

NetDistiller: Empowering Tiny Deep Learning via In-Situ Distillation

Shun Yao Zhang, *Rice University, Houston, TX, 77005, USA*

Yonggan Fu, *Georgia Institute of Technology, Atlanta, GA, 30332, USA*

Shang Wu, *Rice University, Houston, TX, 77005, USA*

Jyotikrishna Dass, *Rice University, Houston, TX, 77005, USA*

Haoran You, *Georgia Institute of Technology, Atlanta, GA, 30332, USA*

Yingyan (Celine) Lin, *Georgia Institute of Technology, Atlanta, GA, 30332, USA*

Abstract—Boosting the task accuracy of tiny neural networks (TNNs) has become a fundamental challenge for enabling the deployments of TNNs on edge devices which are constrained by strict limitations in terms of memory, computation, bandwidth, and power supply. To this end, we propose a framework called NetDistiller to boost the achievable accuracy of TNNs by treating them as sub-networks of a weight-sharing teacher constructed by expanding the number of channels of the TNN. Specifically, the target TNN model is jointly trained with the weight-sharing teacher model via (1) gradient surgery to tackle the gradient conflicts between them and (2) uncertainty-aware distillation to mitigate the overfitting of the teacher model. Extensive experiments across diverse tasks validate NetDistiller’s effectiveness in boosting TNNs’ achievable accuracy over state-of-the-art methods. Our code is available at <https://github.com/GATECH-EIC/NetDistiller>.

The recent record-breaking performance of neural networks (NNs) has spurred their increased application across various scientific and engineering disciplines. In parallel, it is projected that the worldwide number of Internet of Things (IoT) connected devices will reach 30.9 billion units by 2025 [1]. Deploying NN-powered intelligence on numerous IoT devices is of significant importance, as it enables the utilization of edge-collected data for various on-device intelligent functionalities that could potentially revolutionize human life. This tremendously growing demand has given rise to the field of tiny neural networks (TNNs) which have attracted substantially increasing attention. This is because TNNs enable small and inexpensive edge devices to work directly on local data at a lower power and computing cost, leading to both reduced latency and enhanced privacy as it alleviates or even eliminates the necessity of internet connectivity for sharing and centralizing the data on a

cloud server. However, the achievable task performance of TNNs remains unsatisfactory due to their limited model capacity. Hence, improving the task performance of TNNs has become a fundamental challenge for enabling their wide-scale adoption, which is highly desired in numerous real-world edge applications.

To tackle the aforementioned challenge and thereby unlock the potential of TNNs at the edge, there has been an increasing research effort towards boosting their achievable task performance. In particular, it has been shown that training TNNs is fundamentally different from training large NNs. For example, the authors of [2] identify that TNNs suffer from under-fitting due to their limited model capacity, in contrast to large NNs, which are prone to over-fitting. They also observe that data augmentation and regularization techniques, which enhance ImageNet accuracy for over-parameterized large NNs like ResNet50, have a detrimental effect on the accuracy of TNNs, such as MobileNetV2-Tiny [3], which is $174\times$ smaller than ResNet50.

Drawing inspiration from prior arts, we hypothesize that augmenting the model capacity (e.g., channels)

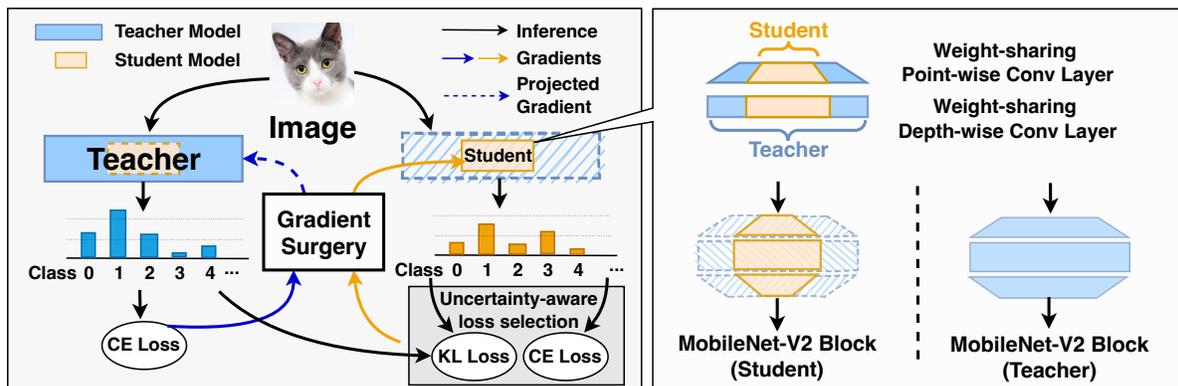


FIGURE 1. An overview of NetDistiller. The target TNN is a student model serving as a sub-network within a weight-sharing teacher model, constructed by expanding the number of channels of the target TNN. The teacher and student models are trained simultaneously while the teacher model is trained with the ground truth labels and the objective for training the student model is input-adaptively decided between an in-situ distillation mechanism and the ground truth labels based on its output uncertainty. To alleviate the gradient conflict issue observed during the training process, the teacher’s gradients are modified via gradient surgery to remove the conflicting components based on the student’s gradients.

