

DOT: A Distillation-Oriented Trainer

Borui Zhao¹ Quan Cui² Renjie Song¹ Jiajun Liang¹

¹MEGVII Technology ²Waseda University

zhaoborui.gm@gmail.com, cui-quan@toki.waseda.jp,
song.renjie@foxmail.com, liangjiajun@megvii.com

Abstract

Knowledge distillation transfers knowledge from a large model to a small one via task and distillation losses. In this paper, we observe a **trade-off** between task and distillation losses, i.e., introducing distillation loss limits the convergence of task loss. We believe that the trade-off results from the insufficient optimization of distillation loss. The reason is: The teacher has a lower task loss than the student, and a lower distillation loss drives the student more similar to the teacher, then a better-converged task loss could be obtained. To break the trade-off, we propose the Distillation-Oriented Trainer (DOT). DOT separately considers gradients of task and distillation losses, then applies a larger momentum to distillation loss to accelerate its optimization. We empirically prove that DOT breaks the trade-off, i.e., both losses are sufficiently optimized. Extensive experiments validate the superiority of DOT. Notably, DOT achieves a **+2.59%** accuracy improvement on ImageNet-1k for the ResNet50-MobileNetV1 pair. Conclusively, DOT greatly benefits the student’s optimization properties in terms of loss convergence and model generalization. <https://github.com/megvii-research/mdistiller>.

1. Introduction

Knowledge distillation [17, 44, 12, 2, 21, 49, 24] has been proved to be an effective manner to transfer knowledge from a heavy (teacher) model to a light (student) one in a wide range of deep learning tasks [40, 14, 35, 7, 4]. Novel learning algorithms have been proposed to achieve better distillation performance [36, 46, 15, 43]. The working mechanism of knowledge distillation also attracts research attention [33, 28, 20, 30, 42, 39, 41, 45]. Yet, the optimization property of knowledge distillation has not been widely investigated, which is also an important perspective to understand KD.

As shown in Figure 1 (left), the typical optimization objective of knowledge distillation is composed of two parts,

a task loss (e.g., the cross-entropy loss) and a distillation loss (e.g., the KL-Divergence [17]). We mainly study how the incremental distillation loss influences the optimization of task loss. Concretely, for an image classification task, we visualize (1) the task loss landscapes and (2) task and distillation loss dynamics during the optimization. As Figure 1 (middle) illustrates, we observe that the distillation loss helps the student converge to flat minima, where the student tends to generalize better due to the robustness of flatter minima [25, 22, 8]. However, as illustrated in Figure 1 (right), introducing distillation loss brings about a **trade-off**. The task loss is not converged as sufficiently as the cross-entropy baseline, although the student’s logits become similar to the teacher’s.

We suppose that the trade-off is somehow *counter-intuitive*. The reason is presented below: the teacher always yields a lower task loss than the student due to the larger model capacity. If the distillation loss is sufficiently optimized, the task loss would also be decreased since the student becomes more similar to the better-performing teacher. We ask: *why is there a trade-off and how to break it?* We attempt to answer this question from the following perspective. The task and distillation loss terms are combined with a simple summation in practical implementations of popular distillation methods [17, 36, 46, 43, 15]. It could make the optimization manner degrade to multi-task learning, where the network attempts to find a balance between the two tasks. As aforementioned, if the distillation loss is sufficiently optimized, *both* losses would be decreased. Thus, we believe that *sufficiently optimizing the distillation loss is the key to breaking the trade-off*.

To this end, we present the Distillation-Oriented Trainer (DOT) which enables the distillation loss to *dominate* the optimization. It separately considers the gradients provided by the task and the distillation losses, then adjusts the optimization orientation by weighing different *momentums*. A larger momentum is applied to distillation loss, while a smaller one is to task loss. It ensures the optimization is dominated by gradients of distillation loss since a larger momentum accumulates larger and more consistent

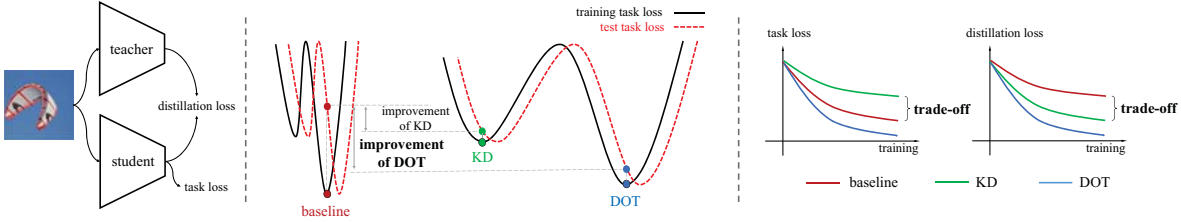


Figure 1: *Left:* the framework of knowledge distillation (KD). KD introduces an extra distillation loss, transferring knowledge from the teacher model. *Middle:* a conceptual sketch of flat and sharp local minima [22, 13]. The Y-axis denotes the loss value, and the X-axis the network parameters. The considerable sensitivity of the training loss at sharp minima damages the generalization on test data. In this paper, we discover that knowledge distillation (KD) benefits the student baseline (CE) with flatter minima but unexpectedly limits the convergence. See Figure 2 and 3. *Right:* the task and distillation loss dynamics. It suggests that introducing KD brings about a trade-off between the task and distillation losses. See Figure 3. To address this trade-off issue and achieve better performance, we propose Distillation-Oriented Trainer (DOT). Our DOT breaks the trade-off and leads the student to ideal minima of both great flatness and convergence.

gradients than a smaller one. In this way, DOT ensures a sufficient optimization of distillation loss. We validate the effectiveness of DOT from three perspectives. (1) As illustrated in Figure 1 (right), we prove that DOT successfully breaks the trade-off between task and distillation losses. (2) As illustrated in Figure 1 (middle), our DOT achieves more generalizable and flatter minima, empirically supporting the benefits of DOT. (3) DOT improves performances of popular distillation methods without bells and whistles, achieving new SOTA results.

