

When Ensembling Smaller Models is More Efficient than Single Large Models

Dan Kondratyuk, Mingxing Tan, Matthew Brown, Boqing Gong
Google AI

{dankondratyuk, tanmingxing, mtbr, bgong}@google.com

Abstract

*Ensembling is a simple and popular technique for boosting evaluation performance by training multiple models (e.g., with different initializations) and aggregating their predictions. This approach is commonly reserved for the largest models, as it is commonly held that increasing the model size provides a more substantial reduction in error than ensembling smaller models. However, we show results from experiments on CIFAR-10 and ImageNet that ensembles can outperform single models with both higher accuracy and requiring fewer total FLOPs to compute, even when those individual models' weights and hyperparameters are highly optimized. Furthermore, this gap in improvement widens as models become large. This presents an interesting observation that output diversity in ensembling can often be more efficient than training larger models, especially when the models approach the size of what their dataset can foster. Instead of using the common practice of tuning a single large model, one can use **ensembles as a more flexible trade-off between a model's inference speed and accuracy**. This also potentially eases hardware design, e.g., an easier way to parallelize the model across multiple workers for real-time or distributed inference.*

1. Introduction

Neural network ensembles are a popular technique to boost the performance of a model's metrics with minimal effort. The most common approach in current literature involves training a neural architecture on the same dataset with different random initializations and averaging their output activations [4]. This is known as *ensemble averaging*, or a simple type of *committee machine*. For instance, on image classification on the ImageNet dataset, one can typically expect a 1-2% top-1 accuracy improvement when ensembling two models this way, as demonstrated by AlexNet [6]. Evidence suggests averaging ensembles works because each model will make some errors independent of one another due to the high variance inherent in neural networks with millions of parameters [3, 9, 2].

For ensembles with more than two models, accuracy can increase further, but with diminishing returns. As such, this technique is typically used in the final stages of model tuning on the largest available model architectures to slightly increase the best evaluation metrics. However, this method can be regarded as impractical for production use-cases that are under latency and size constraints, as it greatly increases computational cost for a modest reduction in error.

One may expect that increasing the number of parameters in a single network should result in higher evaluation performance than an ensemble of the same number of parameters or FLOPs, at least for models that do not overfit too heavily. After all, the ensemble network will have less connectivity than the corresponding single network. But we show cases where there is evidence to the contrary.

In this paper, we show that we can consistently find averaged ensembles of networks with fewer FLOPs and yet higher accuracy than single models with the same underlying architecture. This is true even for families of networks that are highly optimized in terms of its accuracy to FLOPs ratio. We also show how this gap widens as the number of parameters and FLOPs increase. We demonstrate this trend with a family of ResNets on CIFAR-10 [13] and EfficientNets on ImageNet [12].

The results of this finding imply that a large model, especially a model that is so large and begins to overfit to a dataset, can be replaced with an ensemble of a smaller version of the model for both higher accuracy and fewer FLOPs. This can result in faster training and inference with minimal changes to an existing model architecture. Moreover, as an added benefit, the individual models in the ensemble can be distributed to multiple workers which can speed up inference even more and potentially ease the design of specialized hardware.

Lastly, we experiment with this finding by varying the architectures of the models in ensemble averaging using neural architecture search to study if it can learn more diverse information associated with each model architecture. Our experiments show that, surprisingly, we are unable to improve over the baseline approach of duplicating the same architecture in the ensemble in this manner. Several factors

could be attributed to this, including the choice of search space, architectural features, and reward function. With this in mind, either more advanced methods are necessary to provide gains based on architecture, or it is the case that finding optimal single models would be more suitable for reducing errors and FLOPs than searching for different architectures in one ensemble.

