

On the Efficacy of Knowledge Distillation

Jang Hyun Cho
 Cornell University
 jc2926@cornell.edu

Bharath Hariharan
 Cornell University
 bh497@cornell.edu

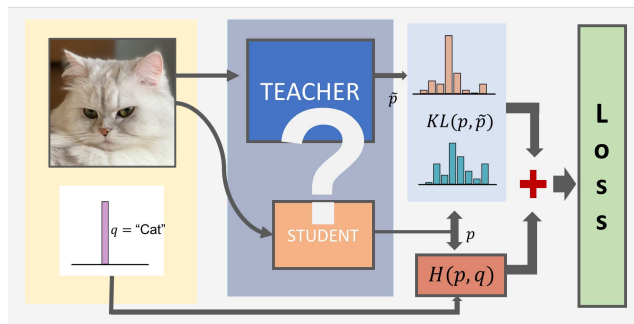


Figure 1. An illustration of standard knowledge distillation. Despite widespread use, an understanding of when the student can learn from the teacher is missing.

Abstract

In this paper, we present a thorough evaluation of the efficacy of knowledge distillation and its dependence on student and teacher architectures. Starting with the observation that more accurate teachers often don't make good teachers, we attempt to tease apart the factors that affect knowledge distillation performance. We find crucially that larger models do not often make better teachers. We show that this is a consequence of mismatched capacity, and that small students are unable to mimic large teachers. We find typical ways of circumventing this (such as performing a sequence of knowledge distillation steps) to be ineffective. Finally, we show that this effect can be mitigated by stopping the teacher's training early. Our results generalize across datasets and models.

1. Introduction

The past few years have seen dramatic improvements in visual recognition systems, but these improvements have been driven by deeper and larger convolutional network architectures. The large computational complexity of these architectures has limited their use in many downstream applications. As such, there has been a lot of recent research

on achieving the same or similar accuracy with smaller models. Some of this work has involved building more efficient neural network families [13, 7], pruning weights from larger neural networks [10], quantizing existing networks to use fewer bits for weights and activations [22] and distilling knowledge from larger networks into smaller ones [12, 1].

The last of these, knowledge distillation, is a general-purpose technique that at first glance is widely applicable and complements all other ways of compressing neural networks [20]. The key idea is to use soft probabilities (or 'logits') of a larger "teacher network" to supervise a smaller "student" network, in addition to the available class labels. These soft probabilities reveal more information than the class labels alone, and can purportedly help the student network learn better.

The appeal of this approach is in its apparent generality: *any* student can learn from *any* teacher. But does knowledge distillation fulfill this promise of generality? Unfortunately, in spite of the recent interest in variants of knowledge distillation [17, 24, 23, 21, 27, 26, 5, 14], an empirical answer to this question is missing. Prior experiments have typically looked at a small number of carefully chosen architectures, with the implicit assumption that conclusions will generalize across student or teacher architectures. However, there are a few isolated reports of failed experiments with knowledge distillation that suggest that this might not be true. For example, Zagoruyko and Komodakis observe that they are "unable to achieve positive results with knowledge distillation on ImageNet" [26]. What characterizes this, and other experiments where knowledge distillation does not seem to improve performance? Are there student-teacher combinations that perform better? And finally, is there something we can do to improve performance for other combinations?

In this paper, we seek to answer these questions. We find that in general, the teacher accuracy is a poor predictor of the student's performance. Larger teachers, though they are more accurate by themselves, do not necessarily make for better teachers. We explore the reasons for this and demonstrate that as the teacher grows in capacity and accuracy, the student often finds it difficult to emulate the teacher (resulting in a high KL divergence from the teacher logits even

during training). We show that this issue cannot be mitigated by solutions suggested in prior work, such as using a sequence of knowledge-distillation steps to increase the student accuracy. Finally, we find an effective solution to the problem: regularizing the teacher by stopping the training of the teacher early, and stopping knowledge distillation close to convergence to allow the student to fit the training loss better. Our solution is simple to implement and effective across the board at improving the efficacy of knowledge distillation.

