GyF

This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Knowledge distillation: A good teacher is patient and consistent

Lucas Beyer^{*} Xiaohua Zhai^{*} Amélie Royer^{*†} Larisa Markeeva^{*‡} Rohan Anil Alexander Kolesnikov^{*} Google Research, Brain Team

{lbeyer,xzhai,akolesnikov}@google.com

Abstract

There is a growing discrepancy in computer vision between large-scale models that achieve state-of-the-art performance and models that are affordable in practical applications. In this paper we address this issue and significantly bridge the gap between these two types of models. Throughout our empirical investigation we do not aim to necessarily propose a new method, but strive to identify a robust and effective recipe for making state-of-the-art large scale models affordable in practice. We demonstrate that, when performed correctly, knowledge distillation can be a powerful tool for reducing the size of large models without compromising their performance. In particular, we uncover that there are certain implicit design choices, which may drastically affect the effectiveness of distillation. Our key contribution is the explicit identification of these design choices, which were not previously articulated in the literature. We back up our findings by a comprehensive empirical study, demonstrate compelling results on a wide range of vision datasets and, in particular, obtain a state-of-the-art ResNet-50 model for ImageNet, which achieves 82.8% top-1 accuracy.

1. Introduction

Large-scale vision models currently dominate many areas of computer vision. Recent state-of-the-art models for image classification [6, 22, 39, 41, 48], object detection [7, 26] or semantic segmentation [52] push model size to the limits allowed by modern hardware. Despite their impressive performance, these models are rarely used in practice due to high computational costs. Instead, practitioners typically use much smaller models, such as ResNet-50 [22] or MobileNet [14], which are order(s) of magnitude cheaper to run. According to the download counts of five BiT models



Figure 1. We demonstrate that distillation works the best when we train *patiently* for a large number of epochs and provide *consistent* image views to teacher and student models (green and blue lines). This can be contrasted to a popular setting of distilling with precomputed teacher targets (black line), which works much worse.

from *Tensorflow Hub*, the smallest ResNet-50 [11] model has been downloaded for significantly more times than the larger ones. As a result, many recent improvements in vision do not translate to real-world applications.

To address this problem, we concentrate on the following task: given a specific application and a large model that performs very well on it, we aim to compress the model to a smaller and more efficient architecture without compromising performance. There are two widely used paradigms that target this task: model pruning [18] and knowledge distillation [12]. *Model pruning* reduces the large model's size by stripping away its parts. This procedure can be restrictive in practice: first, it does not allow changing the model family,

^{*}equal contribution

[†] work done at Google, while being a PhD student in IST Austria.

[‡] work done at Google, while being a PhD student in Skoltech.

say from a ResNet to a MobileNet. Second, there may be architecture-dependent challenges, *e.g.* if the model uses group normalization [46], pruning channels may result in the need to dynamically re-balance channel groups.

Instead, we concentrate on the knowledge distillation approach which does not suffer from these drawbacks. The idea behind knowledge distillation is to "distill" a teacher model, in our case a large and cumbersome model or ensemble of models, into a small and efficient student model. This works by forcing the student's predictions (or internal activations) to match those of the *teacher*, thus naturally allowing a change in the model family as part of compression. We closely follow the original distillation setup from [12] and find it surprisingly effective when done right: We interpret distillation as a task of matching the *functions* implemented by the teacher and student, as illustrated in Figure 2. With this interpretation, we discover two principles of knowledge distillation for model compression. First, teacher and student should process the exact same input image views or, more specifically, same crop and augmentations. Second, we want the functions to match on a large number of support points to generalize well. Using an aggressive variant of mixup [51], we can generate support points outside the original image manifold. With this in mind, we experimentally demonstrate that consistent image views, aggressive augmentations and very long training schedules are the key to make model compression via knowledge distillation work well in practice.

Despite the apparent simplicity of our findings, there are multiple reasons that may commonly prevent researchers (and practitioners) from making the design choices that we suggest. **First**, it is tempting to precompute the teacher's activations for an image offline once to save compute, especially for very large teachers. As we will show, this fixed teacher approach does not work well. **Second**, knowledge distillation is also commonly used in different contexts (other than model compression), where authors recommend different or even opposite design choices [40,48,50], see Figure 2. **Third**, knowledge distillation requires an atypically large number of epochs to reach best performance, much more than commonly used for supervised training. Finally, choices which may look suboptimal in training of regular length often end up being best for long runs, and vice-versa.

