
Sparse Transfer Learning via Winning Lottery Tickets

Rahul S. Mehta*
Palantir Technologies
rahulm@palantir.com

Abstract

The recently proposed Lottery Ticket Hypothesis of Frankle and Carbin (2019) suggests that the performance of over-parameterized deep networks is due to the random initialization seeding the network with a small fraction of favorable weights. These weights retain their dominant status throughout training – in a very real sense, this sub-network “won the lottery” during initialization. The authors find sub-networks via unstructured magnitude pruning with 85-95% of parameters removed that train to the same accuracy as the original network at a similar speed, which they call *winning tickets*. In this paper, we extend the Lottery Ticket Hypothesis to a variety of transfer learning tasks. We show that sparse sub-networks with approximately 90-95% of weights removed achieve (and often exceed) the accuracy of the original dense network in several realistic settings. We experimentally validate this by transferring the sparse representation found via pruning on CIFAR-10 to SmallNORB and FashionMNIST for object recognition tasks.

1 Introduction

Deep neural networks (DNNs) achieve state-of-the-art performance for a wide array of machine learning and pattern recognition problems. Conventional wisdom has long held that increasing the depth of a neural network improves its expressive power. This overparameterization also seems to *improve* the generalization ability of deep networks, which is at odds with traditional machine learning methods. Despite these benefits, increased depth also complicates optimization; deep networks encounter issues such as convergence failures, and are often prohibitively expensive in resource-constrained settings. Conversely, techniques for removing unnecessary weights from networks such as pruning (Le Cun et al., 1990; Han et al., 2015; Li et al., 2016) can eliminate up to 85-90% of parameters with no effect, but still require training the full model.

Recently, Frankle and Carbin (2019) proposed the *lottery ticket hypothesis* (LTH), which states that there exists a sparse, *trainable* sub-network within a dense feed-forward network that reaches the same validation accuracy as the original network when trained from scratch.

The Lottery Ticket Hypothesis A randomly-initialized, dense feed-forward neural network contains a subnetwork that is initialized such that – when trained in isolation – it can match the test accuracy of the original network after training for at most the same number of iterations.

Formally, consider a dense feed-forward network $f(x; \theta)$ with initial parameters $\theta_0 \sim \mathcal{D}_\theta$. Over the course of optimization, f reaches minimum validation loss ℓ at iteration j , with a test accuracy α .

*Palantir Technologies, 15 Little W. 12th Street, New York, NY, 10004. This research was supported by Princeton University, Amazon Research, and the Arora Group. This work was completed while the author was a student at Princeton University.

Also consider the same network f re-trained from $f(\mathbf{x}; m \odot \theta_0)$, for some $m \in \{0, 1\}^{|\theta|}$ (it is clear that as long as $\|m\|_0 > 0$, $\|m \odot \theta\|_0 < |\theta|$), and suppose that it reaches its minimum validation loss ℓ' at iteration j' with test accuracy α' .² The Lottery Ticket Hypothesis formally states that $\exists m \in \{0, 1\}^{|\theta|}$ such that $j' \leq j$ (*commensurate training time*) and $\alpha' \geq \alpha$ (*commensurate accuracy*). We call the sub-network, defined by m , a *winning ticket*. All winning tickets can be found by simple unstructured magnitude pruning (Han et al., 2015). We note that while this procedure still requires that the dense network be fully-trained, the sparse sub-network can be re-trained from scratch to reach the same test accuracy or even *exceed* that of the original network in the same number of training steps.

Subsequently, Frankle et al. (2019) found that such winning tickets can be scaled up to large visual recognition datasets such as ImageNet. However, this line of research still leaves one important question open; does the structure of winning tickets depend on the r they were trained on? Frankle and Carbin (2019) allude to this possibility when discussing the implications of their conjecture, and hypothesize that winning tickets may be transferrable between tasks.

Contributions This paper focuses on investigating whether or not the sparse representation found via unstructured magnitude pruning can be trained in isolation on a *distinct* target task. In particular, we show the following results:³

- We pose the *Ticket Transfer Hypothesis*, modifying the original LTH for transfer learning problems. We formalize this statement in Conjecture 2.1. In doing so, we resolve several inconsistencies between the original statement of the conjecture and the transfer learning setting.
- We show that sparse representations with as few as 5% of parameters remaining found via pruning on CIFAR-10 can be transferred to SmallNORB and FashionMNIST when only the dense feed-forward layers are fine-tuned on the target task (which mirrors many realistic transfer learning scenarios).
- We repeat the above experiments, this time fine-tuning the *entire* network; we find similar winning tickets with up to 95% of parameters removed that reach or exceed the performance of the original network. We also investigate the efficacy of the *iterative* and *one-shot* pruning methods introduced by Han et al. (2015).

