-1 accuracy of the student network on CIFAR-100.

Teacher	ResNet-56	ResNet-110	WRN-40-2
Acc	72.34	74.31	75.61
Student	ResNet-20	ResNet-32	WRN-40-1
Acc	69.06	71.14	71.98
Vanilla KD	70.66	73.08	73.54
MACs	41.6M	70.4M	84.7M
Time	10s	15s	17s
Global-T	71.19	73.52	73.93
MACs	41.6M	70.4M	84.7M
Time	10s	15s	17s
Instance-T	71.32	73.61	74.10
MACs	41.7M	70.5M	84.8M
Time	11s	17s	18s

Table 2: Comparison of global and instance-wise CTKD with various backbones on CIFAR-100. “Time”: The time required for one epoch of training.

In the following experiments, we mainly use the global version as the default scheme. We demonstrate the effectiveness of the instance-wise temperature method in Table 2. To get a better understanding of our method, we describe the training procedure in Algorithm 1.

## Experiments

We evaluate our CTKD on various popular neural networks e.g., VGG (Simonyan and Zisserman 2014), ResNet (He et al. 2016) (abbreviated as RN), Wide ResNet (Zagoruyko and Komodakis 2016) (WRN), ShuffleNet (Zhang et al. 2018; Ma et al. 2018) (SN) and MobileNet (Howard et al. 2017; Sandler et al. 2018) (MN). As an easy-to-use plug-in technique, we applied our CTKD to the existing distillation frameworks including vanilla KD (Hinton, Vinyals, and Dean 2015), PKT (Passalis and Tefas 2018), SP (Tung and Mori 2019), VID (Ahn et al. 2019), CRD (Tian, Krishnan, and Isola 2019), SRRL (Yang et al. 2021) and DKD (Zhao et al. 2022). The evaluations are made in comparison to state-of-the-art approaches based on standard experimental settings. All results are reported in means (standard deviations) over 3 trials.

**Dataset.** The CIFAR-100 dataset consists of colored natural images with  $32 \times 32$  pixels. The training and testing sets contain 50K and 10K images, respectively. ImageNet-2012 (Deng et al. 2009) contains 1.2M images for training, and 50K for validation, from 1K classes. The resolution of input images after pre-processing is  $224 \times 224$ . MSCOCO (Lin et al. 2014) is an 80-category general object de-




Figure 3: The curves of distillation loss during training. Our adversarial distillation technique makes the optimization process harder than the vanilla method as expected.

tection dataset. The train2017 split contains 118k images, and the val2017 split contains 5k images.

**Implementation details.** All details are attached in supplement due to the page limit.

## Main Results

**CIFAR-100 classification.** Table 1 shows the top-1 classification accuracy on CIFAR-100 based on eleven different teacher-student pairs. We can observe that all different student networks benefit from our method and the improvement are quite significant in some cases.

Fig. 3 shows the loss curves of vanilla KD and CTKD. During training, the temperature module is optimized to maximize the distillation loss, which satisfies the condition in Eqn. (9). While the student is optimized to minimize the distillation loss, which plays a leading role in this mini-max game. So the overall losses still show a downward trend. As shown in Fig. 3, the distillation loss of CTKD is higher than the vanilla method, proving the effectiveness of adversarial temperature distillation. We can observe that the distillation loss of CTKD is higher than that of vanilla KD, proving the effect of the adversarial operation.

Fig. 4 demonstrates that representations of our method are more separable than vanilla KD, proving that CTKD benefits the discriminability of deep features. Fig. 5 shows the learning curves of temperature during training. Compared to fixed temperature distillation, our curriculum temperature method achieves better results via an effective dynamic mechanism.

**Global and instance-wise temperature.** Table 2 shows the top-1 classification accuracy and computational efficiency (MACs, Time) of the global and instance-wise versions. Since the instance-wise method introduces an addi-