More importantly, our research brings new insights into the knowledge distillation community. We show great potential for a better optimization manner of knowledge distillation. To the best of our knowledge, we provide the first attempt to understand the working mechanism of knowledge distillation from the optimization perspective.

2. Related Work

Knowledge distillation. Ideas correlated to distillation can date back to [1, 5, 26, 27], and the knowledge distillation concept has become widely known since the application in compressing a heavy network into a light one [17]. Following representative works can be divided into two categories, *i.e.*, distilling knowledge from logits [17, 3, 48, 10] and intermediate features [36, 46, 31, 16, 15, 43]. More and more attempts have been made on understanding how and why knowledge distillation helps the network learning [33, 28, 20, 30, 42, 39, 45, 41] recently. KD [17] conjectures that the improvement comes from network predictions on incorrect classes. [28] explains distillation from a privileged information perspective. [33] studies distillation with linear classifiers and proposes that the success of distillation owes to data geometry, optimization bias, and strong monotonicity. [20] explains the practice of mixing task loss and distillation loss with the data inefficiency concept. Few

efforts are made to analyze distillation from the network optimization perspective, while we provide the first attempt in this work.

Flatness of minima. The minima of neural networks has attracted great research attention [22, 18, 19, 8, 13]. An acknowledged hypothesis is that the flatness of converged minima can influence the generalization ability of networks [22], and flatter ones correspond to better generalization ability. The explanation is that flat ones reduce generalization errors since there are random perturbations around the loss landscape and the flatter one is more robust. [13] also proves the correctness of this hypothesis. Thanks to previous efforts in analyzing loss landscapes [25, 11, 32, 9], we are allowed to visualize the converged minima learned by knowledge distillation to understand its working mechanism.

In this work, we investigate the optimization process of the most representative method KD by visualizing the flatness of minima for the first time. We show that introducing distillation loss benefits the student with better generalization flat minima while resulting in higher task loss. This trade-off between the task and distillation losses is counter-intuitive and mostly overlooked. To this end, we study the optimization process and propose a method to break the trade-off and approach ideal converged minima.

3. Revisiting Knowledge Distillation: An Optimization Perspective

3.1. Recap of Knowledge Distillation

The working mechanism of knowledge distillation methods has been explored from various perspectives [33, 28, 20, 30, 42, 39, 45, 41]. Previous works seldom delve into the optimization property of knowledge distillation. In this section, we explore the optimization behavior of knowledge

distillation by studying how the incremental distillation loss influences the optimization property.

We study the most representative knowledge distillation method KD [17] for easy understanding. The practical loss function of KD could be written as:

$$\begin{aligned} \mathcal{L} &= \alpha \mathcal{L}_{\text{CE}}(x, y; \theta) + (1 - \alpha) \mathcal{L}_{\text{KD}}(x, \phi; \theta) \\ &= \alpha H(y, \sigma(p)) + (1 - \alpha) D_{\text{KL}}(\sigma(p/T) \parallel \sigma(q/T)), \end{aligned} \quad (1)$$

where the input and its label are denoted as x and y . θ and ϕ are the parameters of the student and teacher networks. p and q are the output logits from student and teacher networks respectively. H is the cross-entropy loss function, D_{KL} is the KL-divergence, and σ is the softmax function. The temperature T is introduced to soften predictions and arise attention on negative logits.

Eqn. (1) manifests that the practical optimization objective \mathcal{L} is composed of a task loss \mathcal{L}_{CE} and a distillation loss \mathcal{L}_{KD} . We mainly study how the incremental \mathcal{L}_{KD} influences the optimization property. First, we explore the influence of distillation loss on loss landscapes which can be used to measure the generalization ability of the converged minima. Second, we visualize the loss dynamics and reveal a *trade-off* between task and distillation losses along the entire optimization process.

3.2. Loss Landscapes

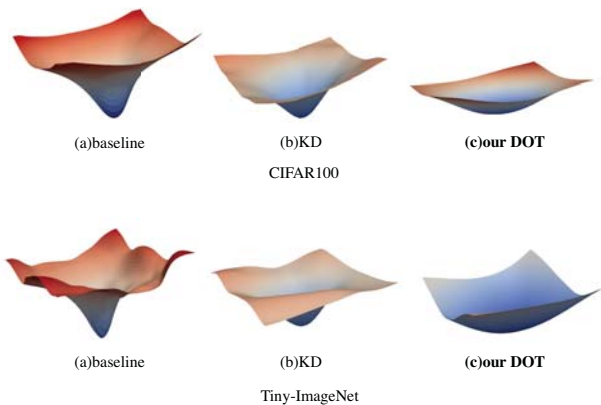


Figure 2: Loss landscapes on CIFAR-100 and Tiny-ImageNet. KD leads the student to flat minima with robust generalization ability. Our DOT achieves much *flatter* minima, further improving the model generalization and distillation performances.

It is well-acknowledged that the model generalization performance is characterized by the flatness of the local minima [22, 18, 19, 8, 13]. As illustrated in Figure 1 (*middle*), sharp minima correspond to a large gap between training and test loss values, *i.e.*, inferior generalization ability, but flat minima tend to reduce the generalization errors.

After respectively training student networks with \mathcal{L}_{CE} (the baseline) and \mathcal{L} (the KD [17]) on CIFAR-100 [23] (typical ResNet32×4-ResNet8×4 pair) and Tiny-ImageNet (typical ResNet18-MobileNetV2 pair), we visualize the task loss landscapes of the converged student networks to study the flatness of local minima in Figure 2. Compared to the baseline trained with only \mathcal{L}_{CE} , optimizing with \mathcal{L} helps the task loss converge to flatter minima, which explains the generalization improvement of student networks. Conversely, training with only the task loss \mathcal{L}_{CE} leads to the sharp minima, resulting in unsatisfactory generalization performance on the test distribution. It suggests that the improvements made by knowledge distillation methods are attributed to enabling the student to converge around flatter minima.