during training enables TNNs to acquire additional knowledge, resulting in improved task accuracy. In vanilla knowledge distillation [4], the knowledge encoded by a large model is transferred to a smaller one by training the small model with the outputs and/or activations of the large model. As a result, the small model is able to achieve a higher accuracy by mimicking the behaviors of the large model. In this work, we advocate a new in-situ knowledge distillation scheme, which is orthogonal to the vanilla one, further boosting the achievable task performance of TNNs. We make the following contributions:

- We are the first to demonstrate that integrating a weight-sharing supernet with in-situ distillation can serve as an effective training recipe for boosting the achievable task performance of TNNs. Specifically, we propose a framework named NetDistiller, which incorporates the target TNN as a student model within a weight-sharing supernet that acts as a teacher model to boost the task performance of the trained TNNs without incurring any inference overhead.
- We identify that vanilla in-situ distillation can cause serious gradient conflicts between the supernet teacher and the sub-network student (i.e., the target TNN). Specifically, we find that up to 50% of the weight gradients in the student model have negative cosine similarities with those of the teacher model. This results in poor convergence when these gradients accumulate on their shared weights. Furthermore, vanilla in-situ distillation tends to induce overfitting in the teacher model, thus diminishing the effectiveness

of our in-situ distillation.

- To alleviate both the two issues identified above, NetDistiller proposes two solutions: (1) remove the conflicting gradients by projecting the teacher’s conflicting gradients onto the normal plane of the student’s gradients, and (2) integrate an uncertainty-aware distillation to dynamically select the student loss function between the Kullback–Leibler divergence and the cross-entropy loss based on the certainty of the student model’s output. These enhancements enable NetDistiller to unleash the promising effectiveness of in-situ distillation in more favorably training TNNs.
- We perform extensive evaluations and ablation studies to validate the effectiveness of our NetDistiller framework for boosting the achievable accuracy of TNNs when compared to the state-of-the-art (SOTA) method. For instance, we observe a 2.3% higher accuracy over NetAug [2], when training the MobileNet-V3-w0.35 model on the ImageNet dataset. We understand that NetDistiller has opened up a new perspective for boosting the achievable task performance of TNNs and enriched the field of knowledge distillation.

RELATED WORK

Efficient / Tiny Neural Networks. Significant progress has been made in designing efficient and mobile-friendly NNs. For example, MobileNets [5] utilize Depth-wise Separable Convolutions, which replace a standard convolution layer with a combination of depthwise convolution and pointwise convolution, demonstrating the potential for reducing computational costs while

maintaining task accuracy. In parallel to manually designed efficient networks and compression schemes, automated machine learning has been successfully used via neural architecture search [6], [7]. In contrast to the above techniques, our proposed framework aims to improve the inference accuracy of TNNs via in-situ distillation where the target TNN architecture is used as a sub-network student of a weight-sharing supernet teacher model.

Knowledge Distillation. Knowledge Distillation (KD) [8] refers to the idea of transferring knowledge acquired by a pre-trained and over-parameterized (teacher) model to a small (student) model that is more suitable for edge deployment. Specifically, the small model usually has insufficient capacity to learn a concise knowledge representation, and KD empowers the student model to learn the exact behavior of the teacher model by mimicking the teacher’s outputs at each level, i.e., soft labels. In this work, we advocate a new in-situ distillation scheme, namely, NetDistiller to enhance the task performance of TNNs. In contrast to vanilla KD where the teacher model is a large NN pre-trained on a different dataset and the student model is a separate smaller model, NetDistiller is an orthogonal approach that incorporates the TNN as a student sub-network within a weight-sharing supernet that acts as the teacher model for in-situ distillation.

