Towards Oracle Knowledge Distillation with Neural Architecture Search

Minsoo Kang,^{1,3*} Jonghwan Mun,^{2,3*} Bohyung Han^{1,3}

¹Computer Vision Lab., ASRI, Seoul National University, Korea

²Computer Vision Lab., POSTECH, Korea

³Neural Processing Research Center (NPRC), Seoul National University, Korea

¹{kminsoo, bhhan}@snu.ac.kr ²jonghwan.mun@postech.ac.kr

Abstract

We present a novel framework of knowledge distillation that is capable of learning powerful and efficient student models from ensemble teacher networks. Our approach addresses the inherent model capacity issue between teacher and student and aims to maximize benefit from teacher models during distillation by reducing their capacity gap. Specifically, we employ a neural architecture search technique to augment useful structures and operations, where the searched network is appropriate for knowledge distillation towards student models and free from sacrificing its performance by fixing the network capacity. We also introduce an oracle knowledge distillation loss to facilitate model search and distillation using an ensemble-based teacher model, where a student network is learned to imitate oracle performance of the teacher. We perform extensive experiments on the image classification datasets—CIFAR-100 and TinyImageNet—using various networks. We also show that searching for a new student model is effective in both accuracy and memory size and that the searched models often outperform their teacher models thanks to neural architecture search with oracle knowledge distillation.

Introduction

Knowledge Distillation (KD) aims to transfer representations from one model to another, where the source plays a role as a *teacher* while the target becomes a *student* mimicking the representations of the teacher. KD is widely used to learn a compact student model with the help of a powerful teacher model, *e.g.*, a very deep neural network or an ensemble of multiple neural networks. Existing algorithms related to KD (Heo et al. 2019; Hinton, Vinyals, and Dean 2015; Park et al. 2019; Romero et al. 2015; Zagoruyko and Komodakis 2017; Zhang et al. 2018) are typically interested in how to improve accuracy by designing an effective training procedure.

Model ensemble (Breiman 1996; Freund, Schapire, and Abe 1999; Guzman-Rivera, Batra, and Kohli 2012; Lee et al. 2017; 2016; 2015; Mun et al. 2018) is a useful technique to boost performance using multiple models trained

independently (or sometimes jointly). When we consider Independent Ensemble (IE) in deep neural networks, multiple neural networks of an identical architecture are trained with different random seeds and the final output is determined by model averaging or majority voting. Although ensemble modeling is effective to achieve high accuracy with moderate amount of effort, its inferences based on simple model averaging or majority voting still have substantial gaps with oracle predictions achieved by the best model selection in the ensemble, and its applicability to resource-hungry system is limited due to the large model size and heavy power consumption.

Although KD with ensemble learning addresses the aforementioned issues partly by transferring the information in an ensemble model to a single network, it is still challenging to train the competitive student compared to the ensemble teacher with a large number of networks. Table 1 presents our observation about the performance of KD with respect to the number of models for ensemble. The accuracy of teacher and student improves gradually in general as the number of models increases while students mostly fail to reach accuracy of teachers and its differences are getting larger. This is partly because a large gap in model capacity between student and teacher hinders learning process of KD as discussed in (Mirzadeh et al. 2019), and the simple objective function to fit the representations of the teacher given by model averaging is not effective to take full advantage of teacher models. In other words, the limited capacity in the student network becomes a bottleneck of KD, which implies that increasing capacity of student models would be beneficial to reduce the performance gap between teacher and student.

We propose an advanced framework for knowledge distillation from ensemble models, which aims to maximize accuracy and efficiency of student networks at the same time. Contrary to the existing methods assuming that a student model is fixed, we adapt its structure and size, and make it more appropriate for knowledge distillation by alleviating the model capacity gap issue. We also address how to effectively extract information from a teacher especially when the teacher is an ensemble of multiple models. To this end, students are made to learn the most accurate information from teacher, which is realized by transferring knowledge from

*Both authors contributed equally

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Table 1: The performance comparison while varying the number of networks in an ensemble-based teacher model on the CIFAR-100 dataset using ResNet-32 and DenseNet-40-12 networks. When the number of ensemble is 1, student and teacher networks are identical.

number of ensemble	ResNet-32			DenseNet-40-12		
	Teacher	Student	T-S	Teacher	Student	T-S
1	69.11	-	-	74.30	-	-
2	73.77	73.84	-0.07	77.47	77.82	-0.35
3	75.57	74.12	1.45	78.70	78.03	0.67
4	76.36	74.10	2.26	79.32	78.16	1.16
5	76.87	74.67	2.20	79.77	78.43	1.34

the optimal ensemble combination for each example. We call this strategy Oracle knowledge Distillation (OD), which encourages student models to achieve oracle accuracy of ensemble teacher models. Since the inherent model capacity gap incurs critical challenges in KD and OD makes the capacity gap larger, we incorporate neural architecture search with oracle knowledge distillation; this strategy facilitates to identify the optimal student model with sufficient capacity, which is well-suited for distillation. In practice, our algorithm searches for a slightly larger model than the backbone student network for effective knowledge distillation, reduces the model capacity gap between student and teacher, and achieves competitive accuracy of the student model.