The rest of the paper is organized as follows. After describing related work (Sec. 2), we first provide some background on knowledge distillation and attention transfer (Sec. 3). We then describe our experimental setting (Sec. 4). In Sec. 5 we discuss our findings and the empirical evidence for each.

2. Related Work

2.1. Knowledge distillation

The notion of training smaller, cheaper models (“students”) to mimic larger ones (“teachers”) is an old one, first described in a seminal paper on model compression by Bucilua et al [2]. This technique can be applied to deep neural networks almost out of the box [1, 12]. In this paper, we use the knowledge distillation framework described by Hinton et al. [12]. A brief description of knowledge distillation is provided in Section 3. The original paper on knowledge distillation experimented with the idea on a few small datasets, but a thorough empirical evaluation of knowledge distillation is missing.

Meanwhile, the focus of past work has been either on improving the quality of knowledge distillation or finding new applications for the idea. On the former direction, prior work has explored adding additional losses on intermediate feature maps of the student to bring them closer to those of the teacher [26, 15, 24]. Zhang et al. train a pair of models, distilling knowledge bidirectionally at every epoch [27]. Tarvainen et al. find that averaging consecutive student models over training steps tend to produce better performing students [21]. Yang et al. modify the loss function of teacher network to be more “tolerant” (that is, by adding more terms to make the model intentionally maintain high energy, benefiting from teacher’s misclassified logits) [23].

A particular approach to improving knowledge distillation is to perform knowledge distillation repetitively (we call it *sequential knowledge distillation* [23, 6, 17]). A particular way of using sequential knowledge distillation is as an alternative to ensembling to increase model accuracy [17, 6]. For example, Furlanello et al. [6] suggest training an ensemble of networks using a sequence of knowledge-distillation steps where a network uses its own previous version as a teacher. Interestingly, our results sug-

gest that this approach *underperforms* an ensemble trained from scratch, and furthermore, such sequential knowledge distillation reduces the ability of the network to act as a teacher. More generally, we find that these methods are highly dependent on the student capacity. In fact we find them ineffective in many cases particularly when student capacity is limited or the dataset is complex.

In terms of applications of knowledge distillation, prior work has found knowledge distillation to be useful for sequence modeling [16, 4], semi-supervised learning [21], domain adaptation [19], multi-modal learning [9] and so on. This wide applicability of the idea of knowledge distillation makes an exhaustive evaluation of knowledge distillation ideas even more important.

3. Background: Knowledge distillation

The key idea behind knowledge distillation is that the soft probabilities output by a trained “teacher” network contains a lot more information about a data point than just the class label. For example, if multiple classes are assigned high probabilities for an image, then that might mean that the image must lie close to a decision boundary between those classes. Forcing a student to mimic these probabilities should thus cause the student network to imbibe some of this knowledge that the teacher has discovered above and beyond the information in the training labels alone.

Concretely, given any input image x the teacher network produces a vector of scores $\mathbf{s}^t(x) = [s_1^t(x), s_2^t(x), \dots, s_K^t(x)]$ that are converted into probabilities: $p_k^t(x) = \frac{e^{s_k^t(x)}}{\sum_j e^{s_j^t(x)}}$. Trained neural networks produce peaky probability distributions, which may be less informative. Hinton et al [12] therefore propose to “soften” these probabilities using temperature scaling [8]:

$$\tilde{p}_k^t(x) = \frac{e^{s_k^t(x)/\tau}}{\sum_j e^{s_j^t(x)/\tau}} \quad (1)$$

where $\tau > 1$ is a hyperparameter.

A student similarly produces a softened class probability distribution, $\tilde{\mathbf{p}}^s(x)$. The loss for the student is then a linear combination of the typical cross entropy loss \mathcal{L}_{cls} and a knowledge distillation loss \mathcal{L}_{KD} :

$$\mathcal{L} = \alpha \mathcal{L}_{cls} + (1 - \alpha) \mathcal{L}_{KD}$$

where $\mathcal{L}_{KD} = -\tau^2 \sum_k \tilde{p}_k^t(x) \log \tilde{p}_k^s(x)$

α and τ are hyperparameters; popular choices are $\tau \in \{3, 4, 5\}$ and $\alpha = 0.9$ [26, 14, 17, 12].

