

Adaptive Distillation: Aggregating Knowledge from Multiple Paths for Efficient Distillation

Sumanth Chennupati
schennupati@wyze.com

Mohammad Mahdi Kamani
mkamani@wyze.com

Zhongwei Cheng
zchen@wyze.com

Lin Chen
lchen@wyze.com

Wyze Labs, Inc.
Kirkland, WA, USA

Abstract

Knowledge Distillation is becoming one of the primary trends among neural network compression algorithms to improve the generalization performance of a smaller student model with guidance from a larger teacher model. This momentous rise in applications of knowledge distillation is accompanied by the introduction of numerous algorithms for distilling the knowledge such as soft targets and hint layers. Despite this advancement in different techniques for distilling the knowledge, the aggregation of different paths for distillation has not been studied comprehensively. This is of particular significance, not only because different paths have different importance, but also due to the fact that some paths might have negative effects on the generalization performance of the student model. Hence, we need to adaptively adjust the importance of each path to maximize the impact of distillation on the student model. In this paper, we explore different approaches for aggregating these different paths and introduce our proposed adaptive approach based on multitask learning methods. We empirically demonstrate the effectiveness of our proposed approach over other baselines in the applications of knowledge distillation for classification, semantic segmentation, and object detection tasks.

1 Introduction

The promising advancements of deep learning models in various AI tasks are dominantly depending on their large and complex model structures, which grants them boosted generalization capabilities on test data. However, the benefits of these achievements are limited in resource-constrained systems such as mobile devices, low power robots, etc. Different model compression algorithms [4, 17, 18, 23, 25, 30, 31] have been proposed to reduce the complexity of such larger models for these systems. Among different compression algorithms, distilling the knowledge from a larger model to a smaller one has shown to be highly

effective and advantageous [22, 48, 57]. Beyond compression capabilities, knowledge distillation has also been a primary motive for different techniques that transfer the knowledge between models such as domain adaptation [36, 47].

With the upsurge in the applications of knowledge distillation in different domains, various methods were introduced for distilling the knowledge from a teacher to a student. These approaches are applied to different parts (*i.e.* distillation paths) of the models such as output logits [22] and hidden layers’ feature maps [59]. Although most of these distillation paths could improve the generalization performance of the student model, a naive combination of these paths might have negative effects on the performance of the student model. For instance, when we use knowledge distillation for an object detector, we can distill the knowledge from their backbone feature maps [75] or its bounding box generator [77]. As it can be inferred from Table 1, naively combining backbone features’ distillation paths with those from the bounding box generator can degrade the performance when compared to the bounding box generator alone. This phenomenon indicates the necessity for a systematic approach towards the aggregation of different paths to gain the most from the knowledge distillation process.

In this paper, we propose an adaptive approach to learn the importance of each path during the distillation and training process. This approach is inspired by multitask learning methods [5], where we consider each path as a separate task that we want to optimize the model for. Using our adaptive approach, we can mitigate the negative effects mentioned above and benefit the most while aggregating from multiple paths of distillation. As it can be seen in Table 1, our proposed Adaptive approach can aggregate these two paths and surpasses both of them in improving the performance of the student model. In addition to our adaptive approach, we propose another baseline method based on multiobjective optimization, where it reduces the aggregation problem to a multi-criteria optimization and intends to find its Pareto stationary solutions. This approach seems to be more effective than naive aggregation methods, but cannot outperform our proposed adaptive distillation.

The main contributions of this paper can be summarized as follows:

- We propose an adaptive distillation approach, inspired by multitask learning methods for efficient aggregation of different distillation paths. We provide a general optimization formulation for this problem and reduce different methods using this formulation.
- We introduce multiobjective optimization for this problem as a baseline approach, which can be more effective than naive aggregation approaches.
- We conduct extensive comparison between our approach and other baseline methods in different tasks of image classification, semantic segmentation, and object detection.

	Path(s)		Measures	
	Feats	Box	mAP	mAP ₅₀
Student			38.1	59.9
Single	✓		39.8	62.3
Single		✓	39.6	62.4
Hand-tuned	✓	✓	39.7	61.4
Adaptive	✓	✓	40.2	62.9
Teacher			42.9	65.7

Table 1: Comparison between a hand-tuned approach versus an adaptive approach on aggregating distillation paths for an object detection task on Cityscapes [9]. Please refer to Section 4 for more details.

2 Related Work

Knowledge distillation has become one of the prevailing approaches in model compression to improve the generalization performance of a smaller student model using the knowledge from a richer teacher model. In this section, we provide an overview of proposals in this domain, as well as methods we utilize to aggregate different paths of knowledge distillation.

Knowledge Distillation The concept of knowledge distillation for neural networks was first introduced by Hinton et al. [22] to distill the knowledge from a teacher to a student model by minimizing the distance between their soft targets. This idea has been expanded to the hidden layers' feature maps by the introduction of Hint layers by Romero et al. [59]. The idea of knowledge distillation and its variants has extensively employed in various problems such as compression [2, 6, 22, 54, 57, 66], knowledge transfer [58, 67, 73], and federated learning [16, 41, 62]. It has also been explored in different tasks such as image classification [22, 53, 72], object detection [6, 10, 66, 75], semantic segmentation [20, 51, 70], and graph neural networks [38, 49] to name but a few. More detailed discussions regarding knowledge distillation approaches and state-of-the-art methods in this domain can be found in [14, 65].