Network Augmentation. The authors in [2] propose network augmentation, known as NetAug, to boost the accuracy of tiny deep learning by alleviating the under-fitting issue. Specifically, NetAug dynamically augments the network during training, incorporating the tiny model as a sub-model within a larger model for auxiliary supervision in addition to its independent functionality. In contrast, NetDistiller provides an alternative scheme to boosting the performance of TNNs via in-situ knowledge distillation. In particular, the TNN in NetDistiller acts as a sub-network (student) in a static weight-sharing supernet (teacher) constructed by expanding the channels of the target TNN. In the experimental results section, we provide a comparative study of the proposed NetDistiller with the SOTA scheme NetAug and find that NetDistiller outperforms NetAug, e.g., achieves 2.3% higher accuracy when training the MobileNet-V3-w0.35 model on the ImageNet dataset.

The NetDistiller FRAMEWORK

Our proposed NetDistiller is a training recipe for boosting the accuracy of TNNs by incorporating the target TNN as a student model (sub-network) within a weight-sharing supernet, which acts as a teacher model. Through in-situ distillation, NetDistiller distills and trans-

fers the knowledge from a supernet teacher model to the sub-network student model which is our target TNN. In this section, we first describe the construction of the weight-sharing supernet from the TNN followed by practical implementation of our in-situ distillation. Next, we describe techniques to resolve the gradient conflicts between the teacher and student models, as well as mitigate the over-fitting issue in the teacher model during the final training stage via uncertainty-aware distillation. Finally, we discuss the training and the inference overheads incurred by NetDistiller.

NetDistiller’s Enabler 0: Constructing the Weight-Sharing Teacher Model

NetDistiller expands the target TNN’s channels to construct a weight-sharing supernet that functions as the teacher model. Thus, the target TNN acts as a sub-network model. Both the student and teacher models share weights across all convolution layers while maintaining their respective Batch-Normalization layers, which account for different running statistics (i.e., means and variances) of their activation values. As a novel training recipe for augmenting the capacity of a target TNN to alleviate under-fitting issues and to boost accuracy, NetDistiller constructs a teacher model with $3\times$ channel numbers as the target TNN. Figure 1 depicts the above-described construction of the weight-sharing teacher model from the target TNN.

NetDistiller’s Enabler 1: In-Situ Distillation

The goal of our in-situ distillation is to stabilize the training of the supernet and improve the performance of sub-networks. In particular, TNNs are more likely to get stuck in local minimums due to insufficient capacity, which limits their performance compared to over-parameterized large NNs [2]. To tackle this, NetDistiller integrates the target TNN as a sub-network student model within a weight-sharing supernet teacher model constructed by expanding the channels of the target NN as demonstrated in Figure 1. To the best of our knowledge, NetDistiller is the first to demonstrate that applying in-situ distillation to a weight-sharing supernet [9] can serve as an effective training recipe for boosting the achievable task performance of TNNs.

Specifically, in-situ distillation leverages the ‘soft labels’ predicted by the supernet as the supervision signals to the sub-network student model during each training iteration while using ground truth labels for the teacher model. Formally, at training iteration n , the supernet parameter W is updated by

$$W^n \leftarrow W^{n-1} + \eta g(W^{n-1}),$$

where η is the learning rate, and:

$$g(W^{n-1}) = \nabla_W \left(\mathcal{L}_{\mathcal{D}}(W) + \mathcal{L}_{stu}([W, W_{stu}]; W^{n-1}) \right) \Big|_{W=W^{n-1}} \quad (1)$$

Here, $\mathcal{L}_{\mathcal{D}}$ is the cross-entropy loss of the supernet teacher on a training dataset \mathcal{D} , W and W_{stu} denote the teacher and student models, respectively, and \mathcal{L}_{stu} is the student loss modulated by uncertainty-aware distillation (introduced in the Enabler 3 section). Additionally, the distillation process in NetDistiller is single-shot, i.e., it is implemented in-situ during training without additional computation and memory cost, unlike two-step vanilla KD where a large model has to be trained first.

NetDistiller’s Enabler 2: Gradient Surgery for Resolving Gradient Conflicts

Considering that the gradients from both the student and the teacher models accumulate on the shared weights, we identify that vanilla in-situ distillation may cause serious gradient conflicts between the supernet teacher and sub-network student (target TNN). Specifically, we find that up to 50% of the student model gradients have a negative cosine similarity with those of the teacher model. Inspired by the PCGrad [10] which performs gradient surgery for multi-task learning, NetDistiller tackles this by projecting the conflicting teacher gradients to the normal plane of student gradients, thereby removing the conflicting components in the teacher gradients and improving the performance of TNNs. Specifically, let ∇l_{stu} and ∇l_{tea} denote the gradients of the student and the teacher models, respectively. We define ϕ as the angle between the above two gradients and g as the final gradient for updating the weights. In order to guarantee the student model training, we project the conflicting teacher’s gradient, $proj(\nabla l_{tea})$, when the cosine similarity $\cos(\phi) = \frac{\nabla l_{stu} \cdot \nabla l_{tea}}{\|\nabla l_{stu}\| \|\nabla l_{tea}\|}$ is negative, which is formulated as follows:

$$g = \nabla l_{stu} + proj(\nabla l_{tea}), \text{ where}$$

$$proj(\nabla l_{tea}) = \begin{cases} \nabla l_{tea} - \frac{\nabla l_{tea}^T \nabla l_{stu}}{\|\nabla l_{stu}\|^2} \nabla l_{stu}, & \text{if } \cos(\phi) < 0 \\ \nabla l_{tea}, & \text{otherwise} \end{cases} \quad (2)$$

NetDistiller’s Enabler 3: Uncertainty-aware Distillation

Since the student and the weight-sharing teacher models are jointly trained from scratch and the over-parameterized teacher model converges faster than the student one, we observe that the supernet teacher

model suffers from over-fitting at the final training stage. In addition, [11] advocates that large models have the largest improvement on samples where the small model is most uncertain. As for certain examples, even those where the small model is not particularly accurate, large models are often unable to improve. Based on these insights, we hypothesize that the teacher model is not always a good teacher during the whole training.

In light of this, we propose a technique called uncertainty-aware distillation (UD) to dynamically select the student loss functions between the Kullback–Leibler (KL) divergence and cross-entropy losses based on the certainty of the student model output (see Figure 1). Specifically, we measure the uncertainty via the entropy of the student outputs. When the entropy of the student output is high (i.e., uncertain), the student is distilled by the weight-sharing teacher via the KL divergence loss, otherwise, the student is trained by the ground truth label via the cross-entropy loss. We formulate this process as follows:

$$\mathcal{L}_{stu} = \begin{cases} \text{KL}(W_{stu}(x), W(x)), & \text{uncertainty} \geq T \\ \text{CE}(W_{stu}(x), y), & \text{otherwise} \end{cases} \quad (3)$$

Here, T denotes the uncertainty threshold; \mathcal{L}_{stu} denotes the student model loss; $\text{KL}()$ and $\text{CE}()$ denote the KL divergence loss and the cross-entropy loss, respectively; x and y denote the input data and the ground truth labels; and uncertainty denotes the entropy of the student model outputs $W_{stu}(x)$.

Analysis of Training and Inference Overhead

In contrast to the two-step distillation process in vanilla KD, NetDistiller performs **one-shot** in-situ distillation of knowledge from the supernet teacher to the sub-network student model without any additional computation and memory cost. Similar to that of NetAug, NetDistiller introduces **zero** extra inference overhead because only the target TNN is used during inference, enabling the deployment of TNNs feasible on resource-constrained edge devices. Despite expanding the target TNN model by $3\times$, we observe a mere 20% increase in the training time of NetDistiller to that of vanilla TNNs.

EXPERIMENTAL RESULTS

Experiment Setup

Models. We benchmark NetDistiller with SOTA TNNs training methods, e.g., NetAug [2], and KD [4], on five commonly adopted TNNs [3], [2], including MobileNet-V2-Tiny, MobileNet-V2 (w0.35 and w1.0), MobileNet-V3, MCUNet (256kb-1mb), and ProxylessNAS (w0.35, w1.0). In particular, w0.35 indicates the models have

TABLE 1. Benchmark NetDistiller with SOTA methods for training TNNs. r160: The input image resolution is 160×160 . w0.35: The model has $0.35 \times$ number of channels than the vanilla one.

Model	MobileNet-V2-Tiny r144	MCUNet r176	MobileNet-V3, r160 w0.35	ProxylessNAS, r160 w0.35	ProxylessNAS, r160 w1.0	MobileNet-V2, r160 w0.35	MobileNet-V2, r160 w1.0
Params	0.75M	0.74M	2.2M	1.8M	4.1M	1.7M	3.5M
MACs	23.5M	81.8M	19.6M	35.7M	164.1M	30.9M	154.1M
Baseline	51.7%	61.5%	58.1%	59.1%	71.2%	56.3%	69.7%
NetAug [2]	53.3%	62.7%	60.3%	60.8%	71.9%	57.8%	70.6%
In-situ	54.1%	62.7%	62.1%	60.7%	71.2%	58.5%	71.2%
In-situ + PCGrad [10]	54.5%	63.4%	62.3%	61.3%	72.5%	59.0%	72.0%
NetDistiller (ours)	54.8%	64.2%	62.6%	61.5%	72.8%	59.3%	72.6%

TABLE 2. Ablation study of channel expansion rates on MobileNet-V2-w0.35 (MBV2-w0.35) and MobileNet-V3-w0.35 (MBV3-w0.35). Different teacher sizes in the first row indicate the channel expansion rates. Considering the limited improvement (0.2%) between $\times 4$ and $\times 3$ teachers on MobileNet-V2-w0.35 model and the training efficiency, the teacher with $\times 3$ size is selected in NetDistiller.