The main contributions of our work are summarized as follows:

- We propose a novel framework for knowledge distillation by incorporating neural architecture search. The proposed algorithm addresses capacity issue in KD and aims to identify the optimal structures and operations with adaptive model sizes.
- Our algorithm introduces a novel oracle knowledge distillation loss, which is particularly useful for an ensemble teacher model. We claim that the student networks mimicking oracle predictions have a potential for achieving higher accuracy than the teacher especially when combined with neural architecture search.
- We demonstrate outstanding performance of the proposed method in diverse settings. We also make a comprehensive analysis about knowledge distillation from ensemble teacher models, including various issues related to model capacity gap, objective function for architecture search, and loss function for knowledge distillation.

The rest of the paper is organized as follows. We first discuss several related works. Then, we describe the details of the proposed framework. Finally, we present extensive experimental results and concludes our paper.

Related Works

KD is originally proposed to learn compact and fast models and has widely been applied to many practical applications including object detection (Chen et al. 2017; Li, Jin, and Yan 2017), face recognition (Luo et al. 2016) and image retrieval (Chen, Wang, and Zhang 2018). The main idea of KD is to transfer information from one model to another, and

it is realized by learning a student network to mimic the output distributions of a teacher network (Hinton, Vinyals, and Dean 2015). Recently, several approaches (Heo et al. 2019; Park et al. 2019; Romero et al. 2015; Zagoruyko and Komodakis 2017) have been proposed to improve performance of KD. They address how to extract information better from teacher networks and deliver it to students using the activations of intermediate layers (Romero et al. 2015), attention maps (Zagoruyko and Komodakis 2017), encoded transportable factors (Kim, Park, and Kwak 2018), decision boundary information obtained by adversarial samples (Heo et al. 2019) or relational information between training examples (Park et al. 2019). Also, instead of transferring information from teacher to student in one direction, (Zhang et al. 2018) proposes a mutual learning strategy, where both models are trained jointly through a bidirectional interactions.

In addition to model compression by learning small and efficient models, the concept of knowledge distillation is often used for other purposes (Chen, Goodfellow, and Shlens 2016; Li and Hoiem 2016; Mun et al. 2018; Noroozi et al. 2018). For example, (Chen, Goodfellow, and Shlens 2016) proposes a framework incrementally learning larger networks from small networks using knowledge distillation while (Li and Hoiem 2016) employs the idea to overcome a catastrophic forgetting issue in continual learning scenarios. MCL-KD (Mun et al. 2018) utilizes distillation to balance between model specialization and generalization within a multiple choice learning framework for visual question answering. On the other hand, (Noroozi et al. 2018) proposes a novel framework of self-supervised learning by distilling learned representations rather than fine-tuning the learned parameters via self supervision for the target tasks.

Our work is different in the following two aspects compared to the existing approaches. First, we address model capacity issue in a student model by searching for an augmented architecture appropriate for KD while prior methods are interested in how to improve the KD given a student model with a fixed capacity. Second, although model ensemble is suitable for KD and effectively constructs a powerful teacher model, existing methods rely on a naïve knowledge transfer from teachers to students via model averaging or majority voting. On the contrary, we employ oracle knowledge distillation, which is based on an intuitive and effective loss function for sophisticated knowledge transfer. Note that the proposed loss function gives a chance to outperform the teachers with simple model averaging for students.

Figure 1: Comparison between standard KD and our proposed OD for the ensemble-based teacher model. In our approach, we train a student network from only the correct models (red arrows) to imitate the oracle predictions of ensemble teacher.

Neural Architecture Search (NAS) is an AutoML technique to identify deep neural network models automatically, which is useful to reduce human effort on manual architecture design. For the purpose, Zoph and Le (Zoph and Le 2017) employ a RNN controller that searches for the optimal models and is trained to maximize the expected reward—accuracy in validation set—using REINFORCE (Williams 1992). While NAS often suffers from huge time complexity for model search, ENAS (Pham et al. 2018) accelerates searching process by sharing the weights in the building blocks of all candidate networks. In addition to the accuracy of target models, MnasNet (Tan et al. 2019) considers inference latency in searching for the optimal model and performs a joint optimization of accuracy and execution time on mobile devices via reinforcement learning. DARTS (Liu, Simonyan, and Yang 2018) and NAO (Luo et al. 2018) are gradient-based algorithms realized by a continuous relaxation of architecture representation. PNAS (Liu et al. 2018) incorporates a progressive search strategy to reduce search cost using sequential model-based optimization. Recently, the effectiveness of NAS on image classification tasks leads to a variety of applications to semantic segmentation (Chen et al. 2018; Liu et al. 2019) and object detection (Chen et al. 2019).

From the NAS perspective, our algorithm is unique because it deals with model capacity problems in the context of knowledge distillation. Note that, instead of searching for a student network architecture from scratch, we start model search from the student network and increase the size of the network to identify the optimal architecture.

Neural Architecture Search with Oracle Knowledge Distillation

We propose a novel knowledge distillation framework using neural architecture search with an oracle knowledge distillation loss, which is designed for transfer learning from ensemble-based teacher models. This is motivated by the fact that knowledge distillation is less effective when the capacity gap (*e.g.*, the number of parameters) between teacher and student is large as discussed in (Mirzadeh et al. 2019). Incorporating oracle knowledge distillation loss aggravates the situation by forcing the student to face more challenging task (*i.e.*, oracle performance of ensemble model). Note that since our goal is to address this capacity issue, we are

interested in increasing model size from the student network contrary to the more common direction of model compression (Ashok et al. 2018). We empirically show that the combination of neural architecture search and oracle knowledge distillation in searching and training networks is particularly helpful to improve both accuracy and efficiency in various scenarios.