Ensemble of Teachers The problem of aggregating different knowledge paths from multiple teachers has been the primary topic of several studies. Most studies investigate how to effectively combine the ensemble of teachers' outputs (logits) for distillation to a student model using weighted averaging [1, 13, 37]. Different approaches have been proposed for finding these weights such as the teachers' confidence score [69] or the similarity between two models' inputs [74]. Zhang et al. [76] introduce the reverse problem, where they can improve the teacher with an ensemble of student models using symmetrical KL divergence. Piao et al. [56] introduce A2dele to combine depth and RGB networks using a confidence-based weights for distillation paths. In some other studies, the problem of an ensemble of teachers on hint layers' feature maps has been studied [45, 55]. However, the number of these proposals is limited due to the challenging nature of this problem because of the non-aligning size of feature maps in different teachers. As it is mentioned by [65], finding a systematic approach to calculate the degree of efficiency for each distillation path in this problem remains open and challenging. To the best of our knowledge this is one of the first attempts to adaptively combine distillation paths with different nature in a training procedure.

Multiobjective Optimization An efficient way to aggregate different distillation paths is to consider the problem as multiobjective optimization, and hence, benefit from approaches proposed in this domain. There are various methods to solve a multiobjective optimization, however, for the sake of efficiency of the approach, we will use a first-order gradient-based method like the ones proposed in [11, 50, 61]. Using these approaches, we can converge to the Pareto stationary of the problem, where no other solution can dominate that solution.

Multitask Learning Another way of aggregating knowledge from different distillation paths is to adjust the importance of each path by scaling their losses. Early works in multitask learning (MTL) [5, 8, 33, 64] use a weighted arithmetic sum of individual task losses. These weights are hand-tuned or manually searched and remain static throughout the training process. Later, Liu et al. [46] and Guo et al. [15] propose to adjust these weights by inspecting the change in loss or difficulty of each task as the training progresses. GradNorm [7] proposes to normalize gradients from different losses to a common scale during backpropagation. Kendall et al. [29] and Leang et al. [39] propose to consider different task weights as parameters and learn them using backpropagation.

3 Methods

Distilling knowledge from a teacher to a student can take place at any stage of a model from initial feature maps to final logits of the output. Not all of them have the same effect on boosting the generalization performance of the student model. In fact, in practice, it can be seen that some of these distillation paths might hurt the generalization of the student models. Hence, it is of paramount importance to take these effects into consideration when updating the parameters based on each of these paths. In this section, we first introduce the problem formulation. Next, using this formulation, we provide few baseline approaches on how to aggregate different paths of distillation during the training procedure.

3.1 Problem Formulation

In order to better formulate the multiple paths distillation problem at hand, we first describe the main learning task. Universally, in many forms of supervised learning tasks, the primary objective is to find an optimal prediction mapping, given a training dataset \mathcal{T} , with N training samples. The mapping is between the input feature space \mathcal{X} and the target space \mathcal{Y} , whether it is a classification, semantic segmentation, or object detection task. In this case, each sample is presented with the intersection of these two spaces denoted by $(\mathbf{x}^{(i)}, y^{(i)}) \in \mathcal{X} \times \mathcal{Y}, i \in [N]$. Deep neural networks aim at representing this mapping using an M -layer neural network model, where each layer l is represented by a parameter set of $\mathbf{w}_l \in \mathbb{R}^{d_l}$ and applies the transformation of $f_l(\cdot; \mathbf{w}_l)$ on its input. The set of all parameters of the network is denoted by $\mathbf{w} = \{\mathbf{w}_1, \dots, \mathbf{w}_M\} \in \mathbb{R}^d$, where $d = \sum_{i \in [M]} d_i$. Thus the main objective of this supervised task is to minimize the empirical risk on the training data, which is:

$$\mathcal{L}(\mathcal{T}; \mathbf{w}) = \frac{1}{N} \sum_{i \in [N]} \ell(\mathbf{x}_i, y_i; \mathbf{w}), \quad (1)$$

where $\ell(\cdot, \cdot; \cdot)$ is the loss on each sample such as cross entropy loss. In knowledge distillation frameworks the ultimate goal is for the student to imitate the teacher’s output features in different layers. For instance, in the primary form of the knowledge distillation using soft targets [22], these outputs are the soft logits of the two models. Whereas in hint layers [59] and attention transfer [34], these features are the middle layers’ outputs. It should be noted that each distillation path only affects a subset of the parameters in the student model (unless for the soft target or equivalent, where all the parameters of the student model are affected). Hence, the general form of knowledge distillation loss for each path between the j -th layer on the student model and the k -th layer on the teacher model can be formulated as:

$$\mathcal{L}_{KD}^{jk}(\mathcal{X}; \mathbf{w}^S, \mathbf{w}^T) = \frac{1}{N} \sum_{i \in [N]} \ell_{KD}^{jk}(\mathbf{x}_i; \mathbf{w}_{\leq j}^S, \mathbf{w}_{\leq k}^T), \quad (2)$$

where \mathbf{w}^S and \mathbf{w}^T are the student and teacher model parameters, respectively. The set of parameters for the layers up to layer j is denoted by $\mathbf{w}_{\leq j}$. The loss for each sample is calculated based on the loss function $\ell_{KD}^{jk}(\cdot, \cdot; \cdot)$. For instance, if the path is the soft target, the loss is the KL-divergence between two soft logits, and if it is a hint layer, the loss could be a simple euclidean distance [59]. For hint layers an *adaptation* layer \mathbf{w}_A might be necessary to match the spatial or channel size of the feature maps on the student model to that of the teacher model. The parameters of this layer will be tuned using the distillation loss defined in Eq (2). Thus, if we consider K paths for distillation with their own defined empirical loss function as in Eq (2), we can create a distillation loss vector denoted by $\mathbf{f}_{KD}(\mathcal{X}; \mathbf{w}^S, \mathbf{w}^T) \in \mathbb{R}^K$. Then, the overall optimization can be written as:




Figure 1: Schematic of different aggregation approaches for knowledge distillation. (a) The single distillation approach. (b) Hyperparameter optimization approach to find the optimal weight for each path before the training. (c) Multiobjective optimization to find the descent direction for multiple paths at every iteration. (d) The proposed adaptive distillation to adaptively learn the weights using the main optimizer of the problem.

$$\min_{\mathbf{w}^S \in \mathbb{R}^{d_S}} \mathcal{L}(\mathcal{T}; \mathbf{w}^S) + \alpha \cdot \mathbf{v}^\top \mathbf{f}_{KD}(\mathcal{X}; \mathbf{w}^S, \mathbf{w}^T) \quad (3)$$

where $\alpha \in [0, 1]$ is the weight indicating the importance of the distillation loss compared to the main empirical loss. $\mathbf{v} \in \mathbb{R}^K$ is the vector indicating the weight for each distillation path. For the current setting, we only consider the linear combination of losses between these paths; however, the nonlinear combination can be investigated in future studies. Next, we will describe how we can combine these losses during the training to distill the knowledge from the teacher model to the student model.

3.2 Equal Weights

The most naive and common form of combining these distillation paths is to consider all of them equally weighted during the optimization. Considering the overall objective of this problem in Eq (3), for this case, we should consider all distillation paths' weights equal, which is $\mathbf{v} = [1, \dots, 1]^\top$. Although this might work in some cases, due to different scales of these losses, in most cases we need to adjust the weights accordingly.

3.3 Hyperparameter Optimization

Another baseline for this problem is to consider the weights for each distillation path as a new hyperparameter that we need to optimize before the training procedure as shown in Figure 1(b). In this case, the weight vector \mathbf{v} is considered to be a hyperparameter that needs to be tuned. This can be done using different hyperparameter optimization approaches such as grid search, random search, or bilevel optimization [12].

3.4 Multiobjective Optimization

If we look closely at the objective of the optimization in Eq (3), it can be evidently reformulated into a multiobjective optimization as shown in Figure 1(c). In multiobjective optimization, the goal is to find a Pareto stationary point, where no other local solution can dominate that solution in any of the objectives in the task. Most of the first-order gradient-based approaches, based on this notion, aim at finding the direction at every step that is not harming any of the objectives until such direction cannot be found. In this way, all the objectives are treated equally and we try to find a Pareto stationary point of the problem.

The main challenge of using multiobjective optimization in our task is that the fundamental goal of the main empirical risk and distillation losses are not the same in the essence, and the former is more important than the latter. A solution is to use preference-based approaches [28, 50], to put more emphasis on the main objective of the learning. Another way is to consider the optimization problem of multiple paths distillation as a multiobjective optimization to find the best weights \mathbf{v} and then combine it with the main empirical risk with the weight of α . In this way, at every iteration, we first find a descent direction for all distillation losses, and then combine it with the gradients of the main learning objective using the weight α . To find the descent direction for distillation losses at every iteration we use approaches introduced in different studies [11, 26, 27, 44, 50, 52, 61].

To do so, for each distillation path, we compute the gradients of the parameters of the student model affected by that loss on the mini-batch ξ . Then for each path i we combine the gradients of the parameters from different layers together in a vector $\mathbf{g}_i(\xi; \mathbf{w}_{\leq j_i}^S, \mathbf{w}_{\leq k_i}^T)$, where j_i and k_i are the layer indices of student and teacher models, respectively, for the i -th distillation path. Since in each distillation path not all of the parameters in the student model are involved, we will use zero gradients for the parameters not involved in a path when gathering all the gradients. Then, by solving the following quadratic optimization, we will find the optimal weights of \mathbf{v} for the current mini-batch as:

$$\mathbf{v}^* \in \arg \min_{\mathbf{v} \in \Delta_K} \left\| \sum_{i \in [K]} v_i \mathbf{g}_i(\xi; \mathbf{w}_{\leq j_i}^S, \mathbf{w}_{\leq k_i}^T) \right\|_2^2, \quad (4)$$

where $\Delta_K = \{p_i | 0 \leq p_i \leq 1, \sum_{i \in [K]} p_i = 1\}$ is a K -dimensional simplex. It has been shown that using the \mathbf{v}^* from Eq (4) the resulting direction is descent for all distillation paths [61].