In our empirical study, we mostly concentrate on compressing the large BiT-ResNet-152x2 from [22] that was pretrained on the ImageNet-21k dataset [36] and fine-tuned to the relevant datasets of interest. We distill it to a standard ResNet-50 architecture [11] (but replace batch normalization with group normalization) on a range of small and mid-sized datasets without compromising accuracy. We also achieve very strong results on the ImageNet [35] dataset: with a total number of 9600 epochs for distillation, we set the new ResNet-50 SOTA 82.8% on ImageNet. This is 4.4% better than the ResNet-50 model from [22], and 2.2% better than the best ResNet-50 model in the literature, which uses a more complex setup [37]. Finally, we demonstrate that our distillation recipe also works when simultaneously compressing and changing the model family, *e.g.* BiT-ResNet architecture to the MobileNet architecture.

2. Experimental setup

In this section, we introduce the experimental setup and benchmarks we use throughout the paper. Given a largescale vision model (the *teacher*, or T) with high accuracy on a particular task, we aim to compress this model to a much smaller model (the *student*, or S) without compromising its performance. Our compression recipe relies on knowledge distillation, as introduced in [12], and a careful investigation of several key ingredients in the training setup.

Datasets, metrics and evaluation protocol. We conduct experiments on five popular image classification datasets: *flowers102* [30], *pets* [32], *food101* [20], *sun397* [47] and *ILSVRC-2012* ("ImageNet") [35]. These datasets span diverse image classification scenarios; In particular, they vary in the number of classes, from 37 to 1000 classes, and total number of training images, from 1020 to 1281167 training images. This allows us to verify our distillation recipe for a broad range of practical settings and ensure its robustness.

As a metric, we always report classification accuracy. For all datasets, we perform design choices and hyperparameters selection using a *validation* split, and report final results on the *test* set. These splits are defined in the appendix E.

Teacher and student models. Throughout the paper, we opt for using pre-trained teacher models from BiT [22], which provides a large collection of ResNet models pre-trained on ILSVRC-2012 and ImageNet-21k datasets, with state-of-the-art accuracy. The only significant differences between BiT-ResNets and standard ResNets is their use of group normalization layer [46] and weight standardization [33], which are used instead of batch normalization [16].

In particular, we concentrate on the BiT-M-R152x2 architecture: a BiT-ResNet-152x2 (152 layers, 'x2' indicates the width multiplier) pretrained on ImageNet-21k. This model demonstrates excellent performance on a variety of vision benchmarks and it is still manageable to run extensive ablation studies with it. It is highly expensive to deploy (requires roughly 10x more compute than the standard ResNet-50), and thus effective compression of this model is of practical importance. For the student's architecture, we use a BiT-ResNet-50 variant, referred to as ResNet-50 for brevity.

Distillation loss. We use the KL-divergence between the teacher's p_t , and the student's p_s predicted class probability vectors as a distillation loss, as was originally introduced in [12]. We do not use any additional loss term with respect



Figure 2. Schematic illustrations of various design choices when doing knowledge distillation. Left: *Teacher* receives a fixed image, while *student* receives a random augmentation. Center-left: *Teacher* and *student* receive independent image augmentations. Center-right: *Teacher* and *student* receive consistent image augmentations. Right: *Teacher* and *student* receive consistent image augmentations plus the input image manifold is extended by including linear segments between pairs of images (known as *mixup* [51] augmentation).

to the original dataset's hard labels:

$$\mathrm{KL}(p_t || p_s) = \sum_{i \in \mathcal{C}} \left[-p_{t,i} \log p_{s,i} + p_{t,i} \log p_{t,i} \right], \quad (1)$$

where C is a set of classes. Also, as in [12], we introduce a temperature parameter T, which is used to adjust the entropy of the predicted softmax-probability distributions before they are used in the loss computation: $p_s \propto \exp(\frac{\log p_s}{T})$ and $p_t \propto \exp(\frac{\log p_t}{T})$.