Knowledge Distillation

KD (Hinton, Vinyals, and Dean 2015) aims to transfer knowledge of a teacher network to a student. This objective is typically achieved by minimizing the distance between the output distributions of student and teacher. In other words, given a ground-truth label $y^{(i)}$ and a representation (logit) of a teacher network $l_t^{(i)}$ for an example $x^{(i)}$, a student network learns a representation $l_s^{(i)}$ that minimizes distillation loss $\mathcal{L}_{KD}$, which is given by

$$\mathcal{L}_{KD}(l_s^{(i)}, l_t^{(i)}, y^{(i)}) = \lambda \mathcal{L}_{CE}(l_s^{(i)}, y^{(i)}) + (1 - \lambda) \mathcal{L}_{KL}(l_s^{(i)}, l_t^{(i)}), \quad (1)$$

where $\mathcal{L}_{CE}(l_s^{(i)}, y^{(i)})$ means Cross-Entropy (CE) loss and $\mathcal{L}_{KL}(l_s^{(i)}, l_t^{(i)})$ denotes a loss term related to Kullback-Leibler (KL) divergence. Each term is further defined as

$$\mathcal{L}_{CE}(l_s^{(i)}, y^{(i)}) = \mathcal{H}(\sigma(l_s^{(i)}), y^{(i)}), \quad (2)$$

$$\mathcal{L}_{KL}(l_s^{(i)}, l_t^{(i)}) = T^2 D_{KL}(\sigma(l_s^{(i)}/T) || \sigma(l_t^{(i)}/T)), \quad (3)$$

where $\mathcal{H}(\cdot, \cdot)$ and $D_{KL}(\cdot, \cdot)$ are the cross-entropy and the KL-divergence functions, respectively while $\sigma(\cdot)$ denotes a softmax function and T is a temperature parameter. The student network is trained to predict the correct labels by $\mathcal{L}_{CE}$ and imitates the output distribution of the teacher network by $\mathcal{L}_{KL}$, where the two losses are balanced by a hyperparameter λ . The softmax distributions of student and teacher networks are softened because the logits are scaled with the temperature $T \geq 1$ in $\mathcal{L}_{KL}$.

Conceptually, any network outperforming student networks can be used as a teacher model. In this work, we consider the ensemble of student networks as the teacher model due to the following two reasons; 1) model ensemble is a straightforward method to improve accuracy of any state-of-the-art network; 2) oracle prediction is available in the ensemble model, which allows students to learn better representation and achieve higher accuracy.

Figure 2: Overview of our algorithm. Given a teacher model based on an ensemble of independently learned multiple student networks, we search for a slightly larger network from a backbone network (*i.e.*, student). LSTM controller provides candidate networks by sampling add-on operations at the end of individual stages in the student. We train the controller by maximizing the expected reward—accuracy of candidates on the validation set—while the candidates are learned with oracle knowledge distillation loss ($\mathcal{L}_{OD}$). We maintain multiple candidate models throughout the optimization process and the best model below the memory constraint is selected as a new student model and it is re-trained from scratch with $\mathcal{L}_{OD}$.

Oracle Knowledge Distillation

Ensemble learning is a powerful technique to improve accuracy by diversifying predictions of multiple models. In general, on the model distillation, a student network is trained to resemble the average predictions of an ensemble model as illustrated in Figure 1(a) due to the absence of model selection capability; for an ensemble teacher model with N networks, the teacher logit $\bar{l}_t^{(i)}$ in Eq. (2) is obtained by averaging the logits of the networks, *i.e.*, $\bar{l}_t^{(i)} = \frac{1}{N} \sum_{j=1}^N l_{t,j}^{(i)}$. However, only a subset of the models may predict correct labels, and, consequently, the average predictions may be incorrect; the ensemble model fails to achieve its oracle accuracy, which can be realized by a proper model selection.

Inspired by this fact, we propose an oracle knowledge distillation loss to improve the performance of distillation from an ensemble teacher. Let $u_j^{(i)}$ be a binary variable about whether the j^{th} model in the ensemble teacher makes a correct prediction for an example $x^{(i)}$. Then, the student network is trained to minimize the following loss given by the ensemble of N networks:

$$\mathcal{L}_{OD} = \begin{cases} \mathcal{L}_{KD}(l_s^{(i)}, \bar{l}_t^{(i)}, y^{(i)}) & \text{if } \sum_{j=1}^N u_j^{(i)} > 0 \\ \mathcal{L}_{CE}(l_s^{(i)}, y^{(i)}) & \text{otherwise} \end{cases}, \quad (4)$$

where

$$\bar{l}_t^{(i)} = \frac{\sum_{j=1}^N u_j^{(i)} l_{t,j}^{(i)}}{\sum_{j=1}^N u_j^{(i)}}.$$

$\mathcal{L}_{OD}$ encourages the student network to achieve the oracle predictions of the ensemble model by mimicking the average predictions of correct models only as depicted in Figure 1(b). In the case that there are no correct models, we make the student network fit to ground-truth labels. Since the accuracy given by oracle predictions is always better than average predictions, the trained student network has a potential to outperform its teacher model that employs model averaging and/or majority voting.