3.5 Adaptive Distillation

Inspired by multitask learning [29, 39], we intend to learn the importance of each distillation path adaptively by considering each path as a separate task in the learning. To do so, in addition to the parameters in Eq (3), we introduce a new set of proxy parameters $\mathbf{z} = [z_1, \dots, z_K]$ to estimate $\mathbf{v} = [e^{-z_1}, \dots, e^{-z_K}]$ as shown in Figure 1(d). Thus, we update the objective of the optimization in Eq (3) as:

$$\min_{\mathbf{w}^S \in \mathbb{R}^{d_S}, \mathbf{z} \in \mathbb{R}^K} \mathcal{L}(\mathcal{T}; \mathbf{w}^S) + \alpha \cdot \underbrace{(\mathbf{v}^\top \mathbf{f}_{KD}(\mathcal{X}; \mathbf{w}^S, \mathbf{w}^T) + \sum_{i \in [K]} z_i)}_{f_d(\mathcal{X}; \mathbf{w}^S, \mathbf{w}^T, \mathbf{z})} \quad (5)$$

where the last two terms define the distillation loss in terms of the model parameters and the proxy parameters \mathbf{z} . Expressing v_i as e^{-z_i} ensures $v_i > 0 \forall z_i \in \mathbb{R}$. The term $\sum_{i \in [K]} z_i$ acts as a regularization to prevent z_i to converge to larger values that decreases v_i , and thereby, vanishes the second loss term. To understand how z_i gets updated, we inspect the gradients of the distillation loss in Eq (5) with respect to z_i on mini-batch of ξ denoted by:

$$\left. \frac{\partial f_d(\mathcal{X}; \mathbf{w}^S, \mathbf{w}^T, \mathbf{z})}{\partial z_i} \right|_{\xi} = \alpha \left[1 - e^{-z_i} \cdot \ell_{KD}^{j_i k_i}(\xi; \mathbf{w}_{\leq j_i}^S, \mathbf{w}_{\leq k_i}^T) \right] \quad (6)$$

Hence, if we have access to the true loss (full batch), based on the first-order optimality condition resulting from Eq (6), we can infer that the optimal value for v_i would be equal to

the inverse of its corresponding loss. This means that scaling the gradients of each loss at every iteration with the inverse of their true loss will give us the optimal results. However, due to the infeasibility of the true loss at every iteration, and since we are using stochastic gradient descent for the optimization, we will use the stochastic gradient in Eq (6) to update z_i values at every iteration. We expect this will converge to the optimal values for each path.

4 Empirical Studies

In this section, we explore knowledge distillation with multiple paths from a teacher using the approaches mentioned in Section 3, in comparison with our proposed adaptive distillation. For most of the experiments, unless specified, we use ResNet50 [19] as a teacher and ResNet18 as a student. We chose attention transfer (AT) [34] for feature distillation paths, and soft target (ST) [22] for logit distillation, as two paths for distillation. For attention transfer, we use the sum of squared differences between attention maps of features from layers 2 through 5 of ResNet50 and ResNet18. We use $\mathbf{v} = [v_{AT}, v_{ST}] = [1, 1]$ for ‘Equal’ weights baseline and chose $\mathbf{v} = [1000, 0.1]$ for ‘Hand-tuned’ weights baseline based on our grid search hyperparameter tuning. Multiobjective and Adaptive methods initialize weights as $\mathbf{v} = [v_{AT}, v_{ST}] = [1, 1]$ and aim to learn the best weights during training. We use $\alpha=1.0$, the relative importance of aggregated distillation losses compared to the main loss. A detailed discussion on the experimental setups with additional results are deferred to Appendix A.

4.1 Results

Image Classification: We perform experiments on CIFAR-10 and CIFAR-100 [35], as well as ImageNet-200 [68] and ImageNet-1K [60] for the image classification task. We evaluate top1 classification error on the validation dataset along with the top1 agreement error as suggested by Stanton et al. [63] for knowledge distillation methods (i.e, classification error between teacher and student). We train our student for 200 epochs with batch size 128 on CIFAR datasets and 100 epochs with batch size 256 on ImageNet-200, ImageNet-1K datasets. We chose an SGD optimizer with an initial learning rate of 0.1 for all datasets. We reduce the learning rate by a factor of 0.1 at [100, 150] epochs for CIFAR-10, [30, 60, 90] for both ImageNet datasets. For CIFAR-100, we reduce the learning rate by 0.2 at [60, 120, 160] epochs. We use random flip augmentation during training in experiments.

As it can be seen in Table 2, we report the performance of teacher and student models without any knowledge distillation, student with single knowledge distillation path using hand-tuned weights, and baselines with multiple knowledge distillation paths which include hand-tuned weights, equal weights, and multiobjective optimization. We observe that adaptive distillation outperforms all other baselines in all three datasets. Among the baseline methods, the multiobjective optimization achieved better results on CIFAR-10 than others while hand-tuned and equal weights achieved better results on CIFAR-100 and ImageNet-200 respectively. Our proposed adaptive distillation method always outperforms the baseline methods and also the teacher model in terms of top1 classification error, as well as top1 agreement error between the teacher and student networks, which demonstrate the efficacy of our proposed aggregation approach for the knowledge distillation process.

Comparisons with state-of-the-arts: In Figure 2, we study the performance of our adaptive distillation methods when added to a combination of existing state-of-the-art distillation methods for image classification task on CIFAR-10, CIFAR-100 and ImageNet-200

	Path(s)		CIFAR-10		CIFAR-100		ImageNet-200		ImageNet-1K	
	AT	ST	top1 err (↓)	top1 agr (↓)	top1 err (↓)	top1 agr (↓)	top1 err (↓)	top1 agr (↓)	top1 err (↓)	top1 agr (↓)
Student			5.19	-	24.84	-	42.34	-	30.10	-
Single	✓		4.51	3.20	22.87	19.04	39.80	35.07	29.42	21.80
Single		✓	5.10	4.06	22.67	19.71	40.24	35.38	29.09	21.66
Hand-tuned	✓	✓	4.75	3.74	21.47	17.40	39.00	32.89	-	-
Equal	✓	✓	5.05	4.23	21.75	18.53	38.25	29.83	28.73	20.52
Multiobjective	✓	✓	4.65	3.63	21.58	18.09	39.75	33.22	-	-
Adaptive	✓	✓	4.39	3.12	20.04	15.89	37.68	29.46	28.39	20.29
Teacher			4.63	-	20.10	-	40.81	-	23.45	-