2. Approaches and Experiments

For our experiments, we train and evaluate convolutional neural networks for image classification at various model sizes and ensemble them. When ensembling, we train the same model architecture independently with random initializations, produce softmax predictions from each model, and calculate a geometric mean¹ μ across the model predictions. For n models, we ensemble them by

$$\mu = (y_1 y_2 \dots y_n)^{\frac{1}{n}} \quad (1)$$

where the multiplication is element-wise for each prediction vector y_i .

We split our evaluation into two main experiments and a third follow-up experiment.

2.1. Image Classification on CIFAR-10

For the first experiment, we train wide residual networks on the CIFAR-10 dataset [13, 5]. We train and evaluate the Wide ResNets at various width and depth scales to examine the relationship between classification accuracy and FLOPs and compare them with the ensembled versions of each of those models. We train 8 models for each scale and ensemble them as described. We select a depth parameter of $n = 16$, increase the model width scales $k \in \{1, 2, 4, 8\}$, and provide the corresponding FLOPs on images with a 32x32 resolution. We use a standard training setup for each model as outlined in [13].

Note that we use smaller models than typically used (e.g., Wide ResNet 28-10) to show that our findings can work on smaller models that are less prone to overfitting.

2.2. Image Classification on ImageNet

To further show that the ensemble behavior as described can scale to larger datasets and more sophisticated models, we apply a similar experiment using EfficientNets on ImageNet [12, 10]. EfficientNet provides a family of models using compound scaling on the network width, network depth, and image resolution, producing models from b0 to b7. We adopt the first five of these for our experiments, training and ensembling up to three of the same model architecture on ImageNet and evaluating on the validation set. We use the

¹Since the softmax applies a transformation in log-space, a geometric mean respects the relationship. We notice slightly improved ensemble accuracy when compared to an arithmetic mean.

original training code and hyperparameters as provided by [12] for each model size with no additional modifications.

3. Results

In this section, we plot the relationship between accuracy and FLOPs for each ensembled model. In cases of single models that are not ensembled, we plot the median accuracy. We observe that the standard deviation of the evaluation accuracy of each model architecture size never exceeds 0.1%, so we exclude it from the results for readability. For models that are ensembled, we vary the number of n trained models and choose the models randomly.

For the first experiment on CIFAR-10, Figure 1 plots a comparison of Wide ResNets with a depth parameter of $n_d = 16$ and width scales $k \in \{1, 2, 4, 8\}$. For clarity in presentation, we show a smaller subset of all the networks we trained. For each network (e.g., “wide resnet 16-8”, which stands for the depth parameter of $n_d = 16$ and the width scale of $k = 8$), we vary the number of models $n \in \{1, 2, \dots, 8\}$ in an ensemble and label it alongside the curve.

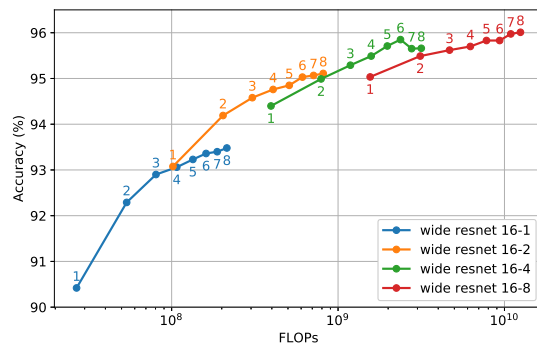


Figure 1. Test accuracy vs. model FLOPs (log-scale) when ensembling models trained on CIFAR-10. Each curve indicates the ensembles of increasing widths for a Wide ResNet n_d-k with a depth of $n = 16$. We show the number of models in each ensemble next to each point.

In the second experiment on ImageNet, Figure 2 plots a comparison of EfficientNets b0 to b5. Notably, we re-train all models using the current official EfficientNet code², but unlike the original paper that uses AutoAugment, here we do not use any specialized augmentation like AutoAugment or RandAugment to better observe the effects of overfitting.