4. Methods

We perform experiments on both CIFAR10 and ImageNet. In each case we keep the student the same and use

multiple teachers of varying capacity to perform knowledge distillation.

CIFAR10 For experiments on CIFAR10, we run each model for 200 epochs using SGD with momentum 0.9 and set the initial learning rate $\gamma = 0.1$, dropping 0.2 every 60 epochs. Standard data augmentation was applied to the dataset. For the hyperparameters regarding knowledge distillation, we stayed consistent with the popular choice ([12], [26]): Temperature $\tau = 4$, $\alpha = 0.9$, and $\beta = 1000$ for attention transfer. The same experiment was repeated 5 times and median, mean, and standard deviation are reported. We consider three different network architectures: ResNet [11], WideResNet [25], and DenseNet [13].

ImageNet For ImageNet experiments we followed Zagoruyko et al. [26] closely since it was the first successful work of knowledge distillation on ImageNet, to the best of our knowledge. We used SGD with nesterov momentum 0.9, initial learning rate $\gamma = 0.1$, weight decay 1×10^{-4} , and dropped learning rate by 0.1 every 30 epochs. As with CIFAR10, we set temperature $\tau = 4$, $\alpha = 0.9$, and $\beta = 1000$ for attention transfer. For ImageNet experiments, we consider ResNet [11].

5. Results

5.1. Bigger models are not better teachers

The idea behind knowledge distillation is that soft probabilities from a trained teacher reflect more about the data than the true label alone. One might expect that as the teacher becomes more accurate, these soft probabilities will capture the underlying class distribution better and thus serve as better supervision to the student. Thus, intuitively, we might expect that bigger, more accurate models might form better teachers.

We first evaluate if this is true on CIFAR10 dataset. In Figure 2, The **red** and **blue** lines shows the accuracy for different student networks trained from different teachers; the left plot varies the “depth” of the teacher while the right plot varies the “width”. From these experiments, we find that the hypothesis that bigger, more accurate models make better teachers is incorrect: although the teacher accuracy continues to rise as the teacher becomes larger (see supplementary for teacher accuracies), the student accuracy rises and then begins to fall. One might wonder if this is an artifact of the CIFAR dataset. We repeated the experiment on ImageNet, with ResNet18 as the student and ResNet18, ResNet34, ResNet50, and ResNet152 as teachers. The results are shown in Table 1. As can be seen, as the teacher becomes larger and more accurate, the student becomes *less accurate*.

Teacher	Teacher Error (%)	Student Error (%)
-	-	30.24
ResNet18	30.24	30.57
ResNet34	26.70	30.79
ResNet50	23.85	30.95

Table 1. Top-1 error rate for various teachers for a ResNet18 student on ImageNet. The first row corresponds to training from scratch.

Student	Teacher	KD Error (%Train)	KD Error (%Test)
WRN28-1	WRN28-3	0.23	4.05
	WRN28-4	0.25	4.53
	WRN28-6	0.23	4.54
	WRN28-8	0.31	4.81
WRN16-1	WRN16-3	1.70	6.32
	WRN16-4	1.69	6.52
	WRN16-6	1.94	6.91
	WRN16-8	1.69	7.01

Table 2. KD error on CIFAR10 for multiple teachers and students. The supplementary shows similar results from teachers with increasing depth.

What might be the reason for this decrease? One possibility is that as the teacher becomes both more confident and more accurate, the output probabilities start resembling more and more a one-hot encoding of the true label, and thus the information available to the student decreases. However, softening the probabilities with high temperature [12] did not change this result (detailed later in Figure 6, Table 10), invalidating this hypothesis. Below, we propose an alternative hypothesis.

5.2. Analyzing student and teacher capacity

There might be two reasons why a larger, more accurate teacher doesn’t lead to better student accuracy:

1. The student is able to mimic the teacher, but this does not improve accuracy. This would suggest a mismatch between the KD loss and the accuracy metric we care about.
2. The student is unable to mimic the teacher, suggesting a mismatch between student and teacher *capacities*.