Training setup. For optimization, we train our models with the Adam optimizer [21] with default parameters, except for the initial learning rate that is part of our hyperparameter exploration. We use a cosine learning rate schedule [27] without warm restarts. We also sweep over the weight decay loss coefficient for all our experiments (for which we use a "decoupled" weight decay convention [28]). To stabilize training we enable gradient clipping with a threshold of 1.0 on the global L2-norm of a gradient. Finally, we use batch size 512 for all our experiments, except for models trained on ImageNet, where we train with batch size 4096. For the remaining hyperparameters, we discuss their sweeping range together with corresponding experiments in the next section.

One additional important component of our recipe is the *mixup* data augmentation strategy [51]. In particular, we introduce a mixup variant in our "function matching" strategy (see Section 3.1.1), in which we use "agressive" mixing coefficients sampled uniformly from [0, 1], which can be seen as an extreme case of the originally proposed sampling from β -distribution.

Unless explicitly specified otherwise, for prepossessing we use an "inception-style" crop [38] and then resize images to a fixed square size. Furthermore, in order to make our extensive analysis computationally feasible (overall we trained dozens of thousands of models), we use relatively low input resolution and resize input images to 128×128 size, except for our ImageNet experiments, that use the standard input 224×224 resolution.

For all our experiments we use Google Cloud TPU accelerators [19]. We also report our batch sizes, epochs or total number of update steps, which allow to estimate resource requirements for any particular experiment of interest. Model code and weights are made publicly available².

3. Distillation for model compression

3.1. Investigating the "consistent and patient teacher" hypothesis

In this section, we perform an experimental verification of our hypothesis formulated in the introduction and visualised in Figure 2, that distillation works best when seen as function matching, *i.e.* when the student and teacher see consistent views of the input images, synthetically "filled" via mixup, and when *student is trained* using long training schedule (i.e. the "teacher" is patient).

To make sure that our findings are robust, we perform a very thorough analysis on four small and medium scale datasets, namely *Flowers102* [30] (1020 training images), *Pets* [32] (3312 training images), *Food101* [20] (about 68k training images), and *SUN397* [47] (76k training images).

In an effort to remove any confounding factors, for each individual distillation setting we sweep over all combinations of learning rates $\{0.0003, 0.001, 0.003, 0.01\}$, weight decays $\{1 \cdot 10^{-5}, 3 \cdot 10^{-5}, 1 \cdot 10^{-4}, 3 \cdot 10^{-4}, 1 \cdot 10^{-3}\}$, and

²https://github.com/google-research/big_transfer



Figure 3. Experimental validation of the "consistency" requirement on the *Flowers102* dataset. Colors match different knowledge distillation design choices as introduced in Figure 2 and Section 3.1.1. Note that while the fixed teacher settings achieve significantly lower distillation loss, they lead to students which do not generalize well. In contrast, consistent teaching and function matching approaches lead to significantly higher student performance. Similar results on more datasets are reported in Appendix C.

distillation temperatures $\{1, 2, 5, 10\}$. In all reported figures, we show every single run as a low opacity curve, and highlight the one with the best final validation accuracy. We provide corresponding test accuracies in Appendix A.

3.1.1 Importance of "consistent" teaching

First, we demonstrate that the consistency criterion, *i.e.* student and teacher seeing the same views, is the only way of performing distillation which reaches peak student performance across all datasets consistently. For this study, we define multiple distillation configurations which correspond to instantiations of all four options sketched in Figure 2, with the same color coding:

- Fixed teacher. We explore several options where the teacher's predictions are constant for a given image (precomputed target). The simplest (and worst) method is fix/rs, where the image is just resized to 224²px for both student and teacher. fix/cc follows a more common approach of using a fixed central crop for the teacher and a mild random crop for the student. fix/ic_ens is a heavy data augmentation approach where the teacher's prediction is the average of 1k inception crops, which we verified to improve the teacher's performance. The student also uses random inception crops. The two latter settings are similar to the input noise strategy from the "noisy student" paper [48].
- Independent noise. We instantiate this common strategy in two ways: ind/rc computes two independent mild random crops for the teacher and student respectively, while ind/ic uses the heavier inception crop instead. A similar setup was used in [40].
- **Consistent teaching.** In this approach, we randomly crop the image only once, either with mild random crop-

ping (same/rc) or heavy inception crop (same/ic), and use this same crop for the input to both the student and the teacher.

• Function matching. This approach extends consistent teaching, by expanding an input manifold of images through mixup (mix), and, again, providing consistent inputs to the student and the teacher. For brevity, we sometimes refer to this approach as "FunMatch".