3.3. Trade-off Between Distillation and Task Losses

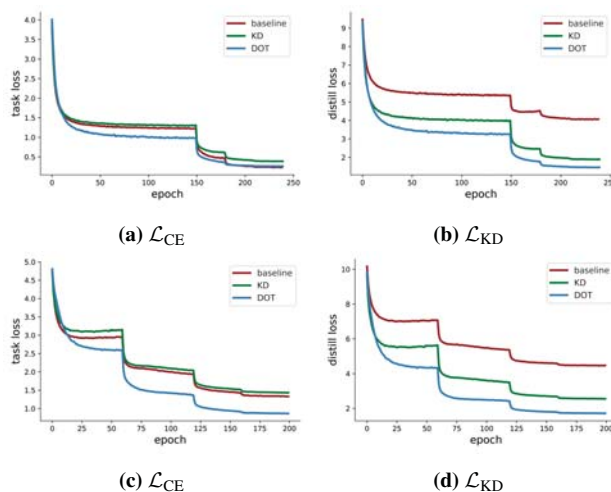


Figure 3: Task and distillation loss curves on CIFAR-100 and Tiny-ImageNet. It suggests that there is a trade-off between task and distillation losses. DOT could break the trade-off, achieving lower task and lower distillation losses *at the same time*.

In Figure 3, we visualize the *training* task and distillation loss dynamics with the progress of optimization. For the KD method, we respectively calculate the task loss and the distillation loss of each epoch. For the baseline, we are only allowed to obtain task loss, so we calculate the KL-Divergence between logits of the student and teacher on-the-fly (without back-propagation). The datasets and the teacher-student pairs are the same as those in Section 3.2. As shown in Figure 3, the introduction of distillation loss (KD) greatly decreases \mathcal{L}_{KD} , because the output logits of the student become more similar to the teacher. However, the task loss \mathcal{L}_{CE} (on training set) is increased. It indicates that the network attempts to find a trade-off between task and distillation losses. and the converged student could be a “Pareto optimum”¹.

¹The trade-off can not be solved by simply tuning the loss weights,

We attempt to explain the trade-off from the perspective of multi-task learning. The targets of both losses are not identical, and learning multiple tasks at the same time makes the optimization difficult [47]. Therefore, it is reasonable that a trade-off exists between task and distillation losses. And the aforementioned observations prove that the task training loss is increased due to the introduction of distillation loss ². We suppose that regarding the optimization of task and distillation losses as multiple tasks is improper for the following reason: The teacher’s training and test losses are both lower than the student’s due to the larger model capacity. Making the student similar to the teacher can help yield *both* lower distillation *and* test losses. Thus, if the distillation loss can be sufficiently optimized, both task loss and distillation loss would be decreased. It inspires us to design an optimization manner where the distillation loss could be more sufficiently optimized.

4. Method

Knowledge distillation benefits the student network with flatter minima, yet introduces a trade-off between the task and the distillation losses. We suppose that the key to breaking the trade-off is making the optimization have a *dominant orientation*, which could reduce gradient conflicts and ensure better network convergence. Since the distillation loss helps the student similar to the teacher, making it dominant (instead of the task loss) could leverage the knowledge and achieve better generalization ability.

4.1. Making KD Dominate the Optimization

Optimizer with momentum. Firstly, we revisit the widely-used Stochastic Gradient Descent (SGD) optimizer. SGD (with momentum) updates the network parameters (denoted as θ) with both the current gradients (denoted as $\mathbf{g} = \nabla_{\theta} \mathcal{L}(\theta)$) and the historical ones. Specifically, SGD maintains a grad buffer named “momentum buffer”(denoted as \mathbf{v}) for network parameters. For every training mini-batch data, SGD updates the momentum buffer by:

$$\mathbf{v} \leftarrow \mathbf{g} + \mu \mathbf{v}, \tag{2}$$

and then the parameters will be updated following the gradient-descent rule:

$$\theta \leftarrow \theta - \gamma \mathbf{v}, \tag{3}$$

where μ and γ denote the momentum coefficient and the learning rate, respectively. Utilizing the momentum buffer can benefit the optimization process with the historical gradients. [34] shows that using momentum can considerably

proofs are presented in Figure 6 of Section 5

²We also conduct experiments with longer training time in the supplement, to validate that the high loss value is not due to inadequate training.

accelerate convergence to a local minimum. Empirically, the momentum coefficient μ is not the larger the better, and 0.9 is the most used value.

Independent momentums for distillation and task losses.

Momentum is a widely-used technique for accelerating gradient descent that accumulates a velocity vector in directions of persistent reduction in the objective across iterations [34]. Under the knowledge distillation framework, setting independent momentums for the grads provided by different losses (*i.e.*, the distillation loss and the task loss) could play an important role in controlling the optimization orientation. Independent momentums enable the loss with the larger momentum to dominate the optimization from two aspects: (1) A large momentum on the distillation loss ensures the optimization orientation is knowledge-transfer-friendly in the initial “transient phase” [6]. (2) A large momentum keeps the historical gradient value undiminished in later training, ensuring the consistency of optimization orientation.

