
FreezeOut: Accelerate Training by Progressively Freezing Layers

Andrew Brock, Theodore Lim, & J.M. Ritchie
School of Engineering and Physical Sciences
Heriot-Watt University
Edinburgh, UK
{ajb5, t.lim, j.m.ritchie}@hw.ac.uk

Nick Weston
Renishaw plc
Research Ave, North
Edinburgh, UK
Nick.Weston@renishaw.com

Abstract

The early layers of a deep neural net have the fewest parameters, but take up the most computation. In this extended abstract, we propose to only train the hidden layers for a set portion of the training run, freezing them out one-by-one and excluding them from the backward pass. Through experiments on CIFAR, we empirically demonstrate that FreezeOut yields savings of up to 20% wall-clock time during training with 3% loss in accuracy for DenseNets, a 20% speedup without loss of accuracy for ResNets, and no improvement for VGG networks.

1 Introduction

Layer-wise pre-training of neural nets [1] [13] has largely been replaced by careful initialization strategies and fully end-to-end training. Successful techniques such as DropOut [11] and Stochastic Depth [4], however, suggest that it is not necessary to have every unit in a network participate in the training process at every training step. Stochastic Depth leverages this to both regularize and reduce computational costs by dropping whole layers at a time (though it requires the use of residual connections [3]), while DropOut's use of masks does not by default result in a computational speedup.

In this work, we are concerned with reducing the time required for training a network by only training each layer for a set portion of the training schedule, progressively "freezing out" layers and excluding them from the backward pass. This technique is motivated by the observation that in many deep architectures, the early layers take up most of the compute budget, but have the fewest parameters and tend to converge to fairly simple configurations (e.g. edge detectors), suggesting that they do not require as much fine tuning as the later layers, where most of the parameters reside. The same idea has been proposed by [9], where it is motivated by the observation that the representations learned by deep nets can be analytically shown to converge from the bottom up.

2 FreezeOut

We propose a simple modification to the standard backprop+SGD pipeline to reduce training time. FreezeOut employs cosine annealing (as proposed by [7]) without restarts (as used by [2]) with a layer-wise schedule, where the first layer's learning rate is reduced to zero partway through training (at t_0), and each subsequent layer's learning rate is annealed to zero some set time thereafter. Once a layer's learning rate reaches zero, we put it in inference mode and exclude it from all future backward passes, resulting in an immediate per-iteration speedup proportional to the computational cost of the layer.

In the simplest version of FreezeOut, each layer L_i starts with a single fixed learning rate α , which anneals to zero at t_i , where t_i is linearly spaced between a user-selected t_0 and the total number of

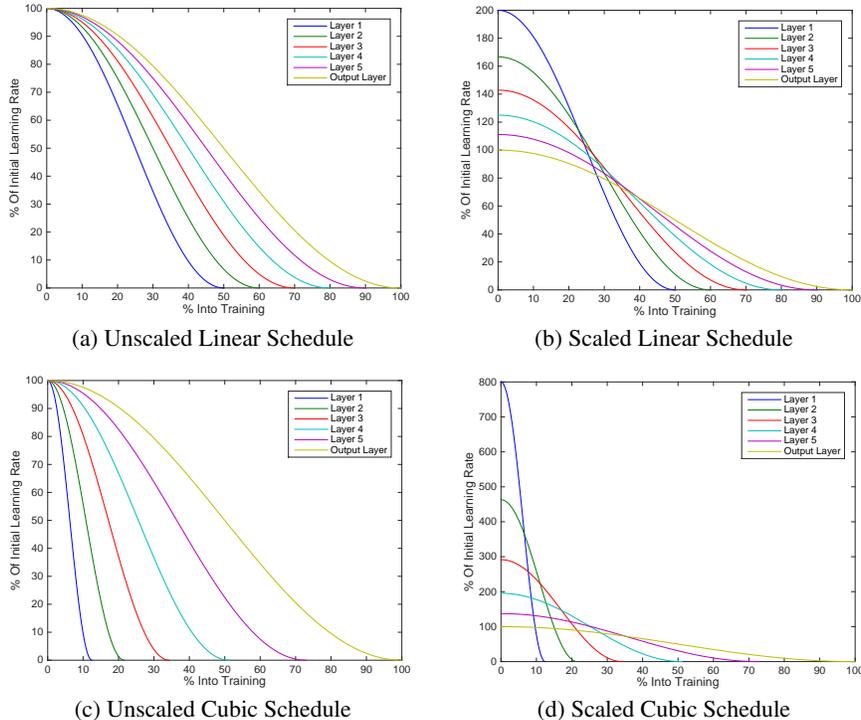


Figure 1: Per-Layer Learning Rate Schedules for a 5-hidden-layer network with $t_0 = 0.5$.

iterations, as shown in Figure 1(a). Each layer’s learning rate at iteration t is thus given as

$$\alpha_i(t) = 0.5\alpha_i(0)(1 + \cos(\pi t/t_i)) \tag{1}$$

We experiment with varying two aspects of this strategy. First, we consider scaling the initial layer-wise learning rate to be $\alpha_i(0) = \alpha/t_i$, where α is the base learning rate (and the learning rate of the final layer). This scaling, shown in Figure 1 (b), causes each layer’s learning curve to integrate to the same value, meaning that each layer travels the same distance in the weight space (modulo its observed gradients and weight dimension) despite the reduced number of steps taken.