Teacher Size	Baseline	$\times 2$	$\times 3$	$\times 4$	$\times 5$
MobileNet-V2-0.35	56.3%	58.0%	58.5%	58.7%	58.3%
MobileNet-V3-0.35	58.1%	61.3%	62.1%	61.8%	61.8%

0.35 times of channels over the vanilla one (i.e., w1.0). Following the model definitions in [2], the channels of w0.35 models are round to products of 8.

Datasets. Following [2], we consider the ImageNet dataset with an input resolution of r144, r160, and r176 for different target TNNs. In the external knowledge distillation experiments, the input resolution for the external teacher is set to r224 to match its pre-training configuration. The object detection experiments are trained on the PASCAL VOC 2007+2012 datasets and evaluated on the PASCAL VOC 2007 validation set with an input resolution of r416.

Evaluation Metrics. We evaluate NetDistiller and the baseline methods in terms of the top-1 accuracy on ImageNet and the average precision at IoU=0.5 (AP50) for the object detection on PASCAL VOC.

Training Setting. Following the training setting in [2], we train TNNs for 180 epochs using an SGD optimizer with a momentum of 0.9 and an initial learning rate of 0.4 with a cosine learning rate scheduler. We adopt a learning rate warm-up for 5 epochs and the gradients are clipped to 1.0 during the whole training period. The label smoothing technique with a factor of 0.1 is adopted when using the ground truth label. For the NetDistiller with uncertainty-aware distillation, we use the same training recipe but increase the training epochs to 360. The uncertainty threshold T is set to 3.75 based on the empirical observation of our ablation study. All the ImageNet experiments are run on 8 GPUs with a batch size of 1024. As a recent paper [2] discovered

TABLE 3. Ablation study of gradient surgery on MobileNet-V2-w0.35 and MobileNet-V3-w0.35 models for 360 epochs. We disable gradient surgery and calculate the cosine-similarity between the two gradients (teacher’s and student’s) of each convolutional layer. The percentage values shown under different epochs reflect the average ratio of the number of layers with negative cosine-similarity (gradient conflicts) w.r.t. the total number of layers in the model.

Epoch	1	90	180	270	360
MobileNet-V2-w0.35	51.5%	40.1%	37.4%	39.4%	38.2%
MobileNet-V3-w0.35	50.1%	45.2%	34.7%	38.5%	37.4%

TABLE 4. Ablation study of different uncertainty-aware distillation thresholds on MobileNet-V2-w0.35 and MobileNet-V2-w1.0. The first row is the thresholds. The uncertainty-aware distillation distills the student model if its output entropy (uncertainty) is higher than the threshold and trains the student model with ground truth labels otherwise.

Model	Uncertainty Threshold		
	2.5	3.75	5.0
MobileNet-V2-w0.35	59.1%	59.3%	58.9%
MobileNet-V2-w1.0	71.9%	72.6%	71.2%

that data augmentation and regularization could be harmful to TNN training, we only utilize standard data augmentations (e.g. random flip, random crop) and disable regularization methods like dropout and drop path. For transfer learning on the object detection task, MobileNet-V2-w0.35 and MobileNet-V3-w0.35 models are connected with a YOLO-v4 head. All the object detection experiments are trained via an SGD optimizer with a momentum of 0.9 and an initial learning rate of $1e-4$ decayed by a cosine learning rate scheduler for 100 epochs with a batch size of 8.

Benchmark with SOTA TNN Training Methods

As shown in Table 1, our observations are as follows: (1) NetDistiller improves up to 4.5% accuracy as compared to all baselines and gains 2.3% higher accuracy than the most competitive baseline NetAug on MobileNet-V3-

TABLE 5. Ablation studies of combining NetDistiller and External knowledge distillation. KD: Distill the target TNN with an external teacher (ImageNet pre-trained ResNet-50). NetDistiller w/o UD: Uncertainty-aware distillation is turned off in the external KD experiments. NetDistiller+KD: The external teacher distills both NetDistiller teacher and student models.