Optimal Model Search for Knowledge Distillation

The capacity issue in a student makes the student network fail to take full advantage of ensemble model. To overcome this challenge, we propose a novel Knowledge Distillation framework with Architecture Search (KDAS) using the proposed oracle knowledge distillation loss designed for ensemble-based teacher models. In the proposed framework, our goal is to find a slightly larger network with sufficient capacity for distillation than the original student models. For the purpose, we perform the operation search from the architecture of a student and select the final model under the memory constraint as illustrated in Figure 2.

Architecture search from backbone models Most architecture search algorithms (Liu et al. 2018; Liu, Simonyan, and Yang 2018; Luo et al. 2018; Zoph and Le 2017) search for an optimal architecture from scratch. However, we perform a backbone-based architecture search, where the architecture of the student network is used as the starting point (*i.e.*, backbone) and we augment a set of operations to the backbone model during the search procedure. This strategy can reduce search space significantly and facilitate to stabilize the training process.

For the efficient and effective neural architecture search to build a larger network from a backbone model, we put add-on operations after the individual stages of the standard convolutional neural networks; modern convolutional neural networks typically consist of a series of identical components, each of which is called by stage, *e.g.*, stacked convolutions in VGGNet (Simonyan and Zisserman 2015) and stacked building blocks in ResNet (He et al. 2016). When human manually designs higher capacity of the network, each stage is often made deeper by adding more operations (He et al. 2016; Huang et al. 2017). Inspired by this convention, we perform operation search at the end of individual stages in the backbone network and identify competitive model with minimal efforts. Assuming that the student

backbone network S has k stages for feature extraction, *i.e.*, $S = \{s_1, s_2, \dots, s_k\}$, the searched network $\hat{S}$ is represented by $\hat{S} = \{s_1, \hat{o}_1, s_2, \hat{o}_2, \dots, s_k, \hat{o}_k\}$, where $\hat{o} = \{\hat{o}_1, \dots, \hat{o}_k\}$ denotes a set of add-on operations. We do not perform operation search in a classifier consisting of the global pooling and fully-connected layers.

Search space We search for a network facilitating distillation based on 7 operations with skip connections: an identity operation, convolutions with filter sizes 3×3 and 5×5 , depthwise-separable convolutions with filter sizes 3×3 and 5×5 , and max pooling and average pooling of kernel size 3×3 . Note that the output channel dimension of each new convolutional layer is set to the input channel dimension of the layer. Overall, there are $2^{\frac{L(L-1)}{2}} \times 7^L$ candidate networks in the whole search space when we add L layers in total to the backbone model.

Optimization We train an LSTM controller to sample an architecture for the oracle knowledge distillation from the predefined search space; learning the LSTM controller is achieved via REINFORCE (Williams 1992) by alternating the following two steps: 1) sampling an architecture using the LSTM controller and training it for the predefined number of iterations with the proposed oracle distillation loss ($\mathcal{L}_{OD}$) and 2) updating the LSTM controller based on the reward of the trained model.

Specifically, let us denote the learnable parameters of the LSTM controller by θ . Then, the controller is trained to maximize the expected reward $J(\theta)$ as follows:

$$J(\theta) \equiv \mathbb{E}_{\mathbf{m} \sim \pi(\mathbf{m}; \theta)} [R(\mathbf{m})], \quad (5)$$

where $\pi(\mathbf{m}; \theta)$ is a policy of the controller and $R(\mathbf{m})$ denotes the reward (*i.e.*, validation accuracy) from a sampled architecture $\mathbf{m}$. Using Monte-Carlo methods, the expected gradient on the sampled architectures is given by

$$\begin{aligned} \nabla_{\theta} J(\theta) &= \mathbb{E}_{\mathbf{m} \sim \pi(\mathbf{m}; \theta)} [R(\mathbf{m}) \nabla_{\theta} \log \pi(\mathbf{m}; \theta)] \\ &\approx \frac{1}{N} \sum_{j=1}^N [R(\mathbf{m}_j) \nabla_{\theta} \log \pi(\mathbf{m}_j; \theta)], \end{aligned} \quad (6)$$

where N is the number of sampled architectures. To reduce the variance, the reward $R(\mathbf{m}_j)$ is replaced by $R(\mathbf{m}_j) - b$, where b is a baseline function given by a moving average of the past rewards. Following ENAS (Pham et al. 2018), we perform an efficient learning scheme; certain layers with their parameters in multiple candidate networks can be shared and the reward is computed only on a batch rather than a whole validation set.

Network selection with memory constraint After convergence of the controller, we select the most accurate and efficient model as follows:

$$\mathbf{m}^* = \underset{\mathbf{m}}{\operatorname{argmax}} R(\mathbf{m}), \quad \text{s.t. } |\mathbf{m}| \leq M, \quad (7)$$

where $|\mathbf{m}|$ denotes the amount of memory spent to store the parameters in a sampled network $\mathbf{m}$ and M is the memory constraint for model selection. The selected model is re-trained from scratch with the proposed oracle knowledge distillation loss $\mathcal{L}_{OD}$.

Experiments

This section presents performance of the proposed algorithm in comparison to existing methods. We also discuss characteristics of our approach obtained from in-depth analysis.