Table 2: Validation results for different knowledge distillation methods on CIFAR-10, CIFAR-100, ImageNet-200 and ImageNet-1K. We report top1 classification error (%) for the student model and top1 agreement error (%) defined in [63] (i.e, classification error between teacher and student). Adaptive distillation achieves the best results, even better than the teacher model, in all datasets. (↓) indicates lower the better.




Figure 2: Comparison between several combinations of existing knowledge distillation methods using Hand-tuned vs Adaptive (ours) methods on (a) CIFAR-10, (b) CIFAR-100 and (c) ImageNet-200 datasets. AT: attention transfer [34], ST: soft target [22], OFD: overhaul of feature distillation [21], FM: feature map [71], and SR: Softmax regression and representation learning [71].

datasets. In addition to Attention Transfer (AT) [34] and Soft Target (ST) [22] distillation methods, we use the recent and advanced distillation methods like Overhaul of Feature Distillation (OFD) [21], Feature Map (FM) [71] and Softmax Regression and Representation Learning (SR) [71]. For each combination, we chose two distillation paths that include distillation at intermediate features (AT, OFD, FM) and logits (ST, SR). We observed that adding adaptive methods to the existing methods consistently improved the performance. This experiment shows that our Adaptive Distillation is orthogonal to other knowledge distillation methods and can be used on top of those approaches to boost their performance.

Ablations: In Table 3, we present how adaptive distillation performs when we add more knowledge distillation paths. We add neural selective transfer (NST) [24] for feature level distillation and regression logits (ℓ_2 -Logit) [3] for output layers in addition to existing paths. We use $\mathbf{v} = [v_{AT}, v_{ST}, v_{NST}, v_{\ell_2\text{-Logit}}] = [100, 0.1, 10, 0.1]$ for hand-tuned baselines. We construct three variants of adaptive distillation methods by choosing [AT, ST], [NST, ℓ_2 -Logit] and [AT, ST, NST, ℓ_2 -Logit]. Finally, we also explore the robustness of adaptive distillation by treating each of the 4 residual layers in the backbone as a unique path for distillation. This results in 5 (4+1) paths for [AT, ST] and [NST, ℓ_2 -Logit] experiments, and 10 (8+2) paths

for [AT, ST, NST, ℓ_2 -Logit]. We refer to these models as ‘Adaptive-layerwise’. For more information on different structures refer to Appendix A.

	Path(s)				CIFAR-100		ImageNet-200	
	AT	ST	NST	ℓ_2 -Logit	top1 err (↓)	top1 agr (↓)	top1 err (↓)	top1 agr (↓)
Student					24.84	-	42.34	-
Single	✓				22.87	19.04	39.80	35.07
Single		✓			22.67	19.71	40.24	35.38
Single			✓		22.45	18.61	40.32	37.48
Single				✓	22.47	19.18	39.95	34.15
Adaptive	✓	✓			20.04	15.89	37.68	29.46
Adaptive			✓	✓	20.49	15.55	38.67	30.71
Adaptive	✓	✓	✓	✓	20.16	15.09	38.29	27.84
Adaptive-layerwise	✓	✓			20.43	15.57	38.32	29.44
Adaptive-layerwise			✓	✓	20.97	16.14	39.50	30.42
Adaptive-layerwise	✓	✓	✓	✓	20.85	15.25	38.34	28.75
Teacher					20.10	-	40.81	-

Table 3: Validation results on CIFAR-100 and ImageNet-200 with additional knowledge distillation paths using neural selective transfer (NST) [24] and (ℓ_2 -Logit) [3].

We observed that our proposed adaptive methods continue to outperform their single baselines in terms of both top1 classification error and top1 agreement error. Adaptive methods with [AT, ST] paths yielded the best results in terms of top1 classification error compared to other counterparts. However, the adaptive method with [AT, ST, NST, ℓ_2 -Logit] achieves the best agreement error, which is the objective of the distillation part. This suggests that shifting the focus from distillation (i.e, reducing the value of α) to main loss could improve top1 classification error. Also, our adaptive methods with [AT, ST, NST, ℓ_2 -Logit] paths performed slightly better than [NST, ℓ_2 -Logit] and worse compared to [AT, ST] paths. This suggests that adaptive distillation can reduce negative effects caused by [NST, ℓ_2 -Logit] to some degree. The same patterns can be observed with layerwise approaches as well.

Semantic Segmentation: For this task, we perform experiments on Cityscapes [9] and ADE20K [78] datasets and evaluate mean intersection over union (mIoU) and mean accuracy (mAcc). We use semantic FPN [32] network with ResNet [19] backbone(s). We initialize our backbone(s) with ImageNet pretrained weights. We use attention transfer (AT) for backbone features and soft target (ST) for outputs. We flatten the semantic segmentation output to generate soft targets. We choose an SGD optimizer with an initial learning rate of 0.01, polynomial learning rate policy (power=0.9) and train for 80K iterations with a batch size of 8 (512x1024 crop) images. Table 5 indicates the performance of knowledge distillation with multiple paths for this task. We observe that whenever the gap between teacher and student models’ performance is large, adaptive distillation outperforms other approaches in this task. This can be inferred from results in ADE 20K dataset, where adaptive distillation outperforms hand-tuned ones. When this gap is small, adaptive distillation can achieve very close performance to the best hand-tuned model, which essentially reduces the time for exhaustive search for best hyperparameters.