Figure 3 shows 10 000 epoch training curves on *Flow-ers102* dataset in all of these configurations. These results clearly show that "consistency" is the key: all "inconsistent" distillation settings plateau at a lower score, while consistent settings increase student performance significantly, with the function matching approach working the best. Furthermore, the training losses show that, for such small datasets, using a fixed teacher leads to strong overfitting. In contrast, function matching never reaches such loss on the training set while generalizing much better to the validation set. Due to space constraints, we show analogous results for other datasets and training durations in Appendix C.

3.1.2 Importance of "patient" teaching

One can interpret distillation as a variant of supervised learning, where labels (potentially soft) are provided by a strong *teacher* model. This is especially true when the *teacher* predictions are (pre)computed for a single image view. This approach inherits all problems of the standard supervised learning, e.g. aggressive data augmentations may distort actual image label, while less aggressive augmentations may cause overfitting.

However, things change if we interpret distillation as function matching, and, crucially, make sure to provide consistent inputs to the *student* and *teacher*. In this case we can be very aggressive with image augmentations: even if an



Figure 4. One needs patience along with consistency when doing distillation. Eventually, the teacher will be matched; this is true across various datasets of different scale.

image view is too distorted, we still will make a progress towards matching the relevant functions on this input. Thus, we can be more opportunistic with augmentations and avoid overfitting by doing aggressive image augmentations and, if true, optimize for very long time until the *student*'s function comes close to the *teacher*'s.

We empirically confirm our intuition in Figure 4, where for each dataset we show the evolution of test accuracy during training of the best function matching student (according to validation), for different amounts of training epochs. The teacher is shown as a red line and is always reached eventually, after a much larger number of epochs than one would ever use in a supervised training setup. Crucially, there is no overfitting even when we optimize for a 1M epochs.

We also trained and tuned two more baselines for reference: training a ResNet-50 from scratch using the dataset original hard labels, as well as transferring a ResNet-50 that was pre-trained on ImageNet-21k. For both of these baselines, we heavily tune learning rate and weight decay as described in Section 3.1. The model trained from scratch using the original labels is substantially outperformed by our student. The transfer model fares much better, but is eventually also outperformed.

Notably, training for a relatively short but common duration of 100 epochs leads to much worse performance than the transfer baseline. Overall, the ResNet-50 student patiently and consistently matches the very strong but much more expensive ResNet-152x2 teacher across the board.

3.2. Scaling up to ImageNet

Based on our insights from the previous sections, we now investigate how the proposed distillation recipe scales to the widely used and more challenging ImageNet dataset [35]. Following the same protocol as before, in Figure 5 (left),

we report student accuracy curves throughout training for three distillation settings: (1) fixed teacher, (2) consistent teaching and (3) function matching. For reference, our base teacher model reaches a top-1 accuracy of 83.0%. Fixed teacher again suffers from long training schedules, and starts overfitting after 600 epochs. In contrast, the consistent teaching approaches continuously improves performance as the training duration increases. From this we can conclude that consistency is a key to make distillation work on ImageNet, similar to the behaviors on the previously discussed small and mid-sized datasets.

Compared to simple consistent teaching, function matching performs slightly worse with short schedules, which likely happens due to underfitting. But when we increase the length of training schedule, the improvement of function matching becomes apparent: for instance with only 1200 epochs, it is able to match the performance of consistent teaching at 4800 epochs, thus saving 75% compute resource. Finally, for the longest run of function matching we experimented on, the vanilla **ResNet-50** student architecture achieves **82.31%** top-1 accuracy on ImageNet.

3.3. Distilling across different input resolutions

So far, we have assumed that both the *student* and *teacher* receive the same standard input resolution of 224px. However, it is possible to pass images of different resolution to the student and the teacher, while still being consistent: one simply has to perform the crop on the original high-resolution image, and subsequently resize it differently for the student and the teacher: their views will be consistent, albeit at different resolutions. This insight can be leveraged for learning from a better, higher resolution, teacher [22, 42], but also for training a smaller, faster student [2]. We investigate both directions: first, following [2], we train a ResNet-50 student



Figure 5. Left: Top-1 accuracy on ImageNet of three distillation setups: (1) fixed teacher; (2) consistent teaching; (3) function matching ("FunMatch"). Light color curves show accuracy throughout training, while the solid scatter plots are the final results. The student with a fixed teacher eventually saturates and overfits to it. Both consistent teaching and function matching do not exhibit overfitting or saturation. Middle: Reducing the optimization cost, via *Shampoo* \mathbb{A} preconditioning; with 1200 epochs, it is able to match the baseline trained for 4800 epochs. Right: Initializing *student* with pre-trained weights improves short training runs, but harms for the longest schedules.

with an input resolution of 160px while leaving the teacher resolution unchanged (224px). This results in a twice faster model, which still achieves remarkable 80.49% top-1 accuracy (see Table 1), compared to the best published 78.8% at this resolution using an array of modifications [2].