Distillation Oriented Trainer. Driven by the analysis above, we present Distillation-Oriented Trainer (DOT) ³. As shown in Figure 4, DOT maintains two individual momentum buffers for the gradients of CE loss and KD loss. The two momentum buffers are denoted as \mathbf{v}_{ce} and \mathbf{v}_{kd} . DOT updates \mathbf{v}_{ce} and \mathbf{v}_{kd} with different momentum coefficients. DOT introduces a hyper-parameter Δ and sets the coefficients for \mathbf{v}_{ce} and \mathbf{v}_{kd} as $\mu - \Delta$ and $\mu + \Delta$, respectively. Given a single mini-batch data, DOT first computes the gradients (denoted as \mathbf{g}_{ce} and \mathbf{g}_{kd}) produced by \mathcal{L}_{CE} and \mathcal{L}_{KD} respectively, then updates the momentum buffers according to:

$$\begin{aligned} \mathbf{v}_{ce} &\leftarrow \mathbf{g}_{ce} + (\mu - \Delta) \mathbf{v}_{ce}, \\ \mathbf{v}_{kd} &\leftarrow \mathbf{g}_{kd} + (\mu + \Delta) \mathbf{v}_{kd}. \end{aligned} \tag{4}$$

Finally, the network parameters are updated with the sum of two momentum buffers:

$$\theta \leftarrow \theta - \gamma (\mathbf{v}_{ce} + \mathbf{v}_{kd}) \tag{5}$$

DOT applies larger momentum to the distillation loss \mathcal{L}_{KD} and smaller momentum to the task loss \mathcal{L}_{CE} . Thus, the optimization orientation could be dominated by the gradients of the distillation loss. DOT better leverages the knowledge from the teacher and mitigates the trade-off problem caused by insufficient optimization ⁴.

4.2. Theoretical Analysis of Gradient

To investigate the difference between DOT and a vanilla trainer, we dissect the gradients of task and distillation

³DOT can be applied to all optimization methods with momentum mechanism. In this paper, we apply DOT to SGD since it is the most common one in the knowledge distillation community.

⁴The algorithm of DOT and more details about DOT’s implementation are attached in the supplement.

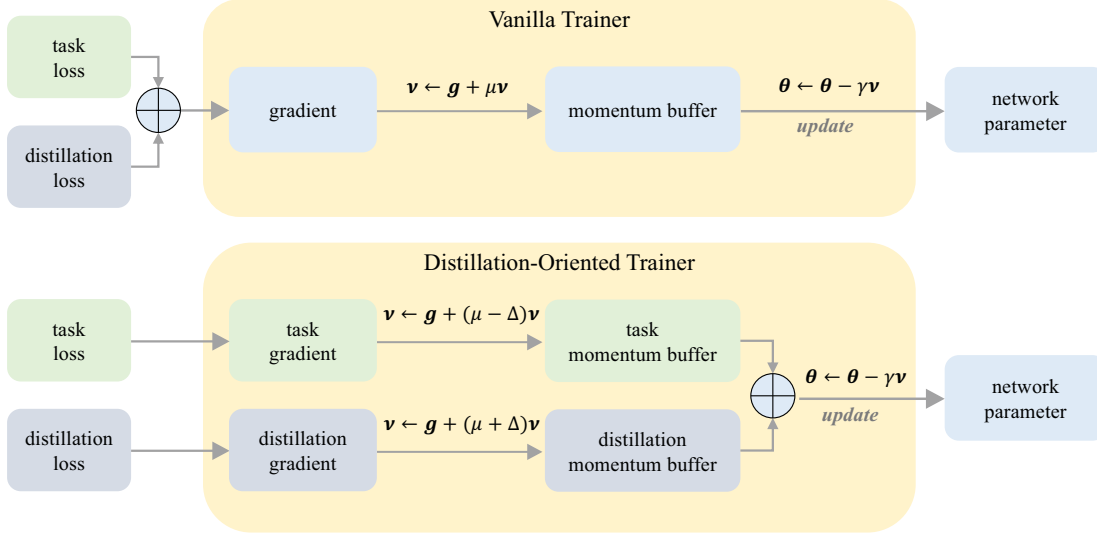


Figure 4: Illustration of a vanilla trainer and our Distillation-Oriented Trainer (DOT). DOT *separately* calculates the gradients of task and distillation losses, then applies *larger momentum* to the distillation gradients and *smaller momentum* to the task ones.

losses as:

$$\begin{aligned} \mathbf{v}_{ce} &= \mathbf{v}^{\text{con}} + \mathbf{v}_{ce}^{\text{incon}}, \\ \mathbf{v}_{kd} &= \mathbf{v}^{\text{con}} + \mathbf{v}_{kd}^{\text{incon}}. \end{aligned} \quad (6)$$

\mathbf{v}^{con} denotes the ‘‘consistent’’ gradient component of both losses. $\mathbf{v}_{ce}^{\text{incon}}$ and $\mathbf{v}_{kd}^{\text{incon}}$ are the ‘‘inconsistent’’ components. For the SGD baseline, the updated momentum buffer can be written as :

$$\begin{aligned} \mathbf{v}_{\text{sgd}} &= \mathbf{g}_{ce} + \mathbf{g}_{kd} + \mu(\mathbf{v}_{ce} + \mathbf{v}_{kd}) \\ &= \mathbf{g}_{ce} + \mathbf{g}_{kd} + \mu(\mathbf{v}_{ce}^{\text{incon}} + \mathbf{v}_{kd}^{\text{incon}} + 2\mathbf{v}^{\text{con}}) \end{aligned} \quad (7)$$

For DOT, the updated momentum buffer is :

$$\begin{aligned} \mathbf{v}_{\text{dot}} &= \mathbf{g}_{ce} + \mathbf{g}_{kd} + (\mu - \Delta)\mathbf{v}_{ce} + (\mu + \Delta)\mathbf{v}_{kd} \\ &= \mathbf{g}_{ce} + \mathbf{g}_{kd} + (\mu - \Delta)(\mathbf{v}_{ce}^{\text{incon}} + \mathbf{v}^{\text{con}}) \\ &\quad + (\mu + \Delta)(\mathbf{v}_{kd}^{\text{incon}} + \mathbf{v}^{\text{con}}) \end{aligned} \quad (8)$$

Comparing Eqn. (7) and Eqn. (8), the difference between the two methods is calculated as follows:

$$\begin{aligned} \mathbf{v}_{\text{diff}} &= \mathbf{v}_{\text{dot}} - \mathbf{v}_{\text{sgd}} \\ &= \Delta(\mathbf{v}_{kd}^{\text{incon}} - \mathbf{v}_{ce}^{\text{incon}}). \end{aligned} \quad (9)$$

Eqn. (9) indicates that when the gradient of task loss conflicts with that of distillation loss, DOT produces gradient $\Delta(\mathbf{v}_{kd}^{\text{incon}} - \mathbf{v}_{ce}^{\text{incon}})$ to accelerate the accumulation of the gradient of the distillation loss. Thus, the optimization is driven by the orientation of distillation loss⁵. The student network would achieve better convergence since the conflict between gradients is alleviated.