2 Revisiting the Lottery Ticket Hypothesis

The original Lottery Ticket Hypothesis is proposed in the context of a single learning task \mathcal{T} ; the dense network is trained to completion on the task, and a sparse sub-network is found via iterative pruning. Then, the sparse sub-network is reset to the original weights and retrained in isolation on \mathcal{T} .

In contrast, in the *transfer learning* setting, the representation learned by a network on a *source* task \mathcal{T}_S is leveraged when fine-tuning the model on a distinct *target* task \mathcal{T}_T . For instance, if we wish to classify a target dataset with relatively few labeled examples, one approach is to transfer the weights learned from a more general task such as object detection on ImageNet, and then fine-tune the representation on the target dataset. In this setting, it is natural to ask if winning tickets found via iterative pruning on a source task \mathcal{T}_S can reach commensurate accuracy when fine-tuned in isolation on a target task \mathcal{T}_T .

Recall that the original statement of the Lottery Ticket Hypothesis specifies that the network can be reset to its original initialization $f(\mathbf{x}; m \odot \theta_0)$ after m is found via pruning. However, this does not make sense in a transfer learning context; if the main motivation is to transfer a learned feature representation from a large dataset such as ImageNet to a task with few labelled examples, resetting the network to $f(\mathbf{x}; m \odot \theta_0)$ eliminates any information learned from the source task; in a sense, the “winning ticket” initialization should be equivalent to a random initialization $f(\mathbf{x}; m \odot \theta')$ for some $\theta' \sim \mathcal{D}_\theta$. Moreover, we can expect that re-initializing the network to any other early weights $\theta_{j'}$ ($j' \ll j$) will adversely affect performance, and prevent the sub-network from reaching the original network’s accuracy. This should be especially evident in the case of only fine-tuning the

²In this context, \odot denotes the element-wise (Hadamard) product.

³We make our source code available at <https://github.com/rahulsmehta/sparsity-experiments>.

fully-connected layers while freezing the convolution layers (recall that j is the iteration of minimum validation loss), since this is essentially equivalent to using a fixed random feature embedding throughout training.

Accordingly, we first informally state our adaption of the Lottery Ticket Hypothesis for transfer learning, which we call the *Ticket Transfer Hypothesis*. We conveniently ignore the matter of initialization below for our informal description of the conjecture; it will be easier to rigorously define this in Conjecture 2.1.

Ticket Transfer Hypothesis (Informal) Given a source and target task (i.e. classification), a randomly-initialized, dense feed-forward neural network contains a subnetwork that can be found while training on the source task such that – when fine-tuned in isolation on a target task – it can match the test accuracy of the original network fine-tuned on the target task in as many training iterations. These sub-networks can always be found with unstructured magnitude pruning.

The major distinction between our variant and the original formulation of the LTH is that the sparse sub-network must be determined using only the *source* dataset; moreover, we require that the sub-network reaches the same accuracy in the same number of training setps when fine-tuned on the target task as the original dense network.

In order to address the matter of initialization, we adapt the concept of *late resetting*. In this approach, instead of re-initializing f to $m \odot \theta_0$, f is reset to some $m \odot \theta_{j'}$ ($j' \ll j$). Frankle et al. (2019) show that late resetting produces winning tickets that can overcome several stability issues and scale to large-scale tasks such as ImageNet. Moreover, in the classic transfer learning setting, the final fitted representation is always the initialization for a pre-trained network fine-tuned on a downstream task. Thus, we consider this to be the analog of late resetting in the transfer learning setting.

We make this intuition precise below; in particular, we resolve the issue of initialization by relaxing the condition that $f(\mathbf{x}; m \odot \theta_0)$ must reach commensurate accuracy on the target task. Consider source and target tasks \mathcal{T}_S and \mathcal{T}_T . Let $f(\mathbf{x}, \theta)$ be a dense, feed-forward network initialized to $\theta = \theta_0 \sim \mathcal{D}_\theta$. Let j denote the iteration at which f reaches its minimum validation loss on \mathcal{T}_S . Further, let $\mathcal{P}_S = \{\theta_i\}_{i=0}^j$ denote the set of parameters through the j th training iteration in the source task \mathcal{T}_S , and let θ_S denote the parameters at the iteration of minimum validation loss. Moreover, suppose that $f(\mathbf{x}; \theta_S)$ reaches its minimum validation loss ℓ and test accuracy α on \mathcal{T}_T at iteration k .