Teacher Acc	ResNet-56 72.34	ResNet-110 74.31	ResNet-110 74.31	WRN-40-2 75.61	WRN-40-2 75.61	ResNet32x4 79.42	ResNet32x4 79.42
Student Acc	ResNet-20 69.06	ResNet-32 71.14	ResNet-20 69.06	WRN-16-2 73.26	WRN-40-1 71.98	ShuffleNet-V1 70.70	ShuffleNet-V2 71.82
PKT	70.85±0.22	73.36±0.15	70.88±0.16	74.82±0.19	74.01±0.23	74.39±0.16	75.10±0.11
+CTKD	71.16±0.08 (+0.31)	73.53±0.05 (+0.17)	71.15±0.09 (+0.27)	75.32±0.11 (+0.52)	74.11±0.20 (+0.10)	74.68±0.16 (+0.29)	75.47±0.19 (+0.37)
SP	70.84±0.25	73.09±0.18	70.74±0.23	74.88±0.28	73.77±0.20	74.97±0.28	75.59±0.15
+CTKD	71.27±0.10 (+0.43)	73.39±0.11 (+0.30)	71.13±0.13 (+0.39)	75.33±0.14 (+0.45)	74.00±0.15 (+0.23)	75.37±0.17(+0.40)	75.82±0.18 (+0.23)
VID	70.62±0.08	73.02±0.10	70.59±0.19	74.89±0.16	73.60±0.26	74.81±0.17	75.24±0.05
+CTKD	70.75±0.11 (+0.13)	73.38±0.24 (+0.36)	71.09±0.24(+0.50)	75.22±0.20 (+0.33)	73.81±0.24 (+0.21)	75.19±0.14 (+0.38)	75.52±0.11 (+0.28)
CRD	71.69±0.15	73.63±0.19	71.38±0.04	75.53±0.10	74.36±0.10	75.13±0.33	75.90±0.15
+CTKD	72.11±0.15 (+0.42)	74.10±0.20 (+0.47)	72.02±0.10 (+0.64)	75.75±0.27 (+0.22)	74.69±0.05 (+0.33)	75.47±0.22 (+0.34)	76.21±0.19 (+0.31)
SRRL	71.13±0.18	73.48±0.16	71.09±0.21	75.69±0.19	74.18±0.03	75.36±0.25	75.90±0.09
+CTKD	71.45±0.15 (+0.32)	73.75±0.30 (+0.27)	71.48±0.14 (+0.39)	75.96±0.06 (+0.27)	74.40±0.13 (+0.22)	75.70±0.22 (+0.34)	76.00±0.22 (+0.10)
DKD	71.43±0.13	73.66±0.15	71.28±0.20	75.70±0.06	74.54±0.12	75.44±0.20	76.48±0.08
+CTKD	71.65±0.24 (+0.27)	74.02±0.29 (+0.36)	71.70±0.10 (+0.42)	75.81±0.14 (+0.11)	74.59±0.08 (+0.05)	75.93±0.29 (+0.49)	76.94±0.04 (+0.46)

Table 3: Top-1 accuracy of the student network on CIFAR-100.

	Teacher	Student	KD	+CTKD	PKT	+CTKD	RKD	+CTKD	SRRL	+CTKD	DKD	+CTKD
Top-1	73.96	70.26	70.83	71.32	70.92	71.29	70.94	71.11	71.01	71.30	71.13	71.51
Top-5	91.58	89.50	90.31	90.27	90.25	90.32	90.33	90.30	90.41	90.42	90.31	90.47

Table 4: Top-1/-5 accuracy on ImageNet-2012. We set ResNet-34 as the teacher and ResNet-18 as the student.




Figure 4: t-SNE of features learned by KD and CTKD.

tional network (i.e., 2-layer MLP) to obtain stronger representation ability, it requires more computational cost than the global version. From this table, we can see that both versions can improve student performance at a negligible additional computational cost. We mainly use the global version in the following experiments.

**Applied to existing distillation works.** As an easy-to-use plug-in technique, CTKD can be seamlessly integrated into existing distillation works. As shown in Table 3, our method brings comprehensive improvements to six state-of-the-art methods based on seven teacher-student pairs. More importantly, CTKD does not incur additional computational costs to the methods since it only contains a lightweight learnable temperature module and a non-parameterized GRL.

**ImageNet-2012 classification.** Table 4 reports the top-1/-5 accuracy of image classification on ImageNet-2012. As a plug-in technique, we also applied our CTKD to four existing state-of-the-art distillation works. The result shows that CTKD can still work on the large-scale dataset effectively.