Second, we vary the strategy for how t_0 relates to the t_i s of the remaining layers. The simple version we explore takes the t_i values determined by the linear scheduling rule, and cubes them with $t_i^{(cubed)} = t_i^{(linear)^3}$, as shown in Figure 1 (c) and (d). This gives more priority (in terms of training time) to later layers than does the linear schedule. As with the linear schedule, we consider an "unscaled" variant where the α_i values are identical, and a "scaled" variant where the α_i values are scaled based on the cubed t_i values. Throughout this paper, we refer to t_i s with respect to their uncubed values, such that a user selected $t_0 = 0.5$ results in a cubed $t_0 = 0.125$.

FreezeOut adds two user decisions—the choice of t_0 and the choice of FreezeOut strategy—to the standard choices of initial learning rate and number of training iterations. In the following section, we empirically investigate the relative merit of each of the four strategies, and provide recommendations for a default configuration, reducing the need for user-tuning of the hyperparameters. FreezeOut is easy to implement in a dynamic graph framework, requiring approximately 15 unique lines of code in PyTorch [8].

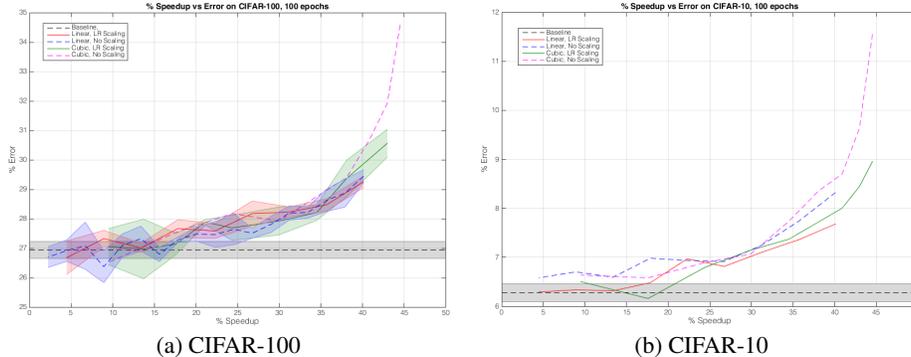


Figure 2: FreezeOut results for $k=12$, $L=76$ DenseNets on CIFAR-100 for 100 epochs. Shaded areas represent one standard deviation from the mean across 2-5 training runs.

3 Experiments

We modify publicly available PyTorch [8] implementations of DenseNets [5]¹, Wide ResNets [14]², and VGG [10] and test each of the four scheduling strategies across a wide range of t_0 values on CIFAR-100 and CIFAR-10.

3.1 DenseNet Experiments

Unless otherwise noted, all DenseNet experiments were performed using DenseNet-BC models with a growth rate of 12 and a depth of 76. We make use of standard data augmentation, train using SGD with Nesterov Momentum [12] and report results on the test set after training is completed. We compare achieved speedups (in terms of observed wall clock time relative to a non-FreezeOut baseline) to performance on the test set, and are primarily interested in the accuracy reduction incurred for a given reduction in training time.

For the speedups presented here, we use the values attained by running each test sweep in a controlled environment, on a GTX1080Ti with no other programs running. The actual speedups we observe in practice are slightly different, as we run the experiments on shared servers, with various types of GPUs, and with many other programs also running.

We found that a quick back-of-the-envelope calculation based on the reduction in operations accurately estimates the obtained speedups. Given c_i , the computational cost for the forward pass through each convolutional layer, and noting that the cost of a full forward-backward pass through that layer is $2c_i$, we calculate the baseline computational cost as the sum of each layer’s cost:

$$C = \Sigma(2c_i \times n_{itr}) \tag{2}$$

with n_{itr} the total number of training iterations. The compute cost for training with FreezeOut is:

$$C_f = \Sigma((1 + t_i) \times c_i \times n_{itr}) \tag{3}$$

and the approximate relative speedup is simply the ratio $1 - C_f/C$. We found that this accurately estimates speedup for cubic scheduling, but underestimates the speedups (e.g. the actual speedup is greater than the calculation gives) for linear scheduling; multiplying the predicted speedup by a correction factor of 1.3 resolves this almost entirely. This estimate neglects computational overhead and the cost of other operations (nonlinearities and BatchNorm), but this seems to be offset by the fact that a frozen layer’s BatchNorm costs are reduced (no longer having to calculate means and variances), and we find that this is generally a reliable estimate of the achieved speedups for a given setting.

Our most-investigated setup is shown in Figure 2(a), where we train on CIFAR-100 for 100 epochs and repeat each experimental setting 2 to 5 times. We also train for a single pass on CIFAR-10 for 100 epochs (Figure 2(b)).