4. Discussion

We draw the following observations from Figures 1 and 2 and particularly highlight the intriguing trade-off between

²The EfficientNet code can be found at <https://github.com/tensorflow/tpu/tree/master/models/official/efficientnet>




Figure 2. Validation accuracy vs. model FLOPs (log-scale) when ensembling models trained on ImageNet. Each curve indicates the ensembles of increasing sizes for a given EfficientNet. We show the number of models in each ensemble next to each point.

accuracy and FLOPs thanks to the ensembling.

First, we can see with no surprise that across the board, as the number of FLOPs increase for a single model, so too does the accuracy. This is also true of the ensembles which essentially multiply the base FLOPs by n for n models.

What is more interesting is that the results show that there can be cases where **ensembles of models with fewer collective FLOPs can achieve higher accuracy than a single larger model**. This is indicated by points that are above and to the left of other points. For instance, an ensemble of eight Wide ResNet 16-2 models achieves the same accuracy of 95% as a much wider Wide ResNet 16-8 at a fraction of the FLOPs (80M vs. 150M). An added benefit is that ensembles can easily be distributed to multiple workers to speed up computation even more.

Increasing the number of models in an ensemble will eventually be hit with diminishing returns, resulting a crossover point where an ensemble of the next largest model provides a better trade-off in terms of accuracy to FLOPs. In CIFAR-10, we observe the optimal ensemble size would be 2-4 models before the accuracy improvement slows down.

Finally, an interesting trend is that for smaller models, we can see that ensembling them has a more difficult time improving over larger single models. But as the models become larger, becoming increasingly likely to be over-parameterized and overfit to the dataset, we can see how ensembling provides a bigger accuracy boost over even larger models. For instance, the ensembles of EfficientNet-b0 do not come close to reaching the same accuracy to FLOPs trade-off as EfficientNet-b1. However, as the models become increasingly large, we see that the ensemble of two EfficientNet-b3 models achieves higher accuracy with fewer FLOPs over EfficientNet-b4, hence a better trade-off than EfficientNet-b4 provides.

Despite EfficientNet’s scaling ability producing highly optimized models, we can still see gaps in performance

where ensembles can perform better under the same number of total FLOPs, especially as the model size grows from b3 onwards. In other words, **ensembling offers an alternative and more effective scaling method than the compound scaling in EfficientNet** when some application scenarios permit the ensembling.

5. Neural Architecture Search (NAS) for Diverse Ensembles

Having noted the observations above, we hypothesize that ensembles can be improved further by varying the architectures of each model in an ensemble rather than duplicating the same architecture. The idea is that different architectures will naturally provide alternative features and therefore may enhance ensemble diversity. This should, in turn, provide improved accuracy at no increase to the number of FLOPs.

5.1. NAS Experiment Setup

To test this hypothesis, we adopt the same NAS framework as MnasNet [11]. We use a search space predicting model depth, width, and convolution type. We also augment the search space to include varying input resolution scales $r \in \{112, 168, 196, 224\}$. As a result, each model provides $m = 50$ hyperparameters to search. Additionally, we expand this to a joint search space to search for an ensemble of models by multiplying the search space n times, one for each model, for a total of nm hyperparameters. Each model is trained individually and ensembled as described in earlier experiments.

We alter the reward function to be penalized by not the total latency of the ensemble, but the maximum latency of all of the models in an ensemble and simulate this latency on a Pixel Phone 1. Assuming that each model can be run in parallel on separate workers, this would require the search to optimize the largest model in the ensemble at any given point to reduce the likelihood of producing ensembles where one model is large and the rest are anemic. Lastly, we train each searched model for 10 epochs before evaluating the accuracy, which is part of the reward, on a held-out set.

5.2. NAS Results

We show the Pareto curves of the ensemble accuracy with respect to model maximum latency across ensembles of size one, two, and three in Figure 3. This plot demonstrates the inherent trade-off between model accuracy and computation speed, with the best models being in the outer edge of the point cloud.