Experimental Setup

Datasets We evaluate our algorithm on the image classification task using CIFAR-100 and TinyImageNet datasets. CIFAR-100 dataset (Krizhevsky 2009) is composed of 50,000 training and 10,000 testing images in 100 classes, where the size of image is 32×32 . TinyImageNet dataset contains 100,000 and 10,000 images from 200 object classes with their size 64×64 for training and validation, respectively. For architecture search, 10% of training images are held out as training-validation set to compute reward. For both datasets, we perform the preprocessing of subtracting means and dividing by standard deviations in individual RGB channels, and employ the standard data augmentation techniques such as random cropping with zero padding and horizontal flipping.

Implementation details We employ publicly available ENAS (Pham et al. 2018) code¹ for neural architecture search implementation in TensorFlow (Abadi et al. 2016). Given a searched network or networks designed by human, we optimize the networks for 300 epochs using SGD with Nesterov momentum (Sutskever et al. 2013) of 0.9, a weight decay of 0.0001 and a batch size of 128. Following (Ian, Zhu, and Gong 2018), the initial learning rate is set to 0.1, and is divided by 10 at 150th and 225th epoch, respectively. We also perform warm-up strategy (He et al. 2016) with learning rate of 0.01 with ResNet-110 until 400th and 900th iterations for CIFAR-100 and TinyImagenet datasets, respectively. For KD and OD, a temperature T is fixed 3 and a balancing factor λ is set to 0. We train all models 3 times and report their average and standard deviation for score reports.

Experimental Results

Main results We perform extensive experiments on the two datasets to investigate the effectiveness of architecture search and oracle knowledge distillation loss. In this experiment, we consider the following types of models, (i) M1: a teacher model of ResNet-32 $\times$ 5, (ii) M2-4: a student model of ResNet-32, (iii) M5-10: two baseline models designed by human, which are ResNet-62 and ResNet-110 having 2x and 4x model size of student, (iv) M11-19: three variants of KDAS models searched by different losses, where we constraint model size up to 2x of student. We train all models except teacher using one out of three loss functions, $\mathcal{L}_{CE}$, $\mathcal{L}_{KD}$ and $\mathcal{L}_{OD}$, as presented in Table 2.

Table 2 presents the results and we can obtain following observations. First, as we claimed earlier, the student networks trained by distillation losses (M3-4) suffer from capacity and complexity issues and fail to achieve competitive accuracy. Second, manually designed networks with higher capacity (*i.e.*, ResNet-62 and ResNet-110) partly address the capacity issue, where their accuracy is proportional

¹<https://github.com/melodyguan/enas>

Table 2: Ablation studies varying combinations of loss in searching ($\mathcal{L}_S$) and training ($\mathcal{L}_T$) networks on CIFAR-100 and TinyImageNet datasets. We employ a ResNet-32 network as the student model and ensemble of five student networks (*i.e.*, ResNet-32 $\times$ 5) as the teacher model whose oracle accuracies (%) are 87.79 and 75.30 in CIFAR-100 and TinyImageNet. The bold-faced and italic numbers mean the best algorithm in accuracy for each model and each dataset, respectively.

	Model	$\mathcal{L}_S$	$\mathcal{L}_T$	CIFAR-100		TinyImageNet		Network identified by
				Accuracy	Memory	Accuracy	Memory	
M1	Teacher	-	$\mathcal{L}_{CE}$	76.87	2.35M	62.59	2.38M	-
M2	Student	-	$\mathcal{L}_{CE}$	69.11 $\pm$ 0.24	0.47M	54.14 $\pm$ 0.65	0.48M	
M3			$\mathcal{L}_{KD}$	74.67 $\pm$ 0.10		58.68 $\pm$ 0.09		
M4			$\mathcal{L}_{OD}$	74.77 $\pm$ 0.02		58.66 $\pm$ 0.25		
M5	ResNet-62	-	$\mathcal{L}_{CE}$	72.06 $\pm$ 0.31	0.96M	58.62 $\pm$ 0.16	0.97M	Man-Made
M6			$\mathcal{L}_{KD}$	76.09 $\pm$ 0.20		61.05 $\pm$ 0.31		
M7			$\mathcal{L}_{OD}$	75.89 $\pm$ 0.19		61.25 $\pm$ 0.14		
M8	ResNet-110	-	$\mathcal{L}_{CE}$	73.77 $\pm$ 0.19	1.73M	60.24 $\pm$ 0.45	1.74M	
M9			$\mathcal{L}_{KD}$	76.77 $\pm$ 0.52		62.03 $\pm$ 0.03		
M10			$\mathcal{L}_{OD}$	76.68 $\pm$ 0.17		62.66 $\pm$ 0.53		
M11	NAS	$\mathcal{L}_{CE}$	$\mathcal{L}_{CE}$	74.55 $\pm$ 0.51	0.97M	62.01 $\pm$ 0.60	0.90M	AutoML
M12			$\mathcal{L}_{KD}$	76.85 $\pm$ 0.33		62.10 $\pm$ 0.17		
M13			$\mathcal{L}_{OD}$	77.05 $\pm$ 0.23		62.57 $\pm$ 0.11		
M14	KDAS (ours)	$\mathcal{L}_{KD}$	$\mathcal{L}_{CE}$	74.56 $\pm$ 0.35	0.93M	62.92 $\pm$ 0.10	0.95M	
M15			$\mathcal{L}_{KD}$	76.97 $\pm$ 0.08		62.34 $\pm$ 0.10		
M16			$\mathcal{L}_{OD}$	77.04 $\pm$ 0.33		62.73 $\pm$ 0.09		
M17	KDAS (ours)	$\mathcal{L}_{OD}$	$\mathcal{L}_{CE}$	75.14 $\pm$ 0.26	0.89M	62.60 $\pm$ 0.11	0.87M	
M18			$\mathcal{L}_{KD}$	76.92 $\pm$ 0.33		62.17 $\pm$ 0.12		
M19			$\mathcal{L}_{OD}$	77.27 $\pm$ 0.11		63.04 $\pm$ 0.17		