Model	Baseline	KD	NetDistiller w/o UD	NetDistiller+KD
MobileNet-V2-Tiny, r144	51.7%	53.7% (+2.0%)	55.5% (+3.8%)	56.1% (+4.4%)
MobileNet-V2-w0.35, r160	56.3%	58.4% (+2.1%)	59.0% (+2.7%)	59.5% (+3.2%)
MobileNet-V3-w0.35, r160	58.1%	61.6% (+3.5%)	62.3% (+4.2%)	62.5% (+4.4%)
ProxylessNAS-w0.35, r160	59.1%	60.8% (+1.7%)	61.3% (+2.2%)	61.9% (+2.8%)

w0.35. (2) In-situ distillation contributes the most: TNNs trained solely through in-situ distillation can either match or surpass the SOTA TNNs training method, NetAug [2]. (3) Among all the models, TNNs gain an accuracy boost of approximately 0.5% after the introduction of PCGrad [10] to mitigate gradient conflicts. (4) Our proposed NetDistiller, which combines the in-situ distillation with PCGrad and uncertainty-aware distillation, proves to be an effective training approach for enhancing the attainable task performance of TNNs.

Ablation Studies of NetDistiller

Channel Expansion Rates of the Teacher Model.

Since NetDistiller expands the channels of TNNs to create a weight-sharing supernet as the teacher model, the channel expansion rate of the teacher model plays a crucial role in determining the effectiveness of the in-situ distillation mechanism. This is particularly important considering that a thin teacher model may have limited capacity, while an excessively wide teacher may not effectively transfer its information. In order to identify the optimal channel expansion rate, we evaluate NetDistiller with $\times 2$, $\times 3$, $\times 4$, and $\times 5$ channel expansion rates on top of two TNNs, MobileNet-V2-w0.35 and MobileNet-V3-w0.35, respectively. As shown in Table 2, we observe that: (1) all four teacher models enhance the accuracy of TNNs, demonstrating the overall efficacy of NetDistiller, and (2) MobileNet-V2-w0.35 and MobileNet-V3-w0.35 achieve the highest accuracy when the channel expansion rates are $\times 4$ and $\times 3$, respectively. To minimize the training overhead caused by the expanded teacher model, we by default set the channel expansion rate to $\times 3$ in NetDistiller.

Quantify the Gradient Conflicts. Because of the weight-sharing mechanism, the joint training of the student and teacher models leads to the accumulation of gradients on the same weights, inevitably resulting in gradient conflicts. To verify the occurrence of gradient conflicts at different training stages, we calculate the ratio of layers with negative cosine similarity, averaged over the validation set, in relation to the total number of layers in the model throughout the training process.

The results in Table 3 demonstrate that the teacher and student models indeed experience gradient conflicts, with as many as 51.5% of layers exhibiting conflicting gradients. To address this issue, the adoption of gradient surgery leads to an accuracy improvement of 0.5% and 0.2% for MobileNet-V2-w0.35 and MobileNet-V3-w0.35, respectively, as shown in Table 1.

Threshold of the Uncertainty-aware Distillation.

Our proposed uncertainty-aware distillation mechanism dynamically selects the objective for the student model, opting for either in-situ distillation or cross-entropy based on the uncertainty of the student model’s outputs. To decide the uncertainty threshold, we validate MobileNet-V2-w0.35 and MobileNet-V2-w1.0 models with uncertainty thresholds of 5.0, 3.75, and 2.5, considering the entropy of ImageNet models is in the range of [1.5, 10] when adopting a label smoothing factor of 0.1. As shown in Table 4, we observe that both models achieve their highest accuracy when the uncertainty threshold is set to 3.75, resulting in accuracy improvements of 0.3% and 0.6%, respectively, as compared to NetDistiller without uncertainty-aware distillation. Consequently, we adopt a default uncertainty threshold of 3.75 when enabling uncertainty-aware distillation, without any additional overhead.

Complementary Knowledge Distillation. A natural baseline for evaluating NetDistiller is standard knowledge distillation. Recent observations have suggested that significant gaps between the teacher and student models can lead to suboptimal knowledge distillation performance [12]. Therefore, we hypothesize that (a) our proposed in-situ distillation serves as a more effective mechanism for enabling TNN training when compared to knowledge distillation, and (b) knowledge distillation is complementary to our method and can be applied concurrently. To validate these hypotheses, we distill the knowledge from an ImageNet pre-trained ResNet-50 for both NetDistiller’s teacher and student models. This process is referred to as external distillation to distinguish it from our in-situ distillation. We then benchmark this approach against (1) vanilla NetDistiller and (2) standard knowledge distillation. As shown in