Object Detection: We perform experiments on Cityscapes [9] and COCO [42] datasets and evaluate mean average precision(mAP). We use RetinaNet [43] with Generalized Focal Loss [40] built using ResNet [19] backbone(s). We use ResNet18 as student backbone and ResNet101 as the teacher backbone. We use feature-based knowledge distillation [75] for features (‘Feats_b’ indicate features from the backbone layers and ‘Feats_p’ indicate features from the pyramid layers) and localization distillation [77] for bounding box generator (re-

	Path(s)		Cityscapes		ADE 20K	
	AT	ST	mIoU (↑)	mAcc (↑)	mIoU(↑)	mAcc (↑)
Student			73.42	81.33	32.84	42.50
Single	✓		75.75	83.59	34.03	43.69
Single		✓	75.25	82.89	33.19	42.79
Hand-tuned	✓	✓	75.67	83.49	33.43	43.66
Equal	✓	✓	74.38	81.77	34.06	42.95
Multiobjective	✓	✓	68.99	77.19	34.01	43.04
Adaptive	✓	✓	75.43	82.51	34.73	44.28
Teacher			74.52	81.86	36.45	46.96

Table 4: Performance of knowledge distillation methods with multiple paths on Cityscapes [9] and ADE 20K [78] datasets for semantic segmentation. (↑) indicates higher the better.

	Path(s)		Cityscapes	COCO	Path(s)		Cityscapes	COCO
	Feats _B	Box	mAP (↑)	mAP (↑)	Feats _P	Box	mAP (↑)	mAP (↑)
Student			38.1	36.1			38.1	36.1
Single	✓		39.6	36.3	✓		39.9	36.8
Single		✓	39.8	36.2		✓	39.8	36.3
Hand-tuned	✓	✓	39.7	36.4	✓	✓	40.1	36.7
Multiobjective	✓	✓	40.0	36.5	✓	✓	40.6	36.8
Adaptive	✓	✓	40.2	37.1	✓	✓	40.9	37.7
Teacher			42.9	42.7			42.9	42.7

Table 5: Performance of knowledge distillation with multiple paths on Cityscapes [9] and COCO [42] datasets for object detection. (↑) indicates higher the better. Feats_B: features from backbone layers, Feats_P: features from pyramid layers, and Box: bounding box generator.

ferred as ‘Box’). We initialize our backbone(s) with COCO pretrained weights. We choose an SGD optimizer with an initial learning rate of (0.01, 0.02) and train for a total of (64, 12) epochs on Cityscapes and COCO datasets. We reduce the learning rate by factor of 10 after 56 epochs on Cityscapes while after 8, 11 epochs on COCO dataset. We use a batch size of 8 (512x1024 crop) images. For hand-tuned weights, we use the hyperparameters suggested by the respective authors. Results on the Table 5 indicate that adaptive distillation and multiobjective outperform Hand-tuned baselines. When ‘Feats_B’ is used for feature distillation along with Bbox distillation, hand-tuned baseline failed to minimize the negative effects caused by distillation on backbone features.

5 Conclusion and Future Directions

In this paper, we have explored different methods for aggregating different paths for efficient knowledge distillation from a teacher to a student. We have proposed an adaptive approach, with which we intend to learn the importance of each distillation path during the training process. The effectiveness of this approach has been corroborated by our extensive empirical studies on various tasks such as classification, semantic segmentation, and object detection. Moreover, we have introduced another baseline based on multiobjective optimization.

Although our approach has been examined on a single teacher with multiple distillation paths, the extension of these methods to multiple teacher can be investigated in future works. Moreover, as another future direction, a theoretical investigation of these approaches can further illuminate the effect of each path on the distillation process.

References

- [1] Zeyuan Allen-Zhu and Yuanzhi Li. Towards understanding ensemble, knowledge distillation and self-distillation in deep learning. *arXiv preprint arXiv:2012.09816*, 2020.
- [2] Anubhav Ashok, Nicholas Rhinehart, Fares Beainy, and Kris M. Kitani. N2n learning: Network to network compression via policy gradient reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2018.
- [3] Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- [4] Davis Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John Guttag. What is the state of neural network pruning? *arXiv preprint arXiv:2003.03033*, 2020.
- [5] Rich Caruana. Multitask learning. *Machine Learning Journal*, 28(1):41–75, 1997.
- [6] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 742–751, 2017.
- [7] Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International Conference on Machine Learning (ICML)*, 2018.
- [8] S. Chennupati, Ganesh Sistu., Senthil Yogamani., and Samir Rawashdeh. Auxnet: Auxiliary tasks enhanced semantic segmentation for automated driving. In *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISAPP)*, pages 645–652. INSTICC, SciTePress, 2019.
- [9] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [10] Xing Dai, Zeren Jiang, Zhao Wu, Yiping Bao, Zhicheng Wang, Si Liu, and Erjin Zhou. General instance distillation for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7842–7851, 2021.
- [11] Jean-Antoine Désidéri. Multiple-gradient descent algorithm (mgda) for multiobjective optimization. *Comptes Rendus Mathématique*, 350(5-6):313–318, 2012.
- [12] Luca Franceschi, Paolo Frasconi, Saverio Salzo, Riccardo Grazi, and Massimiliano Pontil. Bilevel programming for hyperparameter optimization and meta-learning. In *International Conference on Machine Learning (ICML)*, pages 1568–1577. PMLR, 2018.
- [13] Takashi Fukuda, Masayuki Suzuki, Gakuto Kurata, Samuel Thomas, Jia Cui, and Bhuvana Ramabhadran. Efficient knowledge distillation from an ensemble of teachers. In *Interspeech*, pages 3697–3701, 2017.