Second, following [22], we distill a teacher that was finetuned at a resolution of 384px (and attains 83.7% top-1 accuracy), this time leaving the student resolution unchanged, *i.e.* consuming a 224px input image. Compared to the baseline teacher, this provides a modest but consistent improvement across the board, as shown in Table 1.

3.4. Optimization: A second order preconditioner (1) improves training efficiency

We observe that optimization efficacy creates a computational bottleneck for our distillation recipe with "function matching" perspective due to long training schedules. Intuitively, we believe that optimization difficulties stem from the fact that it is much harder to fit a general *function* with multivariate outputs, rather than fixed image-level labels. Thus, we conduct an initial exploration, whether more powerful optimizers can do a much better job at our task.

To this end, we change the underlying optimizer from Adam to Shampoo [1], with the second order preconditioner. In Figure 5 (middle) we observe that Shampoo achieves the same test accuracy reached by Adam at 4800 epochs in just 1200 epochs, and with minimal step time overhead. And, in general, we observe consistent improvement over Adam in all our experimental settings. Experimental details on the Shampoo optimizer are provided in the Appendix D.

3.5. Optimization: A good initialization improves short runs but eventually falls behind

Motivated by transfer learning literature [10, 22] and [37], where a good initialization is able to significantly shorten the training cost and achieve a better solution, we try to initialize the student model with a pre-trained BiT-M-ResNet50 weights and show the results in Figure 5 (right).

The BiT-M initialization improves more than 2% when the distillation duration is short. However, the gap closes when the training schedule is long enough. Our observation is similar to the conclusion of [10]. Starting from 1200 epochs, distilling from scratch matches the BiT-M initializated student, and slightly overtakes it for 4800 epochs.

3.6. Distilling across different model families

Going beyond using different input resolutions for student and teacher, nothing in principle prevents us from using architectures of different families entirely, as our consistent patient teacher approach still applies in this setting. This allows us to efficiently transfer knowledge from stronger and more complex teachers, e.g. ensembles, while keeping the simple architecture of a ResNet50 student, but also transfer the state-of-the-art performance of large ResNet models to more efficient architectures e.g. MobileNet. We demonstrate this via two experiments. First, we use an ensemble of two models as teacher and show that this further improves performance. Second, we train a MobileNet v3 [13] student and obtain the best reported MobileNet v3 model to date.

MobileNet student. We use MobileNet v3 (Large) as a





Figure 6. Distilling *pet* and *sun397* datasets on different data sources. Results indicate that distilling on completely unrelated images works to some extent, even though final results are relatively low. Distilling on "in-domain" data is the best and distilling on related/overlapping images can work reasonably well, but may require extra long training schedule.

Figure 7. Baseline ResNet-50 models trained from scratch with labels vs with the ResNet-152x2 teacher.

student, for most experiments we opt for the variant which uses GroupNorm (with the default of 8 groups) instead of BatchNorm. We do not use any of the training tricks used in the original paper,we simply perform function matching. Our student reaches 74.60% after 300 epochs, and 76.31% after 1200 epochs, resulting in the best published MobileNet v3 model. More results are in the Appendix A.

Ensemble teacher. We now try a better teacher: we create a model which consists of averaging the logits from our default teacher at 224px resolution, and our teacher at 384px resolution from the previous section. This is a different, though closely related, type of teacher which is significantly more powerful but also slower. This teacher's student is better than our default teacher's student at every duration we tried (Appendix A) and, after 9600 epochs, reaches a new state-of-the-art **top-1 ImageNet accuracy of 82.82%**.