We evaluated these hypotheses on CIFAR10 and ImageNet. In Table 2, we show the *KD error* for CIFAR: the fraction of examples for which the student and teacher predictions differ. Odd rows in Table 3 show the KD loss on ImageNet for a ResNet 18 student trained with different teachers. (We show KD error instead of KD loss on CIFAR because of scale issues caused by peaky output distributions).




Figure 2. The error plot of student networks distilled from different teachers on CIFAR10. WideResNet [25] 16-1 (left/red, right/red), 28-1 (right/blue), and DenseNet [13] 40-12 (left/blue) were used as the student networks. Increasing teacher capacity (depth: left, width: right) and thus accuracy does not necessarily increase the accuracy of the student network, indicating that the accuracy of the teacher network alone is not a valid metric to knowledge distillation.

In both cases, the KD Error/Loss is much higher for the largest teacher, which in turn leads to the least accurate student. This suggests that the student *is unable to mimic large teachers* and points to the second hypothesis, namely, the issue is one of mismatched capacity. We therefore posit that on both ImageNet and CIFAR, due to much lower capacity, the student is unable to find a solution in its space that corresponds well to the largest teacher.

5.3. Distillation adversely affects training

Note that knowledge distillation performs particularly poorly on ImageNet, where *all teachers* lead to lower student accuracy than a student trained from scratch (Table 1). While the previous section suggests that the student may not have enough capacity to match a very large teacher, it is still a mystery why *no teacher* improves accuracy on ImageNet. Despite multiple recent papers in knowledge distillation, experiments on ImageNet are rarely reported. The few that do report find that standard setting of knowledge distillation fails on ImageNet [26] or perform an experiment with a small portion of ImageNet [21]. But the reason for this has not been explored.

We dug deeper into the result. Figure 3 shows a comparison of validation accuracy plots between ResNet18 trained from scratch and using knowledge distillation with ResNet34. We find that while the KD loss improves validation accuracy initially, it begins to hurt accuracy towards the end of training.

We hypothesized that because ImageNet is a more challenging problem, the low-capacity student may be in the underfitting regime. The student may not have enough capacity to minimize both the training loss and the knowledge distillation loss, and might end up minimizing one loss (KD loss) at the expense of the other (cross entropy loss), espe-



Figure 3. ImageNet result of error plot of full knowledge distillation and training from scratch. In the figure the student is trained with ResNet34. Knowledge distillation helps initially but starts to hurt accuracy later in training. The same behavior occurs in the plots with different teachers (more plots in Supplementary).

cially towards the end of training.

This hypothesis suggests that we might want to stop the knowledge distillation early in the training process, and do gradient descent only on the cross-entropy loss for the rest of the training. We call this process “Early-stopped” knowledge distillation (“ESKD”) as opposed to standard knowledge distillation (“Full KD”).

Table 3 shows how this version compares to standard knowledge distillation, and also shows the loss values at the end of training. We find that the early-stopped version is better for all three teachers. We also find that, consistent with our hypothesis, the early-stopped version achieves a *lower* training cross-entropy loss and a *higher* KD loss than the baseline version suggesting that the latter models are indeed trading off one loss against the other. Note also that

Teacher	Top-1 Error (%, Test)	CE (Train)	KD (Train)	KD (Test)
ResNet18	30.57	0.146	2.916	3.358
ResNet18 (ES KD)	29.01	0.123	2.234	2.491
ResNet34	30.79	0.145	1.357	1.503
ResNet34 (ES KD)	29.16	0.123	2.359	2.582
ResNet50	30.95	0.146	1.553	1.721
ResNet50 (ES KD)	29.35	0.124	2.659	2.940

Table 3. Early-stopping the knowledge distillation can prevent the student from degrading its classification performance on ImageNet.

this simple trick of stopping knowledge distillation early now gives us the promised benefit of knowledge distillation: all the early-stopped students in Table 3 perform better than a model of similar architecture trained from scratch (30.24% accuracy).

However, early stopping does not change our original observation: that larger, more accurate teachers don’t result in more accurate students. Even with early stopping, we find that the KD loss on the test set *increases* with increasing teacher size, suggesting that the student is still struggling to mimic the teacher, and that it is indeed an issue of student capacity.