**MS-COCO object detection.** We also apply our method to the object detection task. We follow the object detection

	mAP	AP50	AP75	API	APm	APs
T: RN-101	42.04	62.48	45.88	54.60	45.55	25.22
S: RN-18	33.26	53.61	35.26	43.16	35.68	18.96
KD	33.97	54.66	36.62	44.14	36.67	18.71
+CTKD	34.56	55.43	36.91	45.07	37.21	19.08
T: RN-50	40.22	61.02	43.81	51.98	43.53	24.16
S: MN-V2	29.47	48.87	30.90	38.86	30.77	16.33
KD	30.13	50.28	31.35	39.56	31.91	16.69
+CTKD	31.39	52.34	33.10	41.06	33.56	18.15

Table 5: Results on MS-COCO based on Faster-RCNN (Ren et al. 2015)-FPN (Lin et al. 2017): AP evaluated on val2017.

implementation of DKD. As shown in Table 5, our CTKD can further boost the detection performance.

## Ablation Study

In the following experiments, we evaluate the effectiveness of hyper-parameters and components on CIFAR-100. We set ResNet-110 as the teacher and ResNet-32 as the student.

**Curriculum parameters.** Table 6 reports the student accuracy with different  $\lambda_{min}$ ,  $\lambda_{max}$ , and  $E_{loops}$ . Table 7 reports the distillation results with different fixed  $\lambda$ . The training of students needs to gradually increase the learning difficulty. Directly starting with a fixed high-difficulty task will significantly reduce the performance of students, especially when  $\lambda$  is greater than 4. Besides, as shown in the sixth and seventh columns of Table 6, rapidly increasing parameter  $\lambda$  in a short time can also be detrimental to student training. When we smooth the learning difficulty of the student and increase  $E_{loops}$ , the performance can be further improved.

**Curriculum strategy.** In Table 8, we compare the performance of different curriculum strategies. “None $_{\tau=1}$ , 10




Figure 5: The learning curves of temperature during training. The yellow dotted line represents the vanilla distillation method at a specified fixed temperature. The solid blue line represents the dynamic temperature learning process. Our dynamic curriculum temperature outperforms the static method.

$E_{loops}$	$[\lambda_{min}, \lambda_{max}]$				
	[0, 1]	[0, 2]	[0, 5]	[0, 10]	[1, 10]
10 Epoch	<b>73.52</b>	73.16	73.12	73.05	72.58
20 Epoch	73.44	73.48	73.01	73.00	72.88
40 Epoch	73.26	73.40	<b>73.50</b>	73.15	72.95
80 Epoch	73.35	73.46	<b>73.52</b>	73.41	73.12
120 Epoch	73.31	73.39	73.16	73.36	73.04
240 Epoch	73.23	73.29	73.20	73.42	73.08

Table 6: Range of dynamic curriculum  $\lambda$ . Smoothly increasing task difficulty is beneficial to students’ learning.

Fixed $\lambda$	1	2	4	5	10	Curriculum
Acc	73.26	73.36	73.16	72.78	72.82	<b>73.52</b>

Table 7: Training with fixed  $\lambda$ . Compared with curriculum  $\lambda$ , directly training the student with the fixed high-difficulty task (e.g.,  $\lambda > 4$ ) will reduce distillation performance.

$E_{loops}$	Curriculum Strategy			
	None $_{\tau=1}$	None $_{\tau=4}$	Lin $_{[0,1]}$	Cos $_{[0,1]}$
10 Epoch	73.21	73.07	73.31	<b>73.52</b>
20 Epoch	73.24	73.06	73.45	73.44
40 Epoch	73.33	73.07	73.10	73.26

Table 8: Comparison of different curriculum strategies. Cosine curriculum strategy works the best.

AT	CD	ResNet-56	ResNet-110	WRN-40-2	VGG-13
		ResNet-20	ResNet-32	WRN-16-2	VGG-8
		70.66	73.08	74.92	72.98
✓		71.01	73.26	74.99	73.43
✓	✓	71.19	73.52	75.45	73.52

Table 9: Ablation of Adversarial Temperature (AT) module and Curriculum Distillation (CD) strategy. The first row indicates the vanilla distillation performance.

## Conclusion

Epoch” means that during the first 10 epochs of training, we only use vanilla distillation, and set the temperature  $\tau = 1$ . After 10 epochs, we start to train the student with CTKD, and  $\lambda$  is fixed to 1. “Lin $_{[0,1]}$ , 10 Epoch” means that we use CTKD to train the student with a linear increasing strategy. The parameter  $\lambda$  is gradually increased from 0 to 1 in 10 epochs of training. The value of  $\lambda$  remains 1 until the end. From this table, we can see that the cosine curriculum strategy works the best.