⁵Besides, we also conduct a toy experiment for better understanding our DOT in the supplement.

5. Experiments

In this part, we first detail the implementations of all experiments. Second, we empirically validate our motivations and then discuss the main results on popular image classification benchmarks. We also provide analysis (*e.g.*, visualizations) for more insights.

5.1. Experimental Settings

Datasets. We conduct experiments on three image classification datasets, CIFAR-100 [23], Tiny-ImageNet and ImageNet-1k [37]. CIFAR-100 is a well-known image classification dataset that consists of 100 classes. The image size is 32×32 . Training and validation sets are composed of 50k and 10k images, respectively. Tiny-ImageNet consists of 200 classes and the image size is 64×64 . The training set contains 100k images and the validation contains 10k images. ImageNet-1k is a large-scale classification dataset that consists of 1k classes. The training set contains 1.28 million images and the validation contains 50k images. All images are cropped to 224×224 .

Implementations. All models are trained three times and we report the average accuracy.

For *CIFAR-100*, we follow the settings in [43]. We train all models for 240 epochs with learning rates decayed by 0.1 at 150th, 180th, and 210th epoch. The initial learning rate is set as 0.01 for MobileNetV2 [38] and ShuffleNetV2 [29], and 0.05 for ResNet8 \times 4. The batch size is 64 for all models. We use SGD with 0.9 momentum and 0.0005 weight decay as the optimizer. We find Δ ranging from 0.05 to 0.075 works well in all experiments. All models are trained on a single GPU.

For *Tiny-ImageNet*, we follow the settings in [45]. All the models are trained for 200 epochs with learning rates decayed by 0.1 at 60th, 120th, and 160th epoch. The initial learning rate is 0.05 for a 64 batch size. We use SGD with 0.9 momentum and 0.0005 weight decay as the optimizer. The hyper-parameter Δ for our DOT is set to 0.075. All models are trained on 4 GPUs.

For *ImageNet-1k*, we use the standard training recipe in [43]. All models are trained for 100 epochs and the learning rates decayed by 0.1 at 30th, 60th, 90th epoch. The initial learning rate is 0.2 for a 512 batch size. We use SGD with 0.9 momentum and 0.0001 weight decay as the optimizer. The hyper-parameter Δ for our DOT is set to 0.09. All models are trained on 8 GPUs.

5.2. Motivation Validations

The verification of our conjectures and motivations is the most important experiment. In this part, We mainly validate that (1) the distillation loss is the dominant component, (2) the trade-off problem is alleviated, and (3) the student is benefited with reasonable flatter minima compared with the classical KD. Additionally, we also conduct ablation experiments for μ and Δ .

Does KD Loss dominate the optimization? To verify that DOT makes KD loss dominate the optimization, we analyze the gradients of DOT and a vanilla trainer in the optimization process. As shown in Figure 5, DOT leads to the following consequences: (1) The cosine similarity between gradients of distillation and total losses is significantly increased, and (2) cosine similarity between gradients of task and total losses is decreased. It suggests that the optimization orientation prefers distillation gradients after applying our DOT, which proves that DOT enables the distillation loss to dominate the optimization.

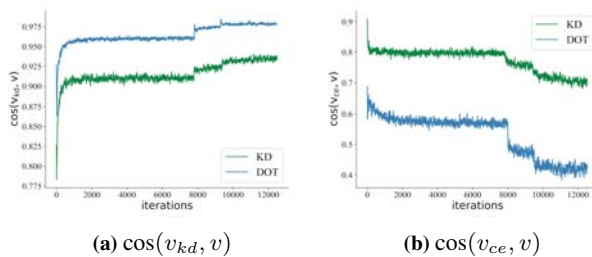


Figure 5: Cosine similarities. We visualize the $\cos(v_{kd}, v)$ and $\cos(v_{ce}, v)$ of each training iteration. It indicates that applying DOT greatly increases $\cos(v_{kd}, v)$, showing the momentum gradients are *dominated* by the distillation component.

Does KD and Task Losses converge better? We visualize the training loss curves to study the convergence of task and distillation losses. As shown in Figure 3, DOT *simultaneously* achieves lower task and distillation losses than the

vanilla trainer and the cross-entropy baseline. It strongly supports our motivation that sufficiently optimized distillation loss contributes to low task loss. Moreover, it also demonstrates that dominating the optimization with distillation loss eliminates the trade-off between task and distillation losses.

Does DOT lead to better minima? As shown in Figure 2 we provide visualizations based on experimental settings in Section 3, which reveal the effect of applying DOT to the representative method KD [17]. Concretely, the loss landscapes of baseline, KD, and DOT on the CIFAR-100 training set are respectively illustrated. It is proven that DOT leads the student to minima with *great flatness*, which is even flatter than the ones of KD. Conclusively, DOT achieves the goal to learn minima with both good flatness and low task loss, as also illustrated in Figure 1 (*middle*).

Δ	0.0	+0.025	+0.05	+0.075	+0.09
top-1	73.33	74.08	74.22	75.12	74.43
Δ	-0.09	-0.075	-0.05	-0.025	0.0
top-1	61.58	68.88	72.77	72.86	73.33

Table 1: Different Δ on CIFAR-100. The results indicate that “distillation-oriented” (positive Δ) could achieve **stable** improvements among different hyper-parameter values while “task-oriented” (negative Δ) leads to performance drops.