to the memory size of network. This result implies that increasing capacity of student networks is helpful for reducing performance gap between student and teacher. It is also noticeable that ResNet-110 in TinyImageNet can achieve marginally higher accuracy than teacher by applying $\mathcal{L}_{OD}$. Third, the networks identified by KDAS with distillation (M15-16, M18-19) are consistently better than Man-Made Networks (MMNs), ResNet-62 and ResNet-110, even with smaller memory size. Fourth, training networks with $\mathcal{L}_{OD}$ shows the best accuracy in most cases of KDAS. One exception is when the model is searched with $\mathcal{L}_{KD}$ but trained with $\mathcal{L}_{CE}$. In this case, the network outperforms the teacher even without distillation, which implies that the teacher would have a negative impact on the student with $\mathcal{L}_{KD}$. However, $\mathcal{L}_{OD}$ still encourages a student to learn better representation and the student outperforms the teacher model (see M14-16 on TinyImageNet). Finally, the full KDAS model (M19) with $\mathcal{L}_{OD}$ in both architecture search and training achieves the best accuracy even including teacher model. This result shows that the combination of architecture search with $\mathcal{L}_{OD}$ is particularly well-suited for improving accuracy of a student network.

Comparison with traditional NAS To validate the necessity of NAS along with knowledge distillation, we conduct NAS without a distillation loss in searching for networks and then train the searched networks with KD and OD losses. As presented in Table 2, the NAS baseline networks (M12-13) perform worse than the architectures identified by KDAS with OD (M18-19). Also, note that even with $\mathcal{L}_{CE}$, the network identified by KDAS (M17) provides better performance than the counterpart (M11).

Figure 3: Accuracies varying memory size of networks given by KDAS on the CIFAR-100 dataset with the backbone student network ResNet-32.

Analysis on varying memory size We analyze performance of student models with different sizes obtained by KDAS on CIFAR-100 dataset. Note that teacher is based on five ResNet-32 models. We search for the student network architectures by varying memory constraint from $1.5\times$ to $4\times$ of the base student model (0.47M). For comparison, we also train MMNs using the standard knowledge distillation loss ($\mathcal{L}_{KD}$), which are selected to be $2\times$ larger than the base student model. Figure 3 shows that KDAS models outperform MMNs in terms of both accuracy and efficiency; these results support the benefit of architecture search for KD. Note that KDAS models larger than 0.89M ($2\times$ larger than the base model) achieve better accuracies than the teacher model thanks to the oracle knowledge distillation loss.

Table 3: Results with various networks on the CIFAR-100 dataset. We use ResNet-218, WideResNet-76-1, WideResNet-28-2, WideResNet-28-4 networks as MMN of student ResNet-110, WideResNet-40-1, WideResNet-16-2 networks, and WideResNet-16-4 networks, respectively. Numbers in bold and italic denote the best and second-best models including the teacher model.

Method	$\mathcal{L}_s$	$\mathcal{L}_T$	ResNet-110		WideResNet-40-1		WideResNet-16-2		WideResNet-16-4	
			Accuracy	Memory	Accuracy	Memory	Accuracy	Memory	Accuracy	Memory
Teacher	-	$\mathcal{L}_{CE}$	79.24	8.67M	77.53	2.85M	77.77	3.52M	79.49	13.86M
Student	-	$\mathcal{L}_{CE}$	73.77 $\pm$ 0.19	1.73M	69.96 $\pm$ 0.15	0.57M	71.16 $\pm$ 0.30	0.70M	75.17 $\pm$ 0.24	2.77M
Student	-	$\mathcal{L}_{KD}$	76.77 $\pm$ 0.52	1.73M	74.72 $\pm$ 0.23	0.57M	75.42 $\pm$ 0.04	0.70M	78.59 $\pm$ 0.34	2.77M
MMN	-	$\mathcal{L}_{KD}$	77.39 $\pm$ 0.21	3.48M	76.48 $\pm$ 0.15	1.15M	76.97 $\pm$ 0.05	1.48M	79.28 $\pm$ 0.16	5.87M
KDAS	$\mathcal{L}_{OD}$	$\mathcal{L}_{OD}$	79.01 $\pm$ 0.28	2.73M	76.70 $\pm$ 0.25	1.14M	77.83 $\pm$ 0.23	1.30M	79.79 $\pm$ 0.24	5.47M

Table 4: Performance comparison with other KD algorithms on the CIFAR-100 dataset. We use a single ResNet-110 network as a teacher model. The bold number means the highest accuracy.