### Adversarial temperature and curriculum distillation.

We evaluate the effectiveness of these two elements as shown in Table 9. The second row means that we only adopt the adversarial temperature technique and use the fixed learning difficulty (i.e., fix  $\lambda$  to 1) to train the student. The results demonstrate that learning the temperature parameter in an adversarial manner can also improve distillation performance. The third row shows that the cooperation of two elements can achieve better results than a single element.

In this paper, we propose a curriculum-based distillation approach, termed Curriculum Temperature for Knowledge Distillation, which organizes the distillation task from easy to hard through a dynamic and learnable temperature. The temperature is learned during the student’s training process with a reversed gradient that aims to maximize the distillation loss (i.e., increase the learning difficulty) between teacher and student in an adversarial manner. As an easy-to-use plug-in technique, CTKD can be seamlessly integrated into existing state-of-the-art knowledge distillation frameworks and brings general improvements at a negligible additional computation cost.

## Acknowledgments

This work was supported by the Young Scientists Fund of the National Natural Science Foundation of China (Grant No.62206134).

## References

- Ahn, S.; Hu, S. X.; Damianou, A.; Lawrence, N. D.; and Dai, Z. 2019. Variational information distillation for knowledge transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9163–9171.
- Bengio, Y.; Louradour, J.; Collobert, R.; and Weston, J. 2009. Curriculum learning. In *International Conference on Machine Learning*, 41–48.
- Caubrière, A.; Tomashenko, N.; Laurent, A.; Morin, E.; Camelin, N.; and Estève, Y. 2019. Curriculum-based transfer learning for an effective end-to-end spoken language understanding and domain portability. *arXiv preprint arXiv:1906.07601*.
- Chandrasegaran, K.; Tran, N.-T.; Zhao, Y.; and Cheung, N.-M. 2022. Revisiting Label Smoothing and Knowledge Distillation Compatibility: What was Missing? In *International Conference on Machine Learning*, 2890–2916. PMLR.
- Chen, D.; Mei, J.-P.; Wang, C.; Feng, Y.; and Chen, C. 2020. Online Knowledge Distillation with Diverse Peers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3430–3437.
- Chen, P.; Liu, S.; Zhao, H.; and Jia, J. 2021. Distilling knowledge via knowledge review. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5008–5017.
- Das, D.; Massa, H.; Kulkarni, A.; and Rekatsinas, T. 2020. An empirical analysis of the impact of data augmentation on knowledge distillation. *arXiv preprint arXiv:2006.03810*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Proceeding of the IEEE Conference on Computer Vision and Pattern Recognition*, 248–255. Ieee.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Duan, Y.; Zhu, H.; Wang, H.; Yi, L.; Nevatia, R.; and Guibas, L. J. 2020. Curriculum deepsf. In *European Conference on Computer Vision*, 51–67. Springer.
- Ganin, Y.; and Lempitsky, V. 2015. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning*, 1180–1189. PMLR.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2672–2680.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Howard, A. G.; Zhu, M.; Chen, B.; Kalenichenko, D.; Wang, W.; Weyand, T.; Andreetto, M.; and Adam, H. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.
- Ji, M.; Shin, S.; Hwang, S.; Park, G.; and Moon, I.-C. 2021. Refine Myself by Teaching Myself: Feature Refinement via Self-Knowledge Distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10664–10673.
- Jin, X.; Peng, B.; Wu, Y.; Liu, Y.; Liu, J.; Liang, D.; Yan, J.; and Hu, X. 2019. Knowledge distillation via route constrained optimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1345–1354.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2017. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.
- Li, G.; Li, X.; Wang, Y.; Zhang, S.; Wu, Y.; and Liang, D. 2022. Knowledge distillation for object detection via rank mimicking and prediction-guided feature imitation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 1306–1313.
- Li, X.; Wang, W.; Hu, X.; Li, J.; Tang, J.; and Yang, J. 2021a. Generalized focal loss v2: Learning reliable localization quality estimation for dense object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11632–11641.
- Li, X.; Wang, W.; Wu, L.; Chen, S.; Hu, X.; Li, J.; Tang, J.; and Yang, J. 2020a. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. In *Advances in Neural Information Processing Systems*.
- Li, Z.; Huang, Y.; Chen, D.; Luo, T.; Cai, N.; and Pan, Z. 2020b. Online Knowledge Distillation via Multi-branch Diversity Enhancement. In *Proceedings of the Asian Conference on Computer Vision*.
- Li, Z.; Ye, J.; Song, M.; Huang, Y.; and Pan, Z. 2021b. Online knowledge distillation for efficient pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11740–11750.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2117–2125.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, 740–755. Springer.
- Liu, J.; Liu, B.; Li, H.; and Liu, Y. 2022. Meta Knowledge Distillation. *arXiv preprint arXiv:2202.07940*.
- Liu, Y.; Chen, K.; Liu, C.; Qin, Z.; Luo, Z.; and Wang, J. 2019. Structured knowledge distillation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2604–2613.
- Ma, N.; Zhang, X.; Zheng, H.-T.; and Sun, J. 2018. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *European Conference on Computer Vision*, 116–131.