Conjecture 2.1 (Ticket Transfer Hypothesis). *Suppose that $f(\mathbf{x}; m \odot \theta^*)$ reaches minimum validation loss ℓ' at iteration k' with test accuracy α' on \mathcal{T}_T . Then there exists $m \in \{0, 1\}^{|\theta|}$, $\theta^* \in \mathcal{P}_S$ such that $\|m\|_0 \ll |\theta|$, $k' \leq k$ (commensurate training time), and $\alpha' \geq \alpha$ (commensurate accuracy).*

We address the issue of initialization by enabling our sparse sub-network to be reset to *any* weights that were encountered during the optimization trajectory of the model. This can be viewed as a generalization of a “winning ticket” initialization – since we would like to transfer the representation to another task, we are satisfied if we can achieve commensurate accuracy with *any* initialization found during training on the source task \mathcal{T}_S . In practice, we consider three initializations; θ_0 (the “winning ticket” initialization), θ_S (“late reset”), and $\theta' \sim \mathcal{D}_\theta$ (random).

3 Transferring Winning Tickets for Object Recognition

In the remainder of this paper, we investigate the Ticket Transfer Hypothesis for object recognition problems. Using unstructured magnitude pruning, we find winning tickets that retain between 5-15% of the original parameters when transferring a sparse the representation found on CIFAR-10 to SmallNORB and FashionMNIST. In this section, we find winning tickets when the convolution layers are frozen while fine-tuning on the target task. Below, we briefly outline our algorithm for transferring winning tickets:

1. Randomly initialize $f(\mathbf{x}; \theta)$ and train the network on the source task \mathcal{T}_S . Let θ_S denote the parameters at the iteration of minimum validation loss.
2. Over k iterations, iteratively prune the network until $p\%$ of the parameters (the desired sparsity level) remain. Let $m \in \{0, 1\}^{|\theta|}$ denote the parameter mask.
3. Reset the pruned network to $f(\mathbf{x}; m \odot \theta_S)$, and train f on the target task \mathcal{T}_T

For our vision experiments, we investigate two widely-used convolutional networks; ResNet18 and VGG19. We also compare the performance of winning ticket and random initializations to two simple fully-connected architectures. For convolutional networks, we use the same configurations for our experiments as [Frankle and Carbin \(2019\)](#). These configurations are summarized in Table 1.

For convolutional networks, we eliminate 20% of parameters from the convolution layers at each iteration using unstructured magnitude pruning. During pruning, we leave the fully-connected layers untouched. This is critical in the context of transfer learning, since in some settings we only fine-tune the fully-connected layer while leaving the pruned convolution layers frozen.

In our various experiments, we use CIFAR-10 ([Krizhevsky, 2009](#)) as our source dataset, and fine-tune our models for two target datasets: FashionMNIST ([Xiao et al., 2017](#)) and SmallNORB ([LeCun et al.](#)), which we will abbreviate as FMNIST and NORB respectively. The source and target task is multi-class classification. The various attributes of each dataset are summarized in Table 2. We apply standard data-augmentation when training and fine-tuning our models: specifically, we augment inputs at training time using random horizontal flips and random 4-pixel pads and crops.

3.1 Winning Tickets with Frozen Convolutions

In this section, we consider fine-tuning winning tickets found for CIFAR-10 to both SmallNORB and FMNIST. We consider the setting where we “freeze” the convolution layers while fine-tuning on the target task, and only train the fully-connected layers.

We also consider two variants of this setting; first, we only fine-tune the single fully-connected layer in ResNet18 after training on the source task. Next, we replace the final fully-connected layer with a 2-layer feed-forward network, which is common practice for a wide variety of transfer learning tasks ([Pan and Yang, 2010](#)). In our experiments, we train sparse sub-networks from 3 different initializations; (1) the *late-reset* initialization, which in the transfer learning context corresponds to the fitted weights θ_S at the end of training on the source task, (2) the *winning ticket* initialization θ_0 , which is the original initialization for training on the source task, and (3) *randomly re-initializing* the weights to some $\theta' \sim \mathcal{D}_\theta$ before fine-tuning.

First, we show results for transferring winning tickets found on CIFAR-10 to SmallNORB. We initially consider the setting in which we only fine-tune the existing dense layer in ResNet18. Intuitively, this treats the pruned convolution layers as a sparse feature representation, and fine-tunes the classifier to predict the labels based on this representation.