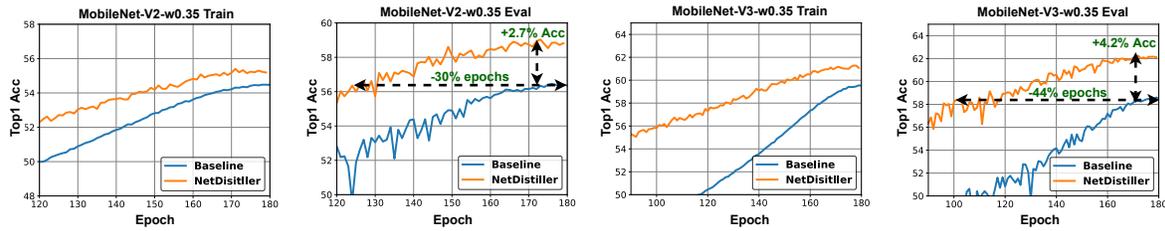


FIGURE 2. Visualizing the training process of NetDistiller and the baselines for MobileNet-V2-w0.35 and MobileNet-V3-w0.35 models. It reveals that NetDistiller significantly enhances both the training and evaluation accuracy of the TNNs. This highlights the substantial potential of NetDistiller in strengthening TNNs and mitigating the underfitting problem.

Table 5, we can observe that (1) vanilla NetDistiller can outperform standard knowledge distillation, e.g., an accuracy improvement of 3.8% and 4.2% on MobileNet-V2-Tiny and MobileNet-V3-w0.35 respectively, thus confirming our hypothesis (a); and (2) knowledge distillation is orthogonal to NetDistiller, as applying knowledge distillation on top of NetDistiller results in 4.4% accuracy improvements on MobileNet-V2-Tiny and MobileNet-V3-w0.35, respectively, validating our hypothesis (b).

Visualization of Training Trajectories

We present the training curves for MobileNet-V2-w0.35 and MobileNet-V3-w0.35 throughout the 180 epochs of training using NetDistiller, as compared to the corresponding standard training baselines, in Figure 2. Notably, we observe that both the training and evaluation accuracy of NetDistiller consistently outperform the corresponding baselines at the same training epoch. For instance, NetDistiller achieves a 2.7% accuracy improvement for MobileNet-V2-w0.35 and an impressive 4.2% accuracy improvement for MobileNet-V3-w0.35 as compared to the baselines. Furthermore, to attain a comparable accuracy, NetDistiller requires fewer training epochs, resulting in a 44% reduction in training epochs, as indicated in Figure 2.

Transfer Learning Study on Object Detection

To evaluate the generality of the representations learned by NetDistiller, we transfer NetDistiller’s trained MobileNet-V2-w0.35 and MobileNet-V3-w0.35 to a downstream object detection task. We then compare their performance against standard pre-trained models, with or without knowledge distillation, on ImageNet. Specifically, we replace the final pooling and linear layers in MobileNet-V2-w0.35 and MobileNet-V3-w0.35 with the YOLO-v4 object detection head. As shown in Table 6, NetDistiller consistently enjoys a better transferability with a higher Average Precision (AP) of

TABLE 6. Transfer learning on object detection tasks using PASCAL VOC 2007+2012 datasets and MobileNet-V2-w0.35 (MBV2) or MobileNet-V3-w0.35 (MBV3) models. Both of the two models are connected with the YOLO-v4 detection head. The model is initialized with NetDistiller pretrained weights on ImageNet. We report the Average Precision at IoU=0.5 (AP50).

Model	Baseline AP50	KD AP50	NetDistiller AP50
MobileNet-v2-w0.35	60.4%	61.1%	62.3%
MobileNet-v3-w0.35	63.6%	62.8%	65.2%

1.9% / 1.6% on MobileNet-V2-w0.35 / MobileNet-V3-w0.35 as compared to the standard training baselines. Note that although pre-trained features on classification tasks may not be necessarily useful for downstream tasks [2], which is also echoed with our results of KD pre-trained MobileNet-V3-w0.35 in Table 6, NetDistiller manages to improve the achievable AP by up to 1.9% on both models. This emphasizes the broad applicability of our method across diverse tasks and datasets.

CONCLUSION

Enhancing the task accuracy of TNNs is a critical challenge in enabling their deployment on resource-constrained IoT devices. Our proposed framework, NetDistiller, addresses this challenge by considering TNNs as sub-networks of a weight-sharing teacher model, achieved by expanding the number of channels in the TNN. By incorporating gradient surgery to handle gradient conflicts and uncertainty-aware distillation to alleviate teacher model overfitting, NetDistiller significantly improves the achievable accuracy of TNNs. Extensive experiments spanning various tasks demonstrate the superior effectiveness of NetDistiller compared to SOTA TNN training schemes. This advancement marks a significant step towards realizing the full potential of TNNs in real-world IoT applications.