μ	0.81	0.825	0.85	0.875	0.9
top-1	73.53	73.66	73.45	73.21	73.33
μ	0.9	0.925	0.95	0.975	0.99
top-1	73.33	73.16	73.14	68.75	37.19

Table 2: Ablation study for different μ on CIFAR-100. It indicates that tuning μ brings no significant performance improvement.

Are independent momentums necessary? To prove that the improvements come from the design of the independent momentum mechanism instead of carefully hyper-parameter tuning, we explore the influence of different hyper-parameter Δ for DOT. As shown in Table 1, setting $\Delta = 0$ means training task and distillation losses with the same momentum (*i.e.*, equals to a vanilla SGD). Applying DOT with $\Delta > 0$ (ranging from 0.025 to 0.09) enables the distillation loss to dominate optimization, which contributes to *stable* performance improvements. It indicates that distillation-oriented optimization is of vital importance. Moreover, we also experiment with $\Delta < 0$ (ranging from -0.09 to -0.025), corresponding to a “task-oriented trainer” where the task loss dominates the optimization. Significant performance drops could be observed in Table 1, which further supports our motivation to make the distillation loss dominant instead of the task loss.

Are improvements attributed to tuning momentum? Without utilizing DOT, we also explore the distillation per-

	teacher	student	KD	AT	CRD	DKD	KD+DOT	CRD+DOT	DKD+DOT
<i>ResNet32×4 as the teacher, ResNet8×4 as the student</i>									
top1	79.42	72.50	73.33	73.44	75.51	76.32	75.12 (+1.79)	75.99 (+0.48)	76.64 (+0.32)
top5	94.58	92.72	93.00	93.06	93.92	93.94	93.46 (+0.46)	94.09 (+0.17)	94.03 (+0.09)
<i>VGG13 as the teacher, VGG8 as the student</i>									
top1	74.64	70.36	72.98	71.43	73.94	74.68	73.77 (+0.79)	74.21 (+0.27)	74.86 (+0.18)
top5	92.60	90.57	92.27	91.94	92.25	92.62	92.33 (+0.05)	92.72 (+0.47)	92.81 (+0.19)
<i>ResNet32×4 as the teacher, ShuffleNetV2 as the student</i>									
top1	79.42	71.82	74.45	72.73	75.65	77.07	75.55 (+1.10)	76.64 (+0.99)	77.41 (+0.34)
top5	94.58	91.77	93.02	93.07	93.71	94.19	93.24 (+0.22)	94.00 (+0.29)	94.13 (-0.06)

(a) CIFAR-100

	teacher	student	KD	AT	CRD	DKD	KD+DOT	CRD+DOT	DKD+DOT
<i>ResNet18 as the teacher, MobileNet-V2 as the student</i>									
top1	62.99	56.28	58.35	57.18	61.18	62.04	64.01 (+5.66)	64.12 (+2.94)	64.60 (+2.56)
top5	83.36	80.32	82.07	81.52	83.13	84.12	84.30 (+2.23)	84.43 (+1.30)	85.38 (+1.26)
<i>ResNet18 as the teacher, ShuffleNetV2 as the student</i>									
top1	62.99	60.78	62.26	62.45	63.97	65.06	65.75 (+3.49)	65.21 (+1.24)	66.21 (+1.15)
top5	83.36	82.49	83.79	83.51	84.70	85.31	85.51 (+1.72)	85.13 (+0.43)	86.16 (+0.85)

(b) Tiny ImageNet

	teacher	student	KD	AT	OFD	CRD	DKD	KD+DOT	DKD+DOT
<i>ResNet34 as the teacher, ResNet18 as the student</i>									
top1	73.31	69.75	71.03	70.69	70.81	71.17	71.70	71.72 (+0.69)	72.03 (+0.33)
top5	91.42	89.07	90.05	90.01	89.98	90.13	90.05	90.30 (+0.25)	90.50 (+0.45)
<i>ResNet50 as the teacher, MobileNetV1 as the student</i>									
top1	76.16	68.87	70.50	69.56	71.25	71.37	72.05	73.09 (+2.59)	73.33 (+1.27)
top5	92.86	88.76	89.80	89.33	90.34	90.41	91.05	91.11 (+1.31)	91.22 (+0.17)

(c) ImageNet-1k

Table 3: Applying our DOT can bring significant performance gains and achieve **new state-of-the-art** distillation results on CIFAR-100, Tiny-ImageNet and ImageNet-1k. Both logit- and feature-based methods could benefit from our DOT. Notably, for the ResNet50-MobileNetV1 pair on ImageNet-1k, DOT achieves a significant **+2.59%** accuracy gain.

formance when training with different momentums μ to verify that the gains mainly come from our design of independent momentums instead of better momentum values. Results in Table 2 show that carefully tuning μ could lead to certain performance gain (73.33% v.s. 73.66%), but not as significant as DOT’s improvement (73.33% v.s. 75.12%).

5.3. Main Results

Following the common setting, we benchmark our DOT on three popular image classification datasets, *i.e.*, CIFAR-100, Tiny-ImageNet, and ImageNet-1k. Additionally, we prove that DOT is compatible with representative distillation methods, and contributes to new state-of-the-art results.

CIFAR-100. Results of CIFAR-100 in Table 3 (a) show that our DOT could contribute to significant performance gains for the classical knowledge distillation⁶. For instance, DOT

improves the classical KD method from 73.33% to 75.12% on the ResNet32×4-ResNet8×4 teacher-student pair. To prove the scalability of DOT, we combine DOT with popular logit-based and feature-based distillation methods. We select DKD [48] and CRD [43] as representative logit-based and feature-based methods, respectively. As shown in Table 3 (a), DOT still succeeds in advancing the performances of evaluated methods, supporting DOT’s practicability.