Student	CE	KD	DML	BSS	TAKD	KDAS (0.91M)
ResNet-62 (0.96M)	71.73 $\pm$ 0.03	74.57 $\pm$ 0.18	72.98 $\pm$ 1.07	73.06 $\pm$ 0.53	75.18 $\pm$ 0.13	75.82 $\pm$ 0.32
ResNet-68 (1.05M)	71.77 $\pm$ 0.06	74.82 $\pm$ 0.09	73.39 $\pm$ 0.70	73.43 $\pm$ 0.21	75.45 $\pm$ 0.12	

Evaluation with other networks To validate the model-agnostic property of our algorithm, various student models have been evaluated on CIFAR-100 dataset. Table 3 summarizes the results from the tested backbone architectures including ResNet-110, WideResNet-40-1, WideResNet-16-2, and WideResNet-16-4 (Zagoruyko and Komodakis 2016), where networks identified by KDAS with OD consistently outperform corresponding MMNs with comparable memory sizes while showing competitive or outperforming performance compared to teacher models.

Comparison with other KD methods Although our work focuses on distillation from ensemble teacher models, we also compare KDAS with other KD approaches including KD (Hinton, Vinyals, and Dean 2015), DML (Zhang et al. 2018), BSS (Heo et al. 2019) and TAKD (Mirzadeh et al. 2019). For the fair comparisons with other KD methods, we follow BSS (Heo et al. 2019) implementation environments² using PyTorch (Paszke et al. 2017). Contrary to the main experiments, we use a single model teacher given by ResNet-110 and train two student networks of ResNet-62 and ResNet-68 for other KD methods while we search for a student network with comparable size to ResNet-62 from a base model of ResNet-32 and train the searched network with KD. For TAKD, we employ ResNet-86 as a teacher assistant (TA) network. Table 4 shows that the model identified by KDAS outperforms other methods in both accuracy and memory size even without the proposed oracle knowledge distillation loss, which is not available in this experiment. This result implies that searching for student architectures is a promising direction for knowledge distillation.

Practical benefit of OD over KD The advantage of OD over KD comes from their different characteristics on training examples. OD may be useless if all the constituent models in an ensemble teacher network predict correct answers for most of training examples. However, in reality, the average training accuracy of all 5 models in the teacher is clearly lower than 100% as seen in Table 5. We also present the dis-

Table 5: Training accuracy of single ResNet-32 network on CIFAR-100 and TinyImageNet datasets. We also present the percentage of training examples in terms of the number of models that predict correctly.

Dataset	# of models that predict correctly					Training Acc.
	1	2	3	4	5	
CIFAR-100	0.5	1.0	2.6	8.9	86.9	94.04
TinyImageNet	6.3	6.8	8.7	14.3	49.6	70.28

tribution of the training examples in terms of the number of models predicting correct answers, and the statistics demonstrates that many training examples in each dataset can take advantage of OD. In particular, the objectives of OD and KD are different for more than half of the training examples (*i.e.*, 50.4%) in TinyImageNet, and it implies that the accuracy gains in OD can be significant compared to KD. Indeed, Table 2 illustrates that the benefit of OD is more salient in TinyImageNet, which is natural because the dataset is more challenging and the trained models are less competitive.

Conclusion

We propose a novel framework of oracle knowledge distillation with neural architecture search especially designed for ensemble teacher models. The proposed framework addresses the capacity and complexity issues, and aims to find a desirable architecture with additional memory, which mimics oracle predictions of ensemble model more accurately. We empirically show that the combination of architecture search and oracle knowledge distillation successfully provides high-performance student models in both accuracy and memory size, and most of the searched networks achieve competitive performances to teacher models. We believe that our framework searching for an optimal student network is a promising research direction of knowledge distillation.

Acknowledgments We truly thank Tackgeun You for helpful discussion. This work was partly supported by Sam-

²https://github.com/bhheo/BSS_distillation

sung Advanced Institute of Technology (SAIT) and Institute for Information & Communications Technology Promotion (IITP) grant funded by the Korea government (MSIT) [2017-0-01778, 2017-0-01780].