5.4. The efficacy of repeated knowledge distillation

If the difference between teacher and student capacities is very large, one possibility is to first distill from the large teacher to an intermediate teacher and then distill to the student, so that each knowledge distillation step has a better match between student and teacher capacity. This notion of sequential knowledge distillation has been proposed in the literature in other contexts. Recently Furlanello et al [6] attempted to train a sequence of models, with the i -th model in the sequence being trained with knowledge distillation with the $i - 1$ -th model as the teacher. They find that such sequential knowledge distillation may improve the performance compared to a model trained from scratch, and ensembling the sequence produces a better model.

We first test this claim on CIFAR with multiple networks and with both knowledge distillation and attention transfer (Table 4). We find that there are several caveats to Furlanello et al.’s result. First, for some models (ResNet 8 and ResNet 14), the last student in the sequence actually underperforms a student model trained from scratch. This suggests that the network architecture heavily determines the success of sequential knowledge distillation. Second, we find that although an ensemble of the student models from the entire sequence outperforms a *single* model trained from scratch, it does not outperform an ensemble of an equal number of models trained from scratch. This

might be because the students obtained through a sequence of knowledge distillation steps may be correlated with each other and therefore may not produce a strong ensemble.

If sequential knowledge distillation does indeed improve the accuracy of a model, a natural question to ask is if the resulting model forms a better teacher. To evaluate this, we conducted the following experiment. We chose WRN16-1 as the student model and WRN16-3 as the teacher (note that this is the optimal teacher for this student as suggested by Figure 2). We then trained the teacher using a sequence of 5 iterations of knowledge distillation. We compared the efficacy of this model as a teacher compared to a teacher trained from scratch. As shown in Table 5, a teacher trained with a sequence of knowledge distillation iterations, though more accurate, is not in fact a better teacher.

As discussed above, we might be interested in a variant of this idea where we first attempt to distill from a “large” model to a “medium” model, and then from the “medium” model to a “small” model. If this worked, it might help us avoid the issue of differing student and teacher capacities. We compare this step-wise knowledge distillation to directly distilling from the large model to the small model, or from the medium model to the small model. This might be a way to get around the effect we observed in Figure 2, where the larger models were not necessarily better teachers. We performed this experiment using WRN16-1 as the small model, WRN16-3 (the optimal teacher for WRN16-1) as the medium model and WRN16-8 as the large model. We find that such a step-wise distillation does not work: it performs almost exactly the same as directly using the large model for distillation with the small model (Table 6). Sequential distillation cannot help make large models better teachers.

We repeat some of these experiments on ImageNet and show the results in Table 7. We use early-stopping when performing knowledge distillation based on results from the previous section. Sequential knowledge distillation is in fact ineffective on ImageNet, too. The best result corresponds to a single knowledge distillation from the “small” model to another “small” model, where “small” is ResNet18, “Med.” is ResNet50, and “Large” is ResNet152. All these results suggest that despite the initial promise of sequential distillation, it is not a panacea and it especially does not help us use a large teacher to train a small student of significantly different capacity.

5.5. Early-stopped teachers make better teachers

In the previous section we have shown that sequential knowledge distillation is ineffective. This might be because it doesn’t address the core problem: the solution the large teacher has found is simply not in the solution space of the small student. The only solution is to find a teacher whose discovered solution is in fact reachable by the student.

Model	# Params	Method	Last Gen. Err.	All Gen. Ensemble Err.	Scratch Err.	Scratch Ensemble Err.
ResNet8	0.07M	AT+KD	13.469	12.786	12.569*	10.176
ResNet14	0.17M	AT+KD	9.226	8.653	9.078*	6.675
WRN16-2	0.69M	KD	6.101	5.181	6.428	4.865
WRN16-2	0.69M	AT+KD	5.696	5.310	6.418	5.003

Table 4. 5 generations of knowledge distillation were done, and the errors for the last generation of distillation sequence (“Last Gen. Err.”), ensemble of all generations (“All Gen. Ensemble Err.”), first generation (“Scratch Err.”), and the ensemble of the same number of scratch models (“Scratch Ensemble Err.”) were reported. Errors with (*) show the cases where repeating knowledge distillation even decreased the performance, and for all models support the claim that repeating knowledge distillation is ineffective.