Tiny-ImageNet. Tiny-ImageNet is a more challenging dataset than CIFAR-100. Results in Table 3 (b) demonstrate that our DOT still achieves *more remarkable* performance gains on such a challenging dataset. DOT greatly improves the top1 accuracy from 58.35% to 64.01% on the ResNet18-MobileNetV2 teacher-student pair, and improves the top1 accuracy from 62.26% to 65.75% on the ResNet18-ShuffleNetV2 pair. DOT also increases the SOTA performance to 64.60% and 66.21% on SOTA methods.

ImageNet-1k. We also conduct experiments on ImageNet-

⁶More pairs on CIFAR-100 can be attached to the supplement.

1k. Experimental results in Table 3 (c) consistently validate the superiority of DOT. Especially for the ResNet50-MobileNetV1 pair, DOT achieves a +2.59% accuracy gain on the classical KD method, outperforming previous state-of-the-art methods. It strongly demonstrates that the optimization of knowledge distillation methods deserves further exploration. Additionally, applying DOT to DKD could further increase the state-of-the-art performance to a new 73.33% milestone.

5.4. More Analysis

In this part, we first investigate whether simply tuning the loss weight α can break the trade-off, then visualize the distillation fidelity for intuitive understanding.

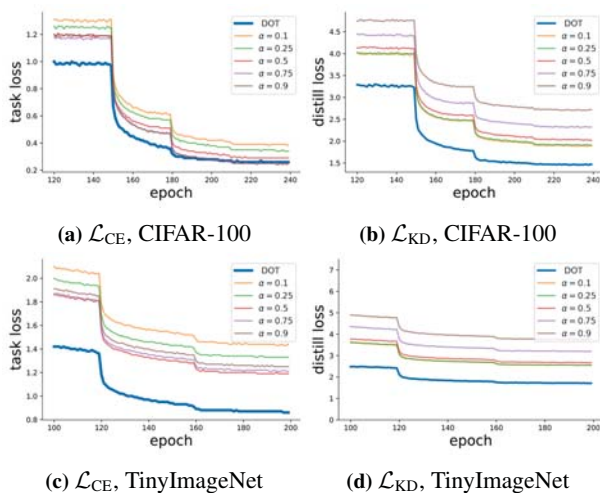


Figure 6: Illustration of loss curves for weighing α . It indicates that simply tuning loss weights still suffers from the trade-off. A larger α leads to a lower task loss but a higher distillation loss. In contrast, our DOT effectively **breaks the trade-off** and achieves both low task and distillation losses, *i.e.*, better convergence.

Adjusting α can not alleviate the trade-off. Intuitively, adjusting α could somehow help the task loss converge better too. As shown in Table 4, Table 5 and Figure 6, we adjust the weight α of the task loss (and the weight of the distillation loss $1 - \alpha$ is correspondingly adjusted) and provide the distillation results and the loss curves. Intuitively, the larger α strengthens the influence of the task loss on the optimization and weakens the influence of the distillation loss. Table 4 and 5 show that weighing α can not exert noticeable influences on final distillation performances. What’s more, Figure 6 suggests that the larger α leads to lower task loss but higher distillation loss, which means the network still *suffers from* the trade-off. Conversely, applying DOT achieves both lower task and distillation losses simultaneously, *i.e.*, the trade-off between task and distillation losses is successfully alleviated.

Distillation fidelity. We visualize the distillation fidelity

α	0.1	0.25	0.5	0.75	0.9	DOT
top-1	73.33	73.56	73.49	73.23	73.19	75.12

Table 4: Different weight (α) for task loss on CIFAR-100. ResNet32×4-ResNet8×4 as the teacher-student pair.

α	0.1	0.25	0.5	0.75	0.9	DOT
top-1	58.35	58.86	59.23	58.36	57.45	64.01

Table 5: Different weight (α) for task loss on Tiny-ImageNet. ResNet18-MobileNetV2 as the teacher-student pair

following [43, 48] for intuitive understanding. Concretely, for ResNet32×4-ResNet8×4 (CIFAR-100) and ResNet18-MobileNetV2 (Tiny-ImageNet) pairs, we calculate the absolute distance between correlation matrices of the teacher and the student. Compared with KD, introducing DOT helps the student output more similar logits to the teacher, resulting in better distillation performance.

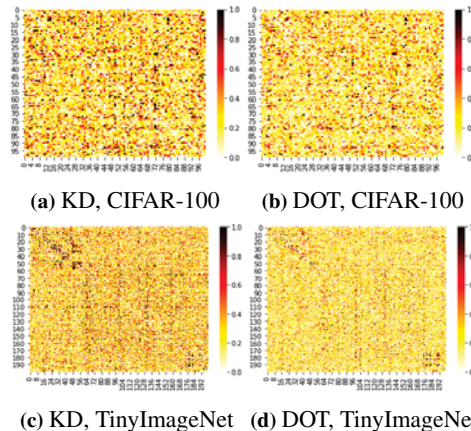


Figure 7: Difference of student and teacher logits. DOT leads to a significantly smaller difference (more similar prediction) than KD.

6. Conclusion

In this paper, we investigate the optimization property of knowledge distillation. We reveal a counter-intuitive phenomenon that introducing distillation loss limits the convergence of task loss, *i.e.*, a trade-off. We conjecture that the key to breaking the trade-off is sufficiently optimizing the distillation loss. To this end, we present a novel optimization method named Distillation-Oriented Trainer (DOT). Extensive experiments validate our motivation and the practical value of DOT. Visualizations show that DOT leads to surprisingly flat minima with both lower task and distillation losses. Additionally, we demonstrate that DOT improves the performance of popular distillation methods. We hope this work can provide valuable experiences for future research in the knowledge distillation community.