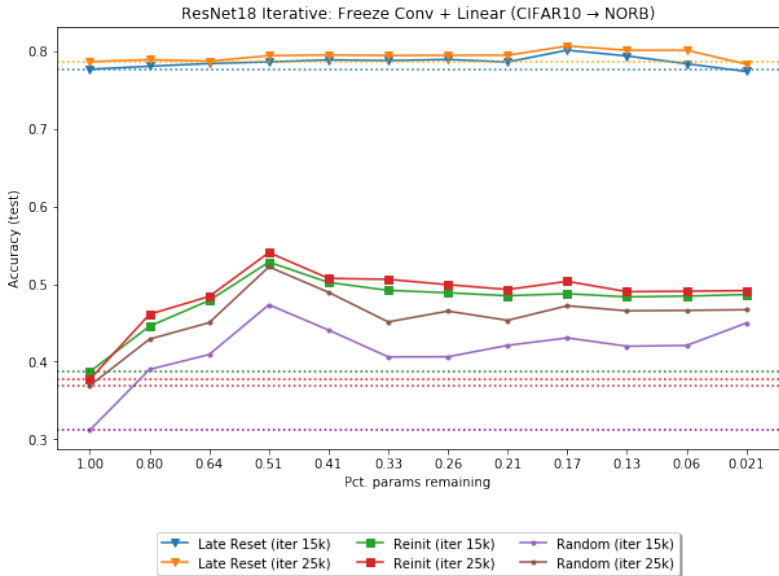


Figure 1: Winning Tickets: CIFAR-10 to NORB (ResNet18, freeze conv. + linear)

Figure 1 summarizes our results. In particular, we note that when fine-tuning from the late-reset initialization, we find winning tickets with as few as 6% of original parameters remaining. The top test accuracy of 80.01% is achieved when fine-tuning a winning ticket with 17% of original parameters retained.

Both the winning ticket and random reinitialization fails to reach commensurate accuracy at all pruning levels. In both cases, though, test accuracy improves as we sparsify the convolution layers. Both the random and winning ticket initializations reach their respective maximum test accuracies with 51% of parameters remaining. Accuracy remains relatively flat at that point, despite pruning more weights in the convolution layers.

Next, we repeat the experiment for ResNet18 after replacing the single fully-connected layer with a 2-layer network before fine-tuning. This increases the pruned parameter count by approximately 50,000 parameters, which is negligible when compared to the size of the convolution layers. Figure 2 summarizes our results.

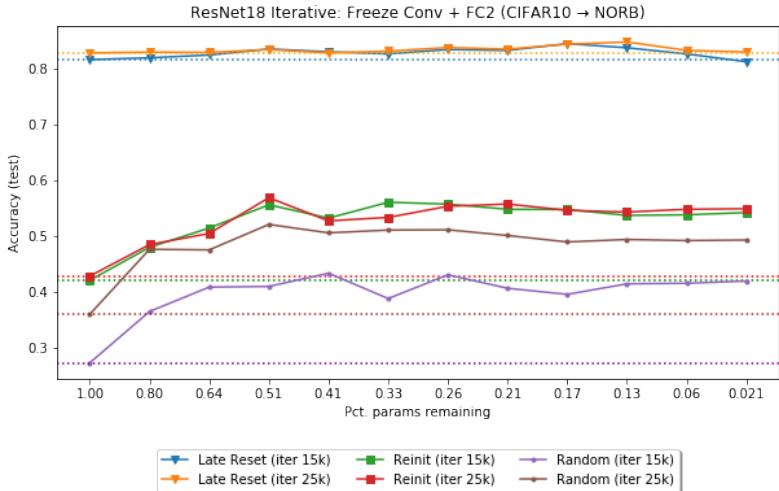


Figure 2: Winning Tickets: CIFAR-10 to NORB (ResNet18, freeze conv. + 2-layer fully-connected)

We demonstrate the same effect as in the linear output layer case; when fine-tuned on the target task when initialized with the fitted weights at the end of the source training task θ_S , we find winning tickets with as few as 6% of parameters retained. Note that while this is at the same sparsity level for the original winning ticket as before, we reach a superior test accuracy of 84.41%. Since we do not prune the dense layers, the acceleration effect described by Arora et al. (2018) likely accounts for the improvement in accuracy due to the overparameterization in the fully-connected layers. This is further supported by empirical evidence presented by Mehta (2018).

We repeat our previous experiment, this time with FashionMNIST as the target dataset. Again, we find winning tickets during training on the source task, and fine-tune the fully-connected layers with the pruned convolution layers frozen. As before, we replace the single output layer in ResNet18 with a 2-layer fully-connected network. Figure 3 summarizes our results.

Again, we observe the same phenomenon as before; the sub-networks that are reset to $m \odot \theta_S$ outperform all other initializations. We find winning tickets with as few as 7% of original parameters remaining. In addition, the previous effect we noticed with respect to reinitialization holds; all sub-networks reset to $m \odot \theta_0$ or $m \odot \theta'$ fail to reach commensurate accuracy. However, as the sparsity in the convolution layers increases, test accuracy improves.

3.2 Reinitialization dynamics with frozen convolutions

A puzzling feature of both experiments above is that although sparse sub-networks never reach commensurate accuracy when trained from the winning ticket initialization $m \odot \theta_0$, the accuracy improves and reaches a plateau as sparsity increases when trained from both the winning ticket as well as random reinitialization.