- [14] Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. Knowledge distillation: A survey. *International Journal of Computer Vision (IJCV)*, 129(6):1789–1819, 2021.
- [15] Michelle Guo, Albert Haque, De-An Huang, Serena Yeung, and Li Fei-Fei. Dynamic task prioritization for multitask learning. In *European Conference on Computer Vision (ECCV)*, pages 282–299. Springer, 2018.
- [16] Farzin Haddadpour, Mohammad Mahdi Kamani, Aryan Mokhtari, and Mehrdad Mahdavi. Federated learning with compression: Unified analysis and sharp guarantees. In *International Conference on Artificial Intelligence and Statistics*, pages 2350–2358. PMLR, 2021.
- [17] Song Han, Jeff Pool, John Tran, and William J. Dally. Learning both weights and connections for efficient neural networks. In *Advances Neural Information Processing Systems (NeurIPS)*, page 1135–1143, Cambridge, MA, USA, 2015.
- [18] Song Han, Huizi Mao, and William J. Dally. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. In Yoshua Bengio and Yann LeCun, editors, *International Conference on Learning Representations (ICLR)*, 2016.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [20] Tong He, Chunhua Shen, Zhi Tian, Dong Gong, Changming Sun, and Youliang Yan. Knowledge adaptation for efficient semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 578–587, 2019.
- [21] Byeongho Heo, Jeesoo Kim, Sangdoon Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1921–1930, 2019.
- [22] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [23] Qiangui Huang, Kevin Zhou, Suyu You, and Ulrich Neumann. Learning to prune filters in convolutional neural networks. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 709–718. IEEE, 2018.
- [24] Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv preprint arXiv:1707.01219*, 2017.
- [25] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. In *British Machine Vision Conference (BMVC)*. British Machine Vision Association, 2014.
- [26] Mohammad Mahdi Kamani. Multiobjective optimization approaches for bias mitigation in machine learning. 2020.

- [27] Mohammad Mahdi Kamani, Farzin Haddadpour, Rana Forsati, and Mehrdad Mahdavi. Efficient fair principal component analysis. *arXiv preprint arXiv:1911.04931*, 2019.
- [28] Mohammad Mahdi Kamani, Rana Forsati, James Z Wang, and Mehrdad Mahdavi. Pareto efficient fairness in supervised learning: From extraction to tracing. *arXiv preprint arXiv:2104.01634*, 2021.
- [29] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [30] Md Fahim Faysal Khan, Mohammad Mahdi Kamani, Mehrdad Mahdavi, and Vijaykrishnan Narayanan. Learning to quantize deep neural networks: A competitive-collaborative approach. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2020.
- [31] Yong-Deok Kim, Eunhyeok Park, Sungjoo Yoo, Taelim Choi, Lu Yang, and Dongjun Shin. Compression of deep convolutional neural networks for fast and low power mobile applications. In Yoshua Bengio and Yann LeCun, editors, *International Conference on Learning Representations, (ICLR)*, 2016.
- [32] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6399–6408, 2019.
- [33] I. Kokkinos. Ubernet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5454–5463, July 2017.
- [34] Nikos Komodakis and Sergey Zagoruyko. Paying more attention to attention: improving the performance of convolutional neural networks via attention transfer. In *International Conference on Learning Representations (ICLR)*, 2017.
- [35] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [36] Jogendra Nath Kundu, Nishank Lakkakula, and R Venkatesh Babu. Um-adapt: Unsupervised multi-task adaptation using adversarial cross-task distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1436–1445, 2019.
- [37] Xu Lan, Xiatian Zhu, and Shaogang Gong. Knowledge distillation by on-the-fly native ensemble. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31. Curran Associates, Inc., 2018.
- [38] Carlos Lassance, Myriam Bontonou, Ghouthi Boukli Hacene, Vincent Gripon, Jian Tang, and Antonio Ortega. Deep geometric knowledge distillation with graphs. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8484–8488. IEEE, 2020.