References

- [1] Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *KDD*, 2006. 2
- [2] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. In *NeurIPS*, 2017. 1
- [3] Jang Hyun Cho and Bharath Hariharan. On the efficacy of knowledge distillation. In *ICCV*, 2019. 2
- [4] Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, 2008. 1
- [5] Mark Craven and Jude Shavlik. Extracting tree-structured representations of trained networks. In *NeurIPS*, 1995. 2
- [6] Christian Darken and John Moody. Towards faster stochastic gradient search. In *NeurIPS*, 1991. 4
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv:1810.04805*, 2018. 1
- [8] Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. Sharp minima can generalize for deep nets. In *ICML*, 2017. 1, 2, 3
- [9] Felix Draxler, Kambis Veschgini, Manfred Salmhofer, and Fred Hamprecht. Essentially no barriers in neural network energy landscape. In *ICML*, 2018. 2
- [10] Tommaso Furlanello, Zachary Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. Born again neural networks. In *ICML*, 2018. 2
- [11] Ian J Goodfellow, Oriol Vinyals, and Andrew M Saxe. Qualitatively characterizing neural network optimization problems. *arXiv:1412.6544*, 2014. 2
- [12] Jianping Gou, Baosheng Yu, Stephen John Maybank, and Dacheng Tao. Knowledge distillation: A survey. *arXiv:2006.05525*, 2020. 1
- [13] Haowei He, Gao Huang, and Yang Yuan. Asymmetric valleys: Beyond sharp and flat local minima. In *NeurIPS*, 2019. 2, 3
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1
- [15] Byeongho Heo, Jeesoo Kim, Sangdoon Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *ICCV*, 2019. 1, 2
- [16] Byeongho Heo, Minsik Lee, Sangdoon Yun, and Jin Young Choi. Knowledge transfer via distillation of activation boundaries formed by hidden neurons. In *AAAI*, 2019. 2
- [17] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *arXiv:1503.02531*, 2015. 1, 2, 3, 6
- [18] Sepp Hochreiter and Jürgen Schmidhuber. Simplifying neural nets by discovering flat minima. In *NeurIPS*, 1994. 2, 3
- [19] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *UAI*, 2018. 2, 3
- [20] Guangda Ji and Zhanxing Zhu. Knowledge distillation in wide neural networks: Risk bound, data efficiency and imperfect teacher. In *NeurIPS*, 2020. 1, 2
- [21] Zijian Kang, Peizhen Zhang, Xiangyu Zhang, Jian Sun, and Nanning Zheng. Instance-conditional knowledge distillation for object detection. In *NeurIPS*, 2021. 1
- [22] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv:1609.04836*, 2016. 1, 2, 3
- [23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 3, 5
- [24] Jogendra Nath Kundu, Siddharth Seth, Anirudh Jamkhandi, Pradyumna YM, Varun Jampani, Anirban Chakraborty, and Venkatesh Babu R. Non-local latent relation distillation for self-adaptive 3d human pose estimation. In *NeurIPS*, 2021. 1
- [25] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *NeurIPS*, 2018. 1, 2
- [26] Jinyu Li, Rui Zhao, Jui-Ting Huang, and Yifan Gong. Learning small-size dnn with output-distribution-based criteria. In *Interspeech*, 2014. 2
- [27] Percy Liang, Hal Daumé III, and Dan Klein. Structure compilation: trading structure for features. In *ICML*, 2008. 2
- [28] David Lopez-Paz, Leon Bottou, Bernhard Scholkopf, and Vladimir Vapnik. Unifying distillation and privileged information. In *ICLR*, 2015. 1, 2
- [29] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet V2: Practical guidelines for efficient cnn architecture design. In *ECCV*, 2018. 5
- [30] Aditya Krishna Menon, Ankit Singh Rawat, Sashank J Reddi, Seungyeon Kim, and Sanjiv Kumar. A statistical perspective on distillation. In *ICML*, 2021. 1, 2
- [31] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *CVPR*, 2019. 2
- [32] Jeffrey Pennington and Yasaman Bahri. Geometry of neural network loss surfaces via random matrix theory. In *ICML*, 2017. 2
- [33] Mary Phuong and Christoph Lampert. Towards understanding knowledge distillation. In *ICML*, 2019. 1, 2
- [34] Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *Ussr computational mathematics and mathematical physics*, 1964. 4
- [35] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 1
- [36] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *ICLR*, 2015. 1, 2
- [37] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. ImageNet large scale visual recognition challenge. *IJCV*, 2015. 5
- [38] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobilenetV2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018. 5

- [39] Zhiqiang Shen, Zechun Liu, Dejie Xu, Zitian Chen, Kwang-Ting Cheng, and Marios Savvides. Is label smoothing truly incompatible with knowledge distillation: An empirical study. In *ICLR*, 2021. 1, 2
- [40] K. Simonyan and A Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 1
- [41] Samuel Stanton, Pavel Izmailov, Polina Kirichenko, Alexander A Alemi, and Andrew G Wilson. Does knowledge distillation really work? In *NeurIPS*, 2021. 1, 2
- [42] Jiayi Tang, Rakesh Shivanna, Zhe Zhao, Dong Lin, Anima Singh, Ed H Chi, and Sagar Jain. Understanding and improving knowledge distillation. *arXiv:2002.03532*, 2020. 1, 2
- [43] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *ICLR*, 2020. 1, 2, 5, 6, 7, 8
- [44] Lin Wang and Kuk-Jin Yoon. Knowledge distillation and student-teacher learning for visual intelligence: A review and new outlooks. *T-PAMI*, 2021. 1
- [45] Li Yuan, Francis EH Tay, Guilin Li, Tao Wang, and Jiashi Feng. Revisiting knowledge distillation via label smoothing regularization. In *CVPR*, 2020. 1, 2, 6
- [46] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2017. 1, 2
- [47] Yu Zhang and Qiang Yang. A survey on multi-task learning. *TKDE*, 2021. 4
- [48] Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *CVPR*, 2022. 2, 7, 8
- [49] Du Zhixing, Rui Zhang, Ming Chang, xishan zhang, Shaoli Liu, Tianshi Chen, and Yunji Chen. Distilling object detectors with feature richness. In *NeurIPS*, 2021. 1