Acknowledgement

The work is supported by the National Science Foundation (NSF) through the CCF program (Award number: 2211815) and by CoCoSys, one of the seven centers in JUMP 2.0, a Semiconductor Research Corporation (SRC) program sponsored by DARPA.

REFERENCES

1. Lionel Sujay Vailshery. Iot and non-iot connections worldwide 2010-2025. <https://www.statista.com/statistics/1101442/iot-number-of-connected-devices-worldwide/>. Accessed: 2021-03-08.
2. Han Cai, Chuang Gan, Ji Lin, and Song Han. Network augmentation for tiny deep learning. In *International Conference on Learning Representations*, 2022.
3. Ji Lin, Wei-Ming Chen, Yujun Lin, Chuang Gan, Song Han, et al. Mucnet: Tiny deep learning on iot devices. *Advances in Neural Information Processing Systems*, 33:11711–11722, 2020.
4. Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015.
5. Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
6. Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.
7. Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *Proceedings of the IEEE/CVF Conference on CVPR*, pages 10734–10742, 2019.
8. Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in NIPS*, volume 27. Curran Associates, Inc., 2014.
9. Dilin Wang, Chengyue Gong, Meng Li, Qiang Liu, and Vikas Chandra. Alphanet: Improved training of supernet with alpha-divergence. In *International Conference on Machine Learning*, pages 10760–10771. PMLR, 2021.
10. Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. *Advances in NIPS*, 33:5824–5836, 2020.
11. Taman Narayan, Heinrich Jiang, Sen Zhao, and Sanjiv Kumar. Predicting on the edge: Identifying where a larger model does better. *arXiv preprint arXiv:2202.07652*, 2022.
12. Jang Hyun Cho and Bharath Hariharan. On the efficacy of knowledge distillation. In *Proceedings of the IEEE/CVF ICCV*, pages 4794–4802, 2019.

Shunyao Zhang is a Ph.D. student at Rice University, Houston, USA. He received his master's degree in Electrical and Computer Engineering from Carnegie Mellon University, Pittsburgh, USA. His research interests are tiny ML and adversarial robustness. Contact him at sz74@rice.edu.

Yonggan Fu is a Ph.D. student at Georgia Institute of Technology. Before that, he obtained his Bachelor's degree from the School of The Gifted Young at the University of Science and Technology of China. His research focus and passion is to develop efficient and robust AI algorithms and co-design the corresponding hardware accelerators toward a triple-win in accuracy, efficiency, and robustness. Contact him at yfu314@gatech.edu.

Shang Wu is a master's student at RICE University where he majored in Electrical and Computer Engineering. He received his bachelor's degree in computer science at George Washington University. His research interests are in efficient ML, robust ML and generative AI. Contact him at sw99@rice.edu.

Jyotikrishna Dass is a Research Scientist at Rice University and manages the activities at Rice Data to Knowledge Program. His research interests are in distributed and parallel machine-learning systems for efficient edge computing. Previously, he was a postdoc at Dr. Yingyan Lin's lab. Dr. Dass received his Ph.D. in Computer Engineering from Texas A&M University. Contact: jdass@rice.edu.

Haoran You is currently a Ph.D. student in the CS Department of Georgia Institute of Technology. He received his bachelor's degree in the advanced class at Huazhong University of Science and Technology and his master's degree at Rice University. He is pursuing his doctoral degree in machine learning and computer architecture realm. His research interests include but are not limited to resource-constrained machine learning, computer vision, deep learning, and algorithm/accelerator co-design. Contact: haoran.you@gatech.edu.

Yingyan (Celine) Lin is currently an Associate Professor in the School of Computer Science and a member of the Machine Learning Center at the Georgia Institute of Technology. She leads the Efficient and Intelligent Computing (EIC) Lab. Her research focuses on developing efficient machine learning techniques via cross-layer innovations, spanning from efficient artificial intelligence (AI) algorithms to AI hardware accelerators and AI chip design, and aims to foster green AI and ubiquitous AI-powered intelligence. She received her Ph.D. degree in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign in 2017. Prof. Lin has received the NSF CAREER Award, IBM Faculty Award, Facebook Faculty Research Award, and the ACM SIGDA Outstanding Young Faculty Award. She has served on the Technical Program committees for various conferences including DAC, ICCAD, MLSys, MICRO, and NeurIPS. She is currently an Associate Editor for the IEEE Transactions on Circuits and Systems II: Express Briefs. Contact her at celine.lin@gatech.edu.