References

- Abadi, M.; Barham, P.; Chen, J.; Chen, Z.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Irving, G.; Isard, M.; et al. 2016. Tensorflow: A System for Large-Scale Machine Learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, 265–283.
- Ashok, A.; Rhinehart, N.; Beainy, F.; and Kitani, K. M. 2018. N2N Learning: Network to Network Compression via Policy Gradient Reinforcement Learning. In *ICLR*.
- Breiman, L. 1996. Bagging Predictors. *Machine learning* 24(2):123–140.
- Chen, G.; Choi, W.; Chen, X.; Han, T. X.; and Chandraker, M. K. 2017. Learning Efficient Object Detection Models with Knowledge Distillation. In *NIPS*.
- Chen, L.-C.; Collins, M.; Zhu, Y.; Papandreou, G.; Zoph, B.; Schroff, F.; Adam, H.; and Shlens, J. 2018. Searching for Efficient Multi-Scale Architectures for Dense Image Prediction. In *NeurIPS*.
- Chen, Y.; Yang, T.; Zhang, X.; Meng, G.; Pan, C.; and Sun, J. 2019. DetNAS: Neural Architecture Search on Object Detection. *arXiv preprint arXiv:1903.10979*.
- Chen, T.; Goodfellow, I. J.; and Shlens, J. 2016. Net2Net: Accelerating Learning via Knowledge Transfer. In *ICLR*.
- Chen, Y.; Wang, N.; and Zhang, Z. 2018. DarkRank: Accelerating Deep Metric Learning via Cross Sample Similarities Transfer. In *AAAI*.
- Freund, Y.; Schapire, R.; and Abe, N. 1999. A Short Introduction to Boosting. *Journal-Japanese Society For Artificial Intelligence* 14(771-780):1612.
- Guzman-Rivera, A.; Batra, D.; and Kohli, P. 2012. Multiple Choice Learning: Learning to Produce Multiple Structured Outputs. In *NIPS*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *CVPR*.
- Heo, B.; Lee, M.; Yun, S.; and Choi, J. Y. 2019. Knowledge Distillation with Adversarial Samples Supporting Decision Boundary. In *AAAI*.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the Knowledge in a Neural Network. In *NIPS Deep Learning and Representation Learning Workshop*.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely Connected Convolutional Networks. In *CVPR*.
- Kim, J.; Park, S.; and Kwak, N. 2018. Paraphrasing Complex Network: Network Compression via Factor Transfer. In *NeurIPS*.
- Krizhevsky, A. 2009. Learning Multiple Layers of Features from Tiny Images. Technical report, Citeseer.
- lan, x.; Zhu, X.; and Gong, S. 2018. Knowledge Distillation by On-the-Fly Native Ensemble. In *NeurIPS*.
- Lee, S.; Purushwalkam, S.; Cogswell, M.; Crandall, D.; and Batra, D. 2015. Why M Heads Are Better Than One: Training A Diverse Ensemble of Deep Networks. *arXiv preprint arXiv:1511.06314*.
- Lee, S.; Prakash, S. P. S.; Cogswell, M.; Ranjan, V.; Crandall, D.; and Batra, D. 2016. Stochastic Multiple Choice Learning for Training Diverse Deep Ensembles. In *NIPS*.
- Lee, K.; Hwang, C.; Park, K.; and Shin, J. 2017. Confident Multiple Choice Learning. In *ICML*.
- Li, Z., and Hoiem, D. 2016. Learning Without Forgetting. In *ECCV*.
- Li, Q.; Jin, S.; and Yan, J. 2017. Mimicking Very Efficient Network for Object Detection. In *CVPR*.
- Liu, C.; Zoph, B.; Neumann, M.; Shlens, J.; Hua, W.; Li, L.-J.; Fei-Fei, L.; Yuille, A.; Huang, J.; and Murphy, K. 2018. Progressive Neural Architecture Search. In *ECCV*.
- Liu, C.; Chen, L.-C.; Schroff, F.; Adam, H.; Hua, W.; Yuille, A.; and Fei-Fei, L. 2019. Auto-Deeplab: Hierarchical Neural Architecture Search for Semantic Image Segmentation. In *CVPR*.
- Liu, H.; Simonyan, K.; and Yang, Y. 2018. DARTS: Differentiable Architecture Search. In *ICLR*.
- Luo, P.; Zhu, Z.; Liu, Z.; Wang, X.; and Tang, X. 2016. Face Model Compression by Distilling Knowledge from Neurons. In *AAAI*.
- Luo, R.; Tian, F.; Qin, T.; Chen, E.; and Liu, T.-Y. 2018. Neural Architecture Optimization. In *NeurIPS*.
- Mirzadeh, S.-I.; Farajtabar, M.; Li, A.; and Ghasemzadeh, H. 2019. Improved Knowledge Distillation via Teacher Assistant: Bridging the Gap Between Student and Teacher. *arXiv preprint arXiv:1902.03393*.
- Mun, J.; Lee, K.; Shin, J.; and Han, B. 2018. Learning to Specialize with Knowledge Distillation for Visual Question Answering. In *NeurIPS*.
- Noroozi, M.; Vinjimoor, A.; Favaro, P.; and Pirsaviash, H. 2018. Boosting Self-Supervised Learning via Knowledge Transfer. In *CVPR*.
- Park, W.; Kim, D.; Lu, Y.; and Cho, M. 2019. Relational Knowledge Distillation. In *CVPR*.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic Differentiation in PyTorch.
- Pham, H.; Guan, M.; Zoph, B.; Le, Q. V.; and Dean, J. 2018. Efficient Neural Architecture Search via Parameters Sharing. In *ICML*.
- Romero, A.; Ballas, N.; Kahou, S. E.; Chassang, A.; Gatta, C.; and Bengio, Y. 2015. FitNets: Hints for Thin Deep Nets. In *ICLR*.
- Simonyan, K., and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*.
- Sutskever, I.; Martens, J.; Dahl, G.; and Hinton, G. 2013. On the Importance of Initialization and Momentum in Deep Learning. In *ICML*.
- Tan, M.; Chen, B.; Pang, R.; Vasudevan, V.; Sandler, M.; Howard, A.; and Le, Q. V. 2019. MnasNet: Platform-Aware Neural Architecture Search for Mobile. In *CVPR*.
- Williams, R. J. 1992. Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning. *Machine learning* 8(3-4):229–256.
- Zagoruyko, S., and Komodakis, N. 2016. Wide residual networks. In *BMVC*.
- Zagoruyko, S., and Komodakis, N. 2017. Paying More Attention to Attention: Improving the Performance of Convolutional Neural Networks via Attention Transfer. In *ICLR*.
- Zhang, Y.; Xiang, T.; Hospedales, T. M.; and Lu, H. 2018. Deep Mutual Learning. In *CVPR*.
- Zoph, B., and Le, Q. 2017. Neural Architecture Search with Reinforcement Learning. In *ICLR*.