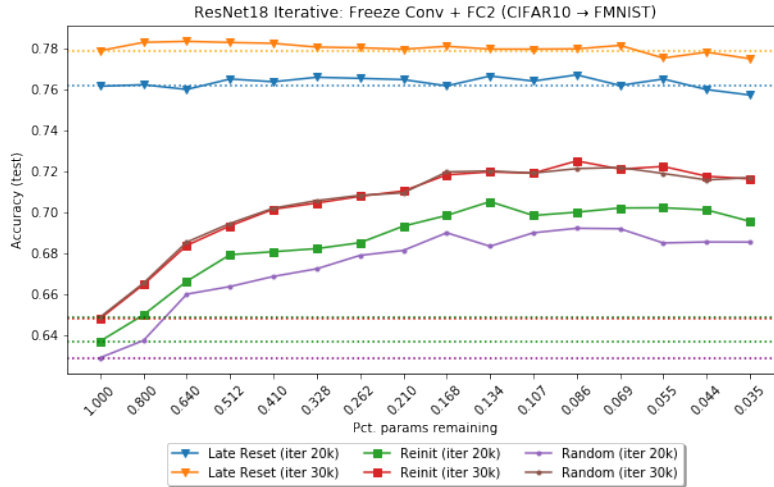


Figure 3: Winning Tickets: CIFAR-10 to FMNIST (ResNet18, freeze conv. + 2-layer fully-connected)

Since the convolution layers are frozen during fine-tuning on the target task, re-initializing these layers either to $m \odot \theta_0$, or $m \odot \theta'$ for any $\theta' \in \mathcal{D}_\theta$ should adversely affect accuracy. This is true, but accuracy also increases as we prune the convolution filters, without any noticeable drop-off.

We hypothesize that as we prune the layers following a reinitialization, we simply approach the performance of the fully-connected layers trained in isolation on the dataset. We experimentally test this by measuring the performance of winning ticket and random initializations relative to two networks; a dense, 2-layer fully-connected network, and a (single-layer) logistic regression classifier.

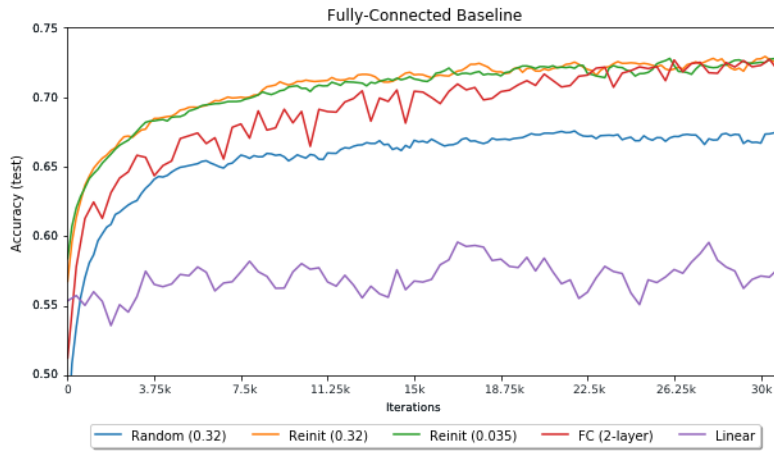


Figure 4: Comparison of CIFAR-10 transferred to FMNIST with 1 and 2-layer networks trained from scratch.

Figure 4 shows our results for this experiment. In particular, we note that the 2-layer network approaches the performance of ResNet18 when the convolution layers are frozen and reset during fine-tuning, while the linear regression baseline hovers between 58-65% accuracy. This confirms our hypothesis; as we increase the sparsity of the model, retraining our pruned network from its original initialization matches the performance of a 2-layer fully-connected network trained from scratch. This suggests that when reinitialized, any winning tickets found in the original deep network will effectively match the performance of training the classification layers in isolation. We note this also holds when we randomly reinitialize our winning ticket.

In the next section, we will show that if we eliminate the restriction that we can only fine-tune the convolution layers, we can recover performant networks trained from the winning ticket initialization.

4 Fine-Tuning the Entire Network

In this section, we relax the previous constraint and also train the convolution layers when fine-tuning on the target task. In particular, we examine the performance of iterative vs. one-shot pruning for ResNet18, and find winning tickets with as few as 7% of parameters remaining for both NORB and FMNIST. We also show that iterative pruning *fails* to produce winning tickets at *any* pruning level for VGG19 when transferring CIFAR-10 to both NORB and FMNIST; for brevity’s sake, we defer the full results to the appendix.

4.1 Transferring Winning Tickets for ResNet18

We begin with our results for ResNet18. Figure 5 shows our results for transferring winning tickets found for CIFAR-10 to FMNIST.