Results show that one-, two-, and three-model ensembles are surprisingly close to one another. The skyline two-model ensembles tend to beat out single models, but only by 1% at best. Skyline three-model ensembles show nearly

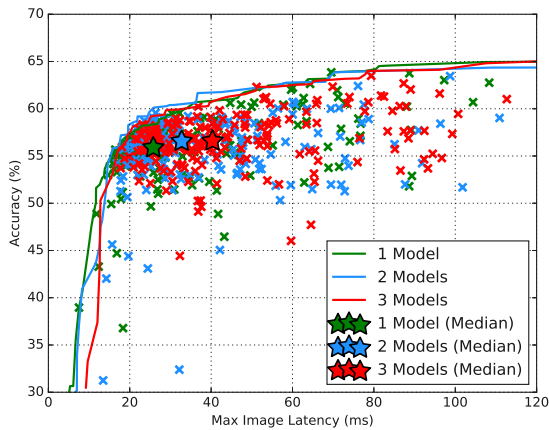


Figure 3. The resulting Pareto curves when searching for architectures across different ensemble sizes (trained for 10 epochs). We indicate the median ensemble accuracy and max latency of each ensemble size as stars in the plot.

identical performance to single models. We see that the median model accuracy does increase as the ensemble size grows, but at the cost of increased maximum latency.

Out of the searched diverse models, we pick the most promising candidates for a target latency. When trained to convergence, we find that two-model and three-model ensembles perform just as well as single models (assuming roughly equal max image latency). Somehow frustratingly, we find that simply duplicating the best single model for a given latency target and ensembling them together provides the best improvement in accuracy.

This experiment presents evidence towards a conclusion that ensembles benefit the most from choosing the most accurate models and not models that are architecturally diverse, at least under our current NAS context. For a fixed computational budget, this corresponds to using the best model architecture across the ensemble. We of course caution that we only have tested this with a simple NAS setup on a single large image classification dataset. This could change with a noisier and smaller dataset, or with more stringent constraints on model losses, regularization, or architectural mechanisms.

6. Related Work

Model ensembling has a long history with many different proposed techniques. Most works in this area come before advancements in deep learning were popularized. For instance, [7] define different subsets of the training data and use cross-validation to divide data into different groups. [1] developed bagging, where a different training set is given for different models to promote diversified feature learning. And [8] is one of the earliest attempts at constructing en-

sembles with different models by changing the number of hidden nodes in each network.

7. Conclusion

We have demonstrated how averaging ensembles can result in higher accuracy with fewer FLOPs than popular single models on image classification. This provides an interesting insight that smaller models can stand to provide great benefit without sacrificing on the accuracy to efficiency trade-offs of larger models. We advocate further inspections into the trade-off of ensembling especially for the applications where distributed inference is plausible.

References

- [1] Leo Breiman. Bagging predictors. *Machine learning*, 24(2):123–140, 1996.
- [2] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [3] Lars Kai Hansen and Peter Salamon. Neural network ensembles. *IEEE transactions on pattern analysis and machine intelligence*, 12(10):993–1001, 1990.
- [4] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [5] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [6] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [7] Anders Krogh and Jesper Vedelsby. Neural network ensembles, cross validation, and active learning. In *Advances in neural information processing systems*, pages 231–238, 1995.
- [8] Derek Partridge and William B Yates. Engineering multiversion neural-net systems. *Neural computation*, 8(4):869–893, 1996.
- [9] Michael P Perrone and Leon N Cooper. When networks disagree: Ensemble methods for hybrid neural networks. Technical report, BROWN UNIV PROVIDENCE RI INST FOR BRAIN AND NEURAL SYSTEMS, 1992.
- [10] 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*, 115(3):211–252, 2015.
- [11] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2820–2828, 2019.
- [12] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *ICML 2019. arXiv preprint arXiv:1905.11946*, 2019.
- [13] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.