- Morerio, P.; Cavazza, J.; Volpi, R.; Vidal, R.; and Murino, V. 2017. Curriculum dropout. In *Proceedings of the IEEE International Conference on Computer Vision*, 3544–3552.
- Park, W.; Kim, D.; Lu, Y.; and Cho, M. 2019. Relational knowledge distillation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3967–3976.
- Passalis, N.; and Tefas, A. 2018. Learning deep representations with probabilistic knowledge transfer. In *European Conference on Computer Vision (ECCV)*, 268–284.
- Peng, B.; Jin, X.; Liu, J.; Li, D.; Wu, Y.; Liu, Y.; Zhou, S.; and Zhang, Z. 2019. Correlation congruence for knowledge distillation. In *Proceedings of the IEEE International Conference on Computer Vision*, 5007–5016.
- Platanios, E. A.; Stretcu, O.; Neubig, G.; Poczos, B.; and Mitchell, T. M. 2019. Competence-based curriculum learning for neural machine translation. *arXiv preprint arXiv:1903.09848*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28: 91–99.
- Romero, A.; Ballas, N.; Kahou, S. E.; Chassang, A.; Gatta, C.; and Bengio, Y. 2014. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*.
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4510–4520.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Sinha, S.; Garg, A.; and Larochelle, H. 2020. Curriculum by smoothing. *Advances in Neural Information Processing Systems*, 33: 21653–21664.
- Tay, Y.; Wang, S.; Tuan, L. A.; Fu, J.; Phan, M. C.; Yuan, X.; Rao, J.; Hui, S. C.; and Zhang, A. 2019. Simple and effective curriculum pointer-generator networks for reading comprehension over long narratives. *arXiv preprint arXiv:1905.10847*.
- Tian, Y.; Krishnan, D.; and Isola, P. 2019. Contrastive representation distillation. *arXiv preprint arXiv:1910.10699*.
- Tung, F.; and Mori, G. 2019. Similarity-preserving knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1365–1374.
- Wang, X.; Chen, Y.; and Zhu, W. 2021. A survey on curriculum learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wu, L.; Tian, F.; Xia, Y.; Fan, Y.; Qin, T.; Jian-Huang, L.; and Liu, T.-Y. 2018. Learning to teach with dynamic loss functions. *Advances in Neural Information Processing Systems*, 31.
- Xiang, L.; Ding, G.; and Han, J. 2020. Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In *European Conference on Computer Vision*, 247–263. Springer.
- Yang, J.; Martinez, B.; Bulat, A.; Tzimiropoulos, G.; et al. 2021. Knowledge distillation via softmax regression representation learning. In *International Conference on Learning Representations*.
- Ye, J.; Ji, Y.; Wang, X.; Ou, K.; Tao, D.; and Song, M. 2019. Student becoming the master: Knowledge amalgamation for joint scene parsing, depth estimation, and more. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2829–2838.
- Yim, J.; Joo, D.; Bae, J.; and Kim, J. 2017. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4133–4141.
- Zagoruyko, S.; and Komodakis, N. 2016. Wide residual networks. *arXiv preprint arXiv:1605.07146*.
- Zhang, X.; Zhou, X.; Lin, M.; and Sun, J. 2018. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6848–6856.
- Zhao, B.; Cui, Q.; Song, R.; Qiu, Y.; and Liang, J. 2022. Decoupled Knowledge Distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11953–11962.
- Zhao, H.; Sun, X.; Dong, J.; Dong, Z.; and Li, Q. 2021. Knowledge distillation via instance-level sequence learning. *Knowledge-Based Systems*, 233: 107519.