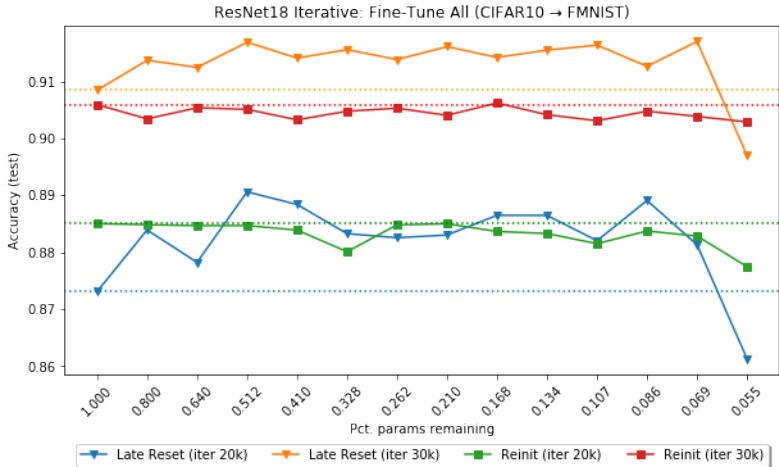


Figure 5: Winning Tickets: CIFAR-10 to FMNIST (ResNet18, fine-tune all, iterative pruning)

When we are able to fine-tune the convolution layers alongside the fully-connected layers, we find winning tickets with 6.9% of parameters remaining for ResNet18 when our network is initialized to $m \odot \theta_S$. After 30k and 50k training steps respectively, the late-reset sub-network exceeds original network’s accuracy at all pruning levels until 6.9%. While the winning tickets we find are the same size as those when we freeze the convolution layers during fine-tuning, we see that fine-tuning the convolutional layers improves the test accuracy of our best winning ticket to 91.7%.

Moreover, we see again that re-initializing the network to its original weights $m \odot \theta_0$ adversely affects the network’s performance on down-stream tasks. After 50k iterations, we find that winning tickets reset to $m \odot \theta_0$ fail to reach commensurate test accuracy at any sparsity level. However, the performance of the re-initialized winning ticket massively improves when we are allowed to fine-tune the convolutional layers.

4.2 Iterative vs. One-Shot Pruning

Frankle and Carbin (2019) show that iterative pruning consistently finds the smallest and most accurate winning tickets for a wide variety of simple models. Accordingly, we investigate the performance of one-shot vs. iterative pruning for finding winning tickets when transferring CIFAR-10 to SmallNORB.

Figure 6 shows our results for transferring winning tickets found for CIFAR-10 with one-shot pruning to NORB. Sub-networks initialized at $m \odot \theta_S$ match commensurate test accuracy to an extent; the




Figure 6: Winning Tickets: CIFAR-10 to NORB (ResNet18, fine-tune all, one-shot pruning)

smallest winning ticket we find is with 21.0% of parameters remaining, and the most accurate is found with 32.8% remaining (with a test accuracy of 89.3%). Moreover, when the network is reset to $m \odot \theta_0$, we fail to find winning tickets at all sparsity levels.

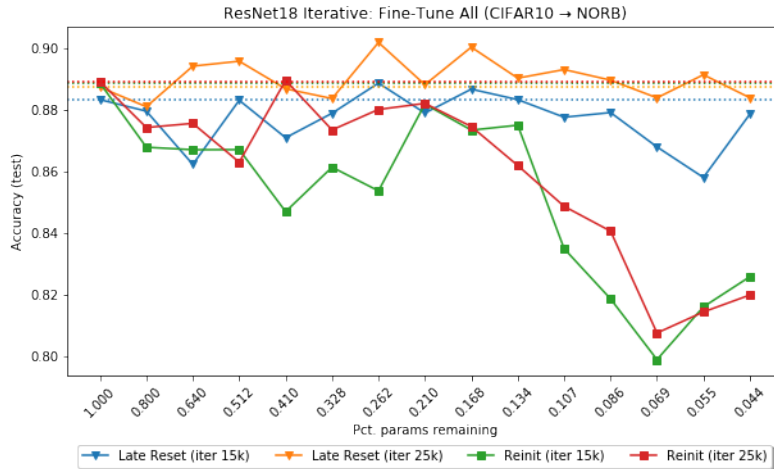


Figure 7: Winning Tickets: CIFAR-10 to NORB (ResNet18, fine-tune all, iterative pruning)

In comparison, Figure 7 shows the results for ResNet18, when winning tickets are found with iterative pruning. The sub-networks reset to $m \odot \theta_S$ again outperform all other initializations, and exceed the commensurate accuracy at nearly all pruning levels. The smallest winning ticket we find has 5.5% of parameters remaining, while the winning ticket that reaches the top validation accuracy has 26.2% of parameters remaining, and reaches a test accuracy of 90.2%. As before, all networks reset to $m \odot \theta_0$ fail to reach commensurate accuracy. Therefore, iterative pruning boosts both sparsity as well as accuracy.