3.7. Comparison to the results from literature.

Now, when we introduced our key experiments, we compare our best ResNet-50 models to the best ResNet-50 models available in the literature, see Table 2. In particular, for 224×224 input resolution we compare against the original ResNet-50 model from [11], BiT-M-ResNet-50 pretrained on ImageNet-21k dataset [36] and previous state-of-the-art model from [37]. For 160×160 input resolution we compare against very recent and competitive model from [2]. We observe that our distillation recipe leads the state-of-the-art performance in both cases and by a significant margin.

3.8. Distilling on the "out-of-domain" data

By looking at knowledge distillation as "function matching", one can draw a reasonable hypothesis that distillation can be done on arbitrary image inputs. In this section we investigate this hypothesis. We conduct experiments on *pets* and *sun397* datasets. We use our distillation recipe to distill *pets* and *sun397* models using out-of-domain images from the *food101* and *ImageNet* datasets and, for the reference results, also run distillation with "in-domain" images from *pets* and *sun397* datasets.

Figure 6 summarizes our results. First we observe that distilling using in-domain data works the best. Somewhat surprisingly, even if the images are completely unrelated, distillation still works to some extent, though results get worse. This, for example, means that the *student* model can learn to classify pets with roughly 30% accuracy by only seeing food images (softly) labeled as breeds of pets. Finally, if distillation images are somewhat related or overlapping with the actual "in-domain" images (e.g. Pets and ImageNet, or sun397 and ImageNet), then results can be as good (or almost as good) as using "in-domain" data, but extra long optimization schedule may be required.

3.9. Finetuning ResNet-50 with augmentations

To make sure that our observed state-of-the-art distillation results are not an artifact of our well-tuned training setup, namely very long schedule and aggressive mixup augmentations, we train corresponding baseline ResNet-50 models. More specifically, we reuse the distillation training setup for supervised training on ImageNet dataset without distillation loss. To further strengthen our baseline, we additionally try SGD optimizer with momentum, which is known to often work better for ImageNet than Adam optimizer.

Results are shown in Figure 7. We observe that training with labels and without distillation loss leads to significantly worse results and starts to overfit for long training schedules. Thus, we conclude that distillation is necessary to make our training recipe work well.

Experiment	300	1200	4800	9600
$T224 \rightarrow S224$	80.30	81.54	82.18	82.31
$T224 \rightarrow S160$	78.17	79.61	N/A	80.49
$T384 \rightarrow S224$	80.46	81.82	82.33	82.64

Table 1. Top-1 test accuracy for different teacher/student input resolutions (rows) and number of training epochs (columns).

Table 2. Comparison of our best and literature ResNet models. The metric is accuracy on ImageNet test split (officially *val* split).

Model	Arch.	Res.	Accuracy
"Revisiting ResNet" [2]	R50	160	78.8%
FunMatch (T224)	R50	160	80.5%
Original ResNet [11]	R50	224	77.2%
BiT-M-R50 [22]	R50	224	78.4%
Meal-v2 [37]	R50	224	80.7%
FunMatch (T384+224)	R50	224	82.8%
"Revisiting ResNet" [2]	R152	224	82.8%

4. Related work

There are many paradigms for compressing neural networks. One of them is **pruning**, where the general idea is to discard parts of the trained model while making it much more efficient and incurring little or no sacrifise in performance. Model pruning comes in many different flavours: it can be unstructured (*i.e.* focus on pruning individual connections) or structured (*i.e.* focus on pruning larger building blocks, e.g. whole channels). It can also come with or without an additional finetuning step, or be iterative or not. Balanced and fair discussion of this topic goes beyond the scope of this paper, so we refer interested reader to recent overviews as a starting point [3,44].

Knowledge distillation [12] is a technique for transferring knowledge from one model (teacher) to another (student), by optimizing a student model to match certain outputs (or intermediate activations) of a *teacher* model. This technique is used in numerous distinct contexts, such as semisupervised learning [40, 48] or even self-supervised learning [8]. In this paper we only consider knowledge distillation as a tool for model compression. The efficiency of distillation has been showcased in numerous works, e.g. [4, 34], under different depth/width patterns of the student and teacher architectures, and even combined with other compression techniques [29]. Notably, MEAL [37] proposes to distill an ensemble of large ResNet teachers to a smaller ResNet student with an adversarial loss and achieves strong results. The main difference of our work to similar works on knowledge distillation for compression, is that our method is simultaneously the simplest and best performing: we do not introduce any new components, but rather discover that correct training setup is sufficient to attain state-of-the art results.