Teacher Training	Teacher Error (%)	Student Error (%)
Scratch	5.34	7.61 (7.68 ± 0.259)
5 KD iterations	4.89	7.79 (7.67 ± 0.19)

Table 5. Sequential knowledge distillation does not make better teachers even when it improves accuracy. The student is WRN16-1, which achieves an error of (8.759 ± 0.129) when trained from scratch. The teacher is WRN16-3.

Training Procedure	Large Error (%)	Medium Error (%)	Small Error (%)
Large→Med.→ Small	4.41	4.80	8.04 (7.99 ± 0.24)
Med. → Small	-	5.34	7.614 (7.68 ± 0.26)
Large → Small	4.41	-	7.98 (8.03 ± 0.14)

Table 6. Using sequential knowledge distillation to distill from a large model (WRN16-8) to a medium model (WRN16-3), and from the latter to a small model (WRN16-1) does not help. The optimal approach still is to distill directly from the medium model to the small model, even though the teacher in this case has lower accuracy.

Training procedure	1st Teacher Error (%)	2nd Teacher Error (%)	Student Error (%)
Large → Small → Small	21.69	29.45	29.41
Med. → Small → Small	23.85	29.35	29.35
Small → Small → Small	30.24	29.01	29.15
Small → Small → Small [23]	-	-	30.12*
Small → Small × 5 [23]	-	-	29.60*

Table 7. Imagenet experiment of sequential early-stopped knowledge distillation (ESKD). “2nd Teacher” is first distilled from “1st Teacher”, and then “Student” is trained with the “2nd Teacher”. The last two lines compare with other variants of sequential knowledge distillation. [*] indicates that the number is inherited from the original paper.

Method	Teacher	Top-1 Error (%)
Scratch	-	30.24
Full KD [12]	ResNet18	30.57
Full KD [12]	ResNet34	30.79
Full KD [12]	ResNet50	30.95
Seq. Full KD [23]	3 Gen.	30.12*
Seq. Full KD [23]	6 Gen.	29.6*
KD+ONE [17]	3 Branches	29.45 ± 0.23*
Full KD + AT [26]	ResNet34	30.94
Full KD + AT [26]	ResNet34	29.3*
ESKD	ResNet18	29.01
ESKD	ResNet34	29.16
ESKD	ResNet50	29.35
ESKD	ResNet152	29.45
ESKD	ResNet34 (50)	29.02
ESKD	ResNet50 (35)	29.05
ESKD	ResNet152 (35)	29.26
Seq. ESKD	L → S → S	29.41
Seq. ESKD	M → S → S	29.35
Seq. ESKD	S → S → S	29.15
ESKD + AT	ResNet34	28.84
ESKD + AT	ResNet34 (50)	28.61

Table 8. Overall result of ImageNet experiments. ESKD: Early-stopped knowledge distillation, AT: Attention transfer [26]. The number inside the parentheses is the total number of epochs if teacher training is early-stopped.*Numbers reported in paper.

We may perform grid-search to find the optimal teacher network architecture, but that is too expensive. Instead, we propose to *regularize* the teacher when training it. In particular, we propose to stop the training of the large teacher early. There is some evidence that a large network trained with only a few epochs behaves as a small network, while still encompassing a greater search space than small network [3, 18]. This method is extremely simple and cheap, since only a third to fourth of the total number of epochs are needed. We evaluate the effectiveness of this idea in both CIFAR10 and ImageNet. Figure 4 plots the error rates vs. total epochs on CIFAR10, where the x-axis represents the total number of epochs each teacher is trained. The same hyperparameters as other CIFAR10 experiments are used for the training teacher, except the total number




Figure 4. CIFAR10 result to examine the effectiveness of the knowledge distillation with early-stopped teachers. For both student types (WRN16-1 and WRN28-1), there are clear “sweet spots” which optimize the performance of student network.

of epochs and the learning rate schedule. For training the teacher network, the learning rate is dropped by 0.2 every $\lfloor \frac{n-5}{3} \rfloor$ where n is the total number of epochs. We chose $n \in \{35, 50, 65, 80\}$. Notice that for both student models (WRN16-1 and WRN28-1), all early-stopped teachers produce better students than the optimal fully-trained teacher (WRN16-3 and WRN28-3).