5 Conclusion

In summary, we extend the Lottery Ticket Hypothesis to a number of transfer learning tasks in vision. We pose the *Ticket Transfer Hypothesis*, and experimentally validate it by showing that sparse sub-networks found with unstructured pruning can be transferred to similar tasks in various settings. Specifically, we find winning tickets both (1) when we fine-tune the fully-connected layers while

freezing the convolution layers, and (2) when we fine-tune the entire network end-to-end. Moreover, we confirm the efficacy of iterative pruning, and show that winning tickets from VGG19 do not transfer well to down-stream tasks.

Acknowledgements

I am grateful to Sanjeev Arora for his advice and guidance over the course of this work. I would also like to thank Jonathan Frankle, Pravesh Kothari, Karthik Narasimhan, Anthony Bak, and Matt Elkherj for many fruitful conversations throughout the preparation of this manuscript. This work was made possible by a generous grant from Amazon Research for our compute infrastructure.

References

- Sanjeev Arora, Nadav Cohen, and Elad Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 244–253, 2018. URL <http://proceedings.mlr.press/v80/arora18a.html>.
- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rJl-b3RcF7>.
- Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M. Roy, and Michael Carbin. The lottery ticket hypothesis at scale. *CoRR*, abs/1903.01611, 2019. URL <http://arxiv.org/abs/1903.01611>.
- Trevor Gale, Erich Elsen, and Sara Hooker. The state of sparsity in deep neural networks. *arXiv preprint arXiv:1902.09574*, 2019.
- 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, 2017.
- Song Han, Jeff Pool, John Tran, and William J. Dally. Learning both weights and connections for efficient neural networks. *CoRR*, abs/1506.02626, 2015. URL <http://arxiv.org/abs/1506.02626>.
- Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical Learning with Sparsity: The Lasso and Generalizations (Chapman & Hall/CRC Monographs on Statistics and Applied Probability)*. Chapman and Hall/CRC, 2015. ISBN 1498712169.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. URL <http://arxiv.org/abs/1512.03385>.
- Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- 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. *CoRR*, abs/1704.04861, 2017. URL <http://arxiv.org/abs/1704.04861>.
- Forrest N. Iandola, Song Han, Matthew W. Moskewicz, Khalid Ashraf, William J. Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <0.5mb model size. *arXiv:1602.07360*, 2016.
- Diederik P. Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick, 2015.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.

- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012. URL <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>.
- Jan Kukačka, Vladimir Golkov, and Daniel Cremers. Regularization for deep learning: A taxonomy, 2018. URL <https://openreview.net/forum?id=SkHkeixAW>.
- Yann Le Cun, John S. Denker, and Sara A. Solla. Optimal brain damage. In *Advances in Neural Information Processing Systems*, pages 598–605. Morgan Kaufmann, 1990.
- Yann LeCun, Fu Jie Huang, Leon Bottou, et al. Learning methods for generic object recognition with invariance to pose and lighting. Citeseer.
- Namhoon Lee, Thalaisyasingam Ajanthan, and Philip Torr. SNIP: single-shot network pruning based on connection sensitivity. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=B1VZqjAcYX>.
- Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *CoRR*, abs/1608.08710, 2016. URL <http://arxiv.org/abs/1608.08710>.
- Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rJlnB3C5Ym>.
- Christos Louizos, Max Welling, and Diederik P. Kingma. Learning sparse neural networks through l_0 regularization, 2017.
- Rahul Mehta. Implicit acceleration by overparameterization: An empirical perspective. Technical report, Princeton University, 2018.
- Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. Variational dropout sparsifies deep neural networks, 2017.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2010.
- Mark B. Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014. URL <http://jmlr.org/papers/v15/srivastava14a.html>.
- Andreas Veit, Michael J Wilber, and Serge Belongie. Residual networks behave like ensembles of relatively shallow networks. In *Advances in Neural Information Processing Systems*, pages 550–558, 2016.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *CoRR*, abs/1708.07747, 2017. URL <http://arxiv.org/abs/1708.07747>.
- Hattie Zhou, Janice Lan, Rosanne Liu, and Jason Yosinski. Deconstructing lottery tickets: Zeros, signs, and the supermask. 2019.




Figure 8: Winning Tickets: CIFAR-10 to NORB (VGG19, fine-tune all, iterative pruning)

A Transferring Winning Tickets for VGG19

We repeat a subset of our experiments for VGG19. We conduct the lottery ticket experiment for both NORB and FMNIST, and examine both the late-reset ($m \odot \theta_S$) and winning ticket ($m \odot \theta_0$) initializations.

Figure 8 shows the results of the lottery ticket experiment for VGG19 when transferring CIFAR-10 to NORB. At all sparsity levels, the sparse sub-networks fail to reach commensurate test accuracy. Networks reset to the winning ticket initialization still perform worse than the late-reset initialization, but accuracy notably degrades as sparsity increases.