Weights quantization [17, 25, 31, 45] and decomposition [5,9,23,43] aim to accelerate and reduce the memory footprint of CNNs by replacing large matrices operations with their lightweight approximations. This line of research is largely orthogonal to this work and can generally be combined with the method from this paper, especially during for the final model deployment stage. We leave exploration of this topic for future research.

Finally, there is a line of work, which approaches our goal (compact and high performing models) from a different angle, by focusing on **altering the architecture** and getting good compact models trained from scratch, so there is no need to compress large models. Some notable examples include ResNeXt [49], Squeeze-and-Excitation Networks [15] and Selective Kernel [24], which propose modifications that improve model accuracy for a fixed compute budget. These improvements are complementary to the research question tackled in this paper and can be compounded.

5. Conclusion

Instead of proposing a new method for model compression, we closely look at the existing common knowledge distillation process and identify how to make it work really well in the context of model compression. Our key findings stem from a specific interpretation of knowledge distillation: we propose to see it as a function matching task. This is not the typical view of knowledge distillation, as normally it is seen as "a strong *teacher* generates better (soft) labels that are useful for training a better and smaller *student* model".

Based on our interpretation we simultaneously incorporate three ingredients: (i) make sure that *teacher* and *student* always get identical inputs, including noise, (ii) introduce aggressive data augmentations to enrich the input image manifold (through mixup) and (iii) use very long training schedules. Even though each component of our recipe may seem trivial, our experiments show that one has to apply all of them jointly to get top results.

We attain very strong empirical results for compressing very large models to the more practical ResNet-50 architecture. We believe that they are very useful from a practical point of view and are a very strong baseline for future research on compressing large-scale models.

Acknowledgements. We thank Daniel Keysers and Frances Hubis for their valuable feedback on this paper; Ilya Tolstikhin and the Google Brain team at large for providing a supportive research environment.

References

- Rohan Anil, Vineet Gupta, Tomer Koren, Kevin Regan, and Yoram Singer. Scalable second order optimization for deep learning. arXiv preprint arXiv:2002.09018, 2021. 6, 11
- [2] Irwan Bello, William Fedus, Xianzhi Du, Ekin D Cubuk, Aravind Srinivas, Tsung-Yi Lin, Jonathon Shlens, and Barret Zoph. Revisiting resnets: Improved training and scaling strategies. arXiv preprint arXiv:2103.07579, 2021. 5, 6, 7, 8
- [3] 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. 8
- [4] Jang Hyun Cho and Bharath Hariharan. On the efficacy of knowledge distillation. In *International Conference on Computer Vision (ICCV)*, 2019. 8
- [5] Emily Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2014. 8
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Computer Vision* (*ICCV*), 2021. 1
- [7] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. arXiv preprint arXiv:2012.07177, 2020. 1
- [8] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. arXiv preprint arXiv:2006.07733, 2020. 8
- [9] Julia Gusak, Maksym Kholiavchenko, Evgeny Ponomarev, Larisa Markeeva, Philip Blagoveschensky, Andrzej Cichocki, and Ivan Oseledets. Automated multi-stage compression of neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019. 8
- [10] Kaiming He, Ross Girshick, and Piotr Dollár. Rethinking imagenet pre-training. arXiv preprint arXiv:1811.08883, 2018.
 6
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Conference* on Computer Vision and Pattern Recognition (CVPR), 2016. 1, 2, 7, 8
- [12] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NeurIPS Deep Learning and Representation Learning Workshop*, 2015. 1, 2, 3, 8
- [13] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 6