Given these promising results, we next turn our attention to ImageNet. We choose $n \in \{35, 50\}$ and learning rate drop schedule of (15, 25, 30) for 35 and (20, 35, 45) for 50. Other hyperparameters and settings are the same with those of the previous ImageNet experiments. Table 8 shows results on ImageNet, where we also compare our results with prior results using knowledge distillation or its variants. Simply early-stopping the knowledge distillation with the largest, fully-trained teacher outperforms most prior work ($\approx 29.45\%$). Our best teachers are the early-stopped ResNet34 and fully-trained ResNet18, ($\approx 29.01\%$) which has ≈ 1.23 point performance gain over the model trained from scratch and $\approx 0.2\%$ from the best known result for this architecture from [17].

Table 8 also shows variants using attention transfer [26],

Method	Teacher	Top-1 Error (%)	Top-5 Error (%)
Scratch	-	47.38	18.51
ESKD	ResNet18	47.09	18.13
Full KD	ResNet34	47.86	18.61
ESKD	ResNet34 (50)	47.14	18.32
Full KD	ResNet50	47.92	18.72
ESKD	ResNet50 (35)	47.02	18.14
ESKD	ResNet152 (35)	47.25	18.25

Table 9. Each student network from ImageNet experiments is fine-tuned to Places-365 dataset for 12 epochs with initial learning rate $\gamma = 0.1$ and weight decaying 10^{-1} every 3 epochs.

Method	Teacher	Top-1 Error (%)	Top-5 Error (%)
Scratch	-	30.24	10.92
Full KD	ResNet18	30.75	11.11
Full KD	ResNet50	30.98	10.20
Full KD	ResNet152	31.27	11.59
ESKD	ResNet18	29.00	9.91
ESKD	ResNet50	29.00	9.76
ESKD	ResNet50 (35)	28.89	9.76

Table 10. Experiments with temperature $\tau = 20$ on ImageNet dataset. High temperature increased overall results for ESKD methods (**lower half**) whereas had no difference for Full KD methods (**upper half**).

an improvement over knowledge distillation. Early stopping of the teacher and of the student are both very compatible with attention transfer, leading to improvements of 1.6 points over the baseline and 0.7 points over the best numbers obtained with attention transfer [26].

5.6. Other factors impacting knowledge distillation

Different configurations In experiments above, we drew both the student and the teacher from the same model family. We now experiment with teachers and students drawn from other, possibly different, model families. Figure 5 shows various combinations of DenseNets and Wide ResNets as students and teachers. Our conclusions, both the inefficacy of knowledge distillation from large teachers and the benefits from early stopping, are apparent in these results.

Impact of α and τ Till now we have fixed the tradeoff between KD and cross entropy, $\alpha = 0.9$ and the temperature $\tau = 4$. Although the standard choice of the temperature is $\tau \in \{3, 4, 5\}$, one might wonder if our conclusions about early stopping are sensitive to these choices. As shown in Figure 6 we find that the early-stopped teacher consistently outperforms the fully-trained teacher across a range of these hyperparameter values on CIFAR10. We further investigate the high temperature case on ImageNet dataset (Table 10); we use $\tau = 20$. High temperature can theoretically mitigate




Figure 5. In all different student and teacher configurations, students trained from early-stopped teacher consistently outperforms those trained from regular teacher. DenseNet40-12 and WideResNet16-1 were used as student network and DenseNet and WideResNet models with varying width and depth were used as teachers (x-axis) [13, 25]. More results are in Supplementary.




Figure 6. CIFAR10 result to examine that the effectiveness of using early-stopped network as a teacher is consistent to different hyperparameter settings. WRN16-8 (ES) has the same/better error compared to the optimal teacher WRN16-3. WRN16-1 is chosen as the student network.

the peakiness of the teacher logits and may result better performance. We find that high temperature does increase the overall performance for early-stopped knowledge distillation (“ESKD”) but had no visible difference for full knowledge distillation (“Full KD”). The early stopped teacher still performed the best.