- [39] Isabelle Leang, Ganesh Sistu, Fabian Bürger, Andrei Bursuc, and Senthil Yogamani. Dynamic task weighting methods for multi-task networks in autonomous driving systems. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–8. IEEE, 2020.
- [40] Xiang Li, Wenhai Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [41] Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. *arXiv preprint arXiv:2006.07242*, 2020.
- [42] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [43] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [44] Xi Lin, Hui-Ling Zhen, Zhenhua Li, Qing-Fu Zhang, and Sam Kwong. Pareto multi-task learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32. Curran Associates, Inc., 2019.
- [45] Iou-Jen Liu, Jian Peng, and Alexander Schwing. Knowledge flow: Improve upon your teachers. In *International Conference on Learning Representations (ICLR)*, 2019.
- [46] Shikun Liu, Edward Johns, and Andrew J Davison. End-to-end multi-task learning with attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1871–1880, 2019.
- [47] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I. Jordan. Unsupervised domain adaptation with residual transfer networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, page 136–144, 2016.
- [48] David Lopez-Paz, Léon Bottou, Bernhard Schölkopf, and Vladimir Vapnik. Unifying distillation and privileged information. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations (ICLR)*, 2016.
- [49] Jiaqi Ma and Qiaozhu Mei. Graph representation learning via multi-task knowledge distillation. In *Advances in Neural Information Processing Systems (NeurIPS), Graph Representation Learning Workshop*, 2019.
- [50] Debabrata Mahapatra and Vaibhav Rajan. Multi-task learning with user preferences: Gradient descent with controlled ascent in pareto optimization. In *International Conference on Machine Learning (ICML)*, pages 6597–6607. PMLR, 2020.
- [51] Umberto Michieli and Pietro Zanuttigh. Knowledge distillation for incremental learning in semantic segmentation. *Computer Vision and Image Understanding*, 205: 103167, 2021.

- [52] Kaisa Miettinen. *Nonlinear multiobjective optimization*, volume 12. Springer Science & Business Media, 2012.
- [53] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. Improved knowledge distillation via teacher assistant. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.
- [54] Asit Mishra and Debbie Marr. Apprentice: Using knowledge distillation techniques to improve low-precision network accuracy. In *International Conference on Learning Representations*, 2018.
- [55] SeongUk Park and Nojun Kwak. Feature-level ensemble knowledge distillation for aggregating knowledge from multiple networks. In *European Conference on Artificial Intelligence*, 2020.
- [56] Yongri Piao, Zhengkun Rong, Miao Zhang, Weisong Ren, and Huchuan Lu. A2dele: Adaptive and attentive depth distiller for efficient rgb-d salient object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9060–9069, 2020.
- [57] Antonio Polino, Razvan Pascanu, and Dan Alistarh. Model compression via distillation and quantization. In *International Conference on Learning Representations (ICLR)*, 2018.
- [58] Ilija Radosavovic, Piotr Dollár, Ross Girshick, Georgia Gkioxari, and Kaiming He. Data distillation: Towards omni-supervised learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4119–4128, 2018.
- [59] Adriana Romero, Samira Ebrahimi Kahou, Polytechnique Montréal, Y. Bengio, Université De Montréal, Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *International Conference on Learning Representations (ICLR)*, 2015.
- [60] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- [61] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 525–536. Curran Associates, Inc., 2018.
- [62] Hyowoon Seo, Jihong Park, Seungeun Oh, Mehdi Bennis, and Seong-Lyun Kim. Federated knowledge distillation. *arXiv preprint arXiv:2011.02367*, 2020.
- [63] Samuel Stanton, Pavel Izmailov, Polina Kirichenko, Alexander A Alemi, and Andrew Gordon Wilson. Does knowledge distillation really work? *arXiv preprint arXiv:2106.05945*, 2021.
- [64] Marvin Teichmann, Michael Weber, Marius Zoellner, Roberto Cipolla, and Raquel Urtasun. Multinet: Real-time joint semantic reasoning for autonomous driving. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 1013–1020. IEEE, 2018.

- [65] Lin Wang and Kuk-Jin Yoon. Knowledge distillation and student-teacher learning for visual intelligence: A review and new outlooks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2021.
- [66] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4933–4942, 2019.
- [67] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A Efros. Dataset distillation. *arXiv preprint arXiv:1811.10959*, 2018.
- [68] Jiayu Wu, Qixiang Zhang, and Guoxi Xu. Tiny imagenet challenge, 2017.
- [69] Liuyu Xiang, Guiguang Ding, and Jungong Han. Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In *European Conference on Computer Vision (ECCV)*, pages 247–263. Springer, 2020.
- [70] J. Xie, Bing Shuai, Jianfang Hu, J. Lin, and W. Zheng. Improving fast segmentation with teacher-student learning. In *British Machine Vision Conference (BMVC)*, 2018.
- [71] Jing Yang, Brais Martinez, Adrian Bulat, and Georgios Tzimiropoulos. Knowledge distillation via softmax regression representation learning. In *International Conference on Learning Representations*, 2021.
- [72] Li Yuan, Francis EH Tay, Guilin Li, Tao Wang, and Jiashi Feng. Revisiting knowledge distillation via label smoothing regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3903–3911, 2020.
- [73] Zhaoyang Zeng, Bei Liu, Jianlong Fu, Hongyang Chao, and Lei Zhang. Wsod2: Learning bottom-up and top-down objectness distillation for weakly-supervised object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8292–8300, 2019.
- [74] Chenrui Zhang and Yuxin Peng. Better and faster: Knowledge transfer from multiple self-supervised learning tasks via graph distillation for video classification. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI*, pages 1135–1141. International Joint Conferences on Artificial Intelligence Organization, 7 2018. doi: 10.24963/ijcai.2018/158.
- [75] Linfeng Zhang and Kaisheng Ma. Improve object detection with feature-based knowledge distillation: Towards accurate and efficient detectors. In *International Conference on Learning Representations (ICLR)*, 2021.
- [76] Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. Deep mutual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4320–4328, 2018.
- [77] Zhaohui Zheng, Rongguang Ye, Ping Wang, Jun Wang, Dongwei Ren, and Wangmeng Zuo. Localization distillation for object detection. *arXiv e-prints*, 2021.
- [78] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.