¹<https://github.com/bamos/densenet.pytorch>

²<https://github.com/xternalz/WideResNet-pytorch>

We observe a clear speedup versus accuracy tradeoff. For every strategy, we observe a speedup of up to 20%, with a maximum relative 3% increase in test error. Lower speedup levels perform better and occasionally outperform the baseline, though given the inherent level of non-determinism in training a network, we consider this margin insignificant.

Whether this tradeoff is acceptable is up to the user. If one is prototyping many different designs and simply wants to observe how they rank relative to one another, then employing higher levels of FreezeOut may be tenable. If, however, one has set one’s network design and hyperparameters and simply wants to maximize performance on a test set, then a reduction in training time is likely of no value, and FreezeOut is not a desirable technique to use.

Based on these experiments, we recommend a default strategy of cubic scheduling with learning rate scaling, using a t_0 value of 0.8 before cubing (so $t_0 = 0.5120$) for maximizing speed while remaining within an envelope of 3% relative error. As a close alternative, we suggest linear scheduling without learning rate scaling, using $t_0 = 0.5$. The user is, of course, free to select parameters to fit a particular point on the design curve.

3.2 Wide ResNet experiments

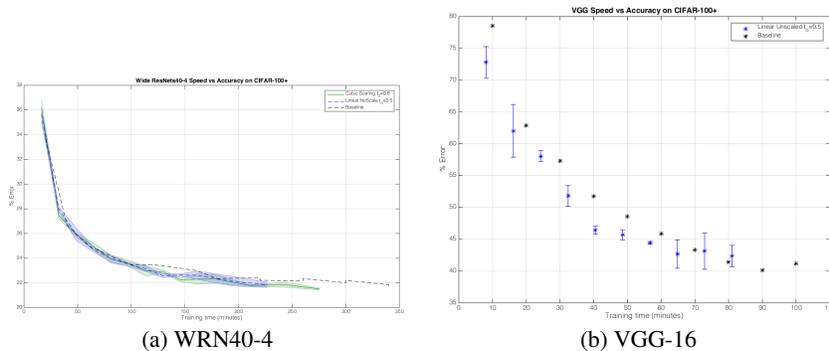


Figure 3: FreezeOut results for WRN40-4 and VGG on CIFAR-100. Shaded areas represent one standard deviation from the mean across 3 training runs for Cubic Scaled and 4 training runs for Linear Unscaled.

We next investigate the feasibility of FreezeOut for use in Residual architectures [3]. We train Wide ResNets [14] with depth of 40 and widening factor of 4, and compare our two recommended strategies against a no-FreezeOut baseline. We vary the number of epochs for which we train, and investigate how FreezeOut performs against a baseline of simply training for fewer epochs (e.g. a network trained with FreezeOut for 100 epochs has approximately the same training time as a network trained without FreezeOut for 80 epochs). The results of this investigation are shown in Figure 3(a). FreezeOut appears to be better suited to Wide ResNets than DenseNets, achieving *higher* accuracy even when trained for the same number of epochs (the final point on the Cubic Scaled curve and the final point on the Baseline curve were both trained for the same number of iterations).

3.3 VGG Experiments

Finally, we investigate the use of FreezeOut on an architecture without skip connections. We employ a variant of the VGG-16 [10] architecture with Batch Normalization [6], no Dropout [11], and 512 units in each of the fully connected layers rather than 4096. We perform a more limited set of experiments using the Linear Unscaled strategy against a baseline without FreezeOut, as shown in Figure 3(b). FreezeOut appears to be less well-suited for use in VGG, suggesting that skip connections (either dense or residual) may be an important element enabling FreezeOut to work for the other architectures we investigate.

4 Conclusion

In this extended abstract, we presented FreezeOut, a simple technique to accelerate neural network training by progressively freezing hidden layers.

References

- [1] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle. Greedy layer-wise training of deep networks. In *NIPS*, 2007.
- [2] X. Gastaldi. Shake-shake regularization of 3-branch residual networks. ICLR 2017 Workshop, 2017.
- [3] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In *ECCV 2016*.
- [4] G. Huang, Y. Sun, Z. Liu, D. Sedra, and K. Q. Weinberger. Deep networks with stochastic depth. In *ECCV 2016*.
- [5] G. Huang, Z. Liu, K.Q. Weinberger, and L. van der Maaten. Densely connected convolutional networks. In *CVPR 2017*, 2017.
- [6] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML 2015*.
- [7] I. Loshchilov and F. Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR 2017*, 2017.
- [8] A. Paszke, S. Gross, and S. Chintala. Pytorch. github.com/pytorch/pytorch, 2017.
- [9] M. Raghu, J. Gilmer, J. Yosinski, and J. Sohl-Dickstein. Svcca: Singular vector canonical correlation analysis for deep understanding and improvement. In *NIPS 2017*, 2017.
- [10] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR 2015*.
- [11] N. Srivastava, G.E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. arXiv Preprint arXiv: 1207.0580, 2012.
- [12] I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. In *ICML 2013*.
- [13] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. In *JMLR 2010*.
- [14] S. Zagoruyko and N. Komodakis. Wide residual networks. arXiv Preprint arXiv: 1605.07146, 2016.