- [14] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861, 2017. 1
- [15] J. Hu, L. Shen, and G. Sun. Squeeze-and-excitation networks. In Conference on Computer Vision and Pattern Recognition (CVPR), pages 7132–7141, 2018. 8
- [16] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learing* (*ICML*), 2015. 2
- [17] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Conference* on Computer Vision and Pattern Recognition (CVPR), pages 2704–2713, 2018. 8
- [18] Steven A. Janowsky. Pruning versus clipping in neural networks. *Phys. Rev. A*, 39, 1989. 1
- [19] Norman P Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, et al. In-datacenter performance analysis of a tensor processing unit. In *Proceedings* of the 44th annual international symposium on computer architecture, 2017. 3
- [20] Parneet Kaur, Karan Sikka, and Ajay Divakaran. Combining weakly and webly supervised learning for classifying food images. arXiv preprint arXiv:1712.08730, 2017. 2, 3
- [21] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 3
- [22] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *International Conference on Computer Vision (ICCV)*, 2020. 1, 2, 5, 6, 8
- [23] Vadim Lebedev, Yaroslav Ganin, Maksim Rakhuba, Ivan Oseledets, and Victor Lempitsky. Speeding-up convolutional neural networks using fine-tuned cp-decomposition. In *International Conference on Learning Representations (ICLR)*, 2014. 8
- [24] Xiang Li, Wenhai Wang, Xiaolin Hu, and Jian Yang. Selective kernel networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [25] Darryl Lin, Sachin Talathi, and Sreekanth Annapureddy. Fixed point quantization of deep convolutional networks. In *International Conference on Machine Learing (ICML)*, pages 2849–2858, 2016. 8
- [26] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. arXiv preprint arXiv:2103.14030, 2021. 1
- [27] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations (ICLR)*, 2017. 3

- [28] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019. 3
- [29] A. K. Mishra and Debbie Marr. Apprentice: Using knowledge distillation techniques to improve low-precision network accuracy. In *International Conference on Learning Representations (ICLR)*, 2018. 8
- [30] Maria-Elena Nilsback and Andrew Zisserman. Delving deeper into the whorl of flower segmentation. In *British Machine Vision Conference (BMVC)*, 2007. 2, 3
- [31] Eunhyeok Park, Junwhan Ahn, and Sungjoo Yoo. Weightedentropy-based quantization for deep neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5456–5464, 2017. 8
- [32] Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 2, 3
- [33] Siyuan Qiao, Huiyu Wang, Chenxi Liu, Wei Shen, and Alan Yuille. Micro-batch training with batch-channel normalization and weight standardization. arXiv preprint arXiv:1903.10520, 2019. 2
- [34] 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. 8
- [35] 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), 2015. 2, 5
- [36] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 2015. 2, 7
- [37] Zhiqiang Shen and Marios Savvides. MEAL V2: Boosting vanilla Resnet-50 to 80%+ top-1 accuracy on ImageNet without tricks. *arXiv preprint arXiv:2009.08453*, 2020. 2, 6, 7, 8
- [38] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 3
- [39] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learing (ICML)*, 2019. 1
- [40] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2017. 2, 4,8
- [41] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic,

and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision. *arXiv preprint arXiv:2105.01601*, 2021. 1

- [42] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. Fixing the train-test resolution discrepancy. In *Confer*ence on Neural Information Processing Systems (NeurIPS), 2019. 5
- [43] Chaoqi Wang, Roger Grosse, Sanja Fidler, and Guodong Zhang. Eigendamage: Structured pruning in the kroneckerfactored eigenbasis. In *International Conference on Machine Learing (ICML)*, pages 6566–6575, 2019. 8
- [44] Huan Wang, Can Qin, Yulun Zhang, and Yun Fu. Emerging paradigms of neural network pruning. arXiv preprint arXiv:2103.06460, 2021. 8
- [45] Hao Wu, Patrick Judd, Xiaojie Zhang, Mikhail Isaev, and Paulius Micikevicius. Integer quantization for deep learning inference: Principles and empirical evaluation. arXiv preprint arXiv:2004.09602, 2020. 8
- [46] Yuxin Wu and Kaiming He. Group normalization. In European Conference on Computer Vision (ECCV), 2018. 2
- [47] Jianxiong Xiao, James Hays, Krista A. Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *Conference on Computer Vision* and Pattern Recognition (CVPR), 2010. 2, 3
- [48] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10687–10698, 2020. 1, 2, 4, 8
- [49] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *Confer*ence on Computer Vision and Pattern Recognition (CVPR), 2017. 8
- [50] Sangdoo Yun, Seong Joon Oh, Byeongho Heo, Dongyoon Han, Junsuk Choe, and Sanghyuk Chun. Re-labeling imagenet: from single to multi-labels, from global to localized labels. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [51] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations (ICLR)*, 2018. 2, 3
- [52] Barret Zoph, Golnaz Ghiasi, Tsung-Yi Lin, Yin Cui, Hanxiao Liu, Ekin D Cubuk, and Quoc V Le. Rethinking pre-training and self-training. *Conference on Neural Information Processing Systems (NeurIPS)*, 2020. 1