Generalizability for transfer learning Although we have seen variations in accuracies on ImageNet, a big aspect of convolutional networks is how well they transfer to other tasks. In the table 9 we examine whether the distilled network can be fine-tuned for classification on Places365 for a variety of students from the previous experiments. The results of transfer learning are consistent with the CIFAR and ImageNet experiments (full KD vs. early-stopped KD, small vs. large teachers, and regular vs. early-stopped teachers), proving that our findings also apply to transfer.

6. Conclusion

In this paper, we have presented an exhaustive study of the factors influencing knowledge distillation. Our key finding is that knowledge distillation is not a panacea and cannot succeed when student capacity is too low to successfully mimic the teacher. We have presented an approach to mitigate this issue by stopping teacher training early, to recover a solution more amenable for the student. Finally we have shown the benefits of this approach on CIFAR10 and ImageNet and also on transfer learning on Places365. We believe that further research into the nuances of distillation are necessary before it can succeed as a general and practical approach.

References

- [1] Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In *NIPS*, 2014.
- [2] Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2006.
- [3] Rich Caruana, Steve Lawrence, and C Lee Giles. Overfitting in neural nets: Backpropagation, conjugate gradient, and early stopping. In *Advances in neural information processing systems*, pages 402–408, 2001.
- [4] Yevgen Chebotar and Austin Waters. Distilling knowledge from ensembles of neural networks for speech recognition. In *Interspeech*, pages 3439–3443, 2016.
- [5] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 742–751. Curran Associates, Inc., 2017.
- [6] Tommaso Furlanello, Zachary C Lipton, Michael Tschanen, Laurent Itti, and Anima Anandkumar. Born again neural networks. In *ICML*, 2018.
- [7] Ariel Gordon, Elad Eban, Ofir Nachum, Bo Chen, Hao Wu, Tien-Ju Yang, and Edward Choi. Morphnet: Fast and simple resource-constrained structure learning of deep networks. In *CVPR*, June 2018.
- [8] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *ICML*, 2017.
- [9] Saurabh Gupta, Judy Hoffman, and Jitendra Malik. Cross modal distillation for supervision transfer. In *CVPR*, pages 2827–2836, 2016.
- [10] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In *ICLR*, 2016.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [12] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015.
- [13] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, 2017.
- [14] Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv preprint arXiv:1707.01219*, 2017.
- [15] Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv preprint arXiv:1707.01219*, 2017.
- [16] Yoon Kim and Alexander M. Rush. Sequence-level knowledge distillation. In *EMNLP*. Association for Computational Linguistics, 2016.
- [17] Xu Lan, Xiatian Zhu, and Shaogang Gong. Knowledge distillation by on-the-fly native ensemble. In *NIPS*, 2018.
- [18] Maren Mahsereci, Lukas Balles, Christoph Lassner, and Philipp Hennig. Early stopping without a validation set. *arXiv preprint arXiv:1703.09580*, 2017.
- [19] Zhong Meng, Jinyu Li, Yifan Gong, and Biing-Hwang Juang. Adversarial teacher-student learning for unsupervised domain adaptation. In *ICASSP*, 2018.
- [20] Asit Mishra and Debbie Marr. Apprentice: Using knowledge distillation techniques to improve low-precision network accuracy. In *ICLR*, 2018.
- [21] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NIPS*, pages 1195–1204, 2017.
- [22] Jiaxiang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, and Jian Cheng. Quantized convolutional neural networks for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4820–4828, 2016.
- [23] Chenglin Yang, Lingxi Xie, Siyuan Qiao, and Alan Yuille. Knowledge distillation in generations: More tolerant teachers educate better students. In *AAAI*, 2018.
- [24] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *CVPR*, 2017.
- [25] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016.
- [26] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2017.
- [27] Ying Zhang, Tao Xiang, Timothy M. Hospedales, and Huchuan Lu. Deep mutual learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.