We observe similar results when we repeat this experiment on FMNIST in Figure 9. We similarly observe that we fail to find winning tickets at all sparsity levels; for both the late-reset and winning ticket initializations, accuracy decreases as we increase the sparsity.

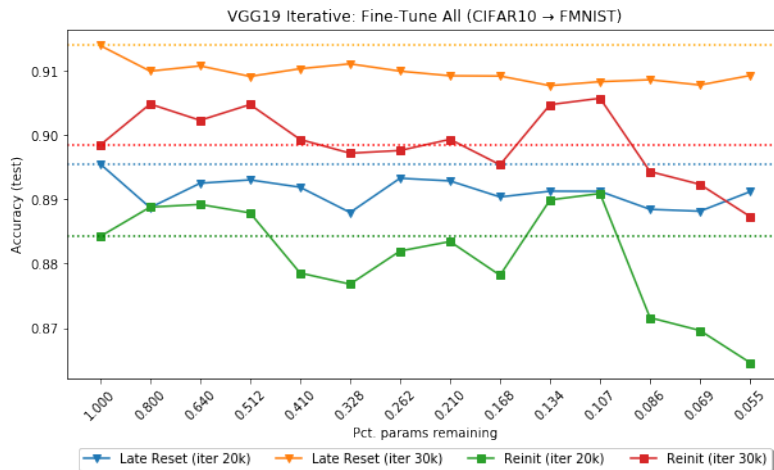


Figure 9: Winning Tickets: CIFAR-10 to FMNIST (VGG19, fine-tune all, iterative pruning)

This corroborates the findings of [Frankle and Carbin \(2019\)](#) and [Frankle et al. \(2019\)](#); they show that variants of ResNet consistently produce smaller and more accurate winning tickets than VGG, and only focus on ResNet-type architectures when scaling winning tickets up to ImageNet. These results provide another piece of evidence that ResNet produces superior winning tickets. In Section C, we draw on the connection between ResNets and ensemble learning identified by [Veit et al. \(2016\)](#) to provide some intuition as to why this might be the case.

B Related Work

We review several lines of research related to our problem. We first summarize various approaches for sparsifying neural networks, including regularization and pruning. We formally state the Lottery Ticket Hypothesis of [Frankle and Carbin \(2019\)](#), and finally provide a high-level overview of transfer learning.

B.1 Network Sparsity

Since initial neural network architectures were proposed, the closely-related problem of reducing the parameter count while maintaining accuracy has been extensively investigated. Benefits include lowering the memory footprint of the model, speeding up prediction, and potentially-improved generalization performance. Sparsification in this regime is also of independent interest, as it can shed light on the underlying dynamics of more complex, overparameterized models. In recent years, three dominant approaches have emerged:

Regularization The concept of regularization has a long history which extends far outside the domain of training deep neural networks. Various forms of regularization were simultaneously invented for solving ill-posed optimization problems and preventing overfitting. One of the earliest forms is known as Tikhonov regularization, which imposes an ℓ_2 weight penalty on model parameters. In classical statistics, this is called *ridge regression*, and is connected to *weight-decay* in the machine learning community. Other forms of “regularized regression” include an ℓ_1 penalty, commonly referred to as LASSO. For an in-depth summary of regularization in classical statistics, see [Hastie et al. \(2015\)](#).

This concept translates naturally to deep neural networks, with a small modification; instead of penalizing the ℓ_2 norm of all weights in the loss function, *decay* the gradient update at each iteration proportional to the ℓ_2 norm of the particular parameter. [Krizhevsky et al. \(2012\)](#) cite a small weight-decay parameter as a crucial component of AlexNet’s state-of-the-art performance at ILSVRC ‘12. The same procedure can be performed with the ℓ_1 norm instead, which imposes sparsity on the fitted weights ([Kukačka et al., 2018](#)).

More recent types of regularization directly seek to impose sparsity during training. The most well-known family of these methods is called *dropout regularization*, which was introduced by [Srivastava et al. \(2014\)](#). This method stochastically “drops out” certain inputs in each layer throughout training, which prevents overfitting and improves test accuracy. Recent extensions of this technique include variational dropout ([Kingma et al., 2015](#); [Molchanov et al., 2017](#)). [Louizos et al. \(2017\)](#) propose ℓ_0 regularization, a related approach which derives a novel technique to circumvent the non-differentiability of the ℓ_0 norm and imposes a strict sparsity constraint during training.

Task-Specific Architectures Another approach to reducing the parameter count of models and improving inference speed is developing streamlined, task-specific architectures. Rather than seeking to sparsify an existing deep network, these models are engineered specifically with parameter savings and inference speed in mind.

Two examples in this category that still remain state-of-the-art for many use-cases include SqueezeNet ([Iandola et al., 2016](#)) and MobileNetV1/V2 ([Howard et al., 2017](#); [Sandler et al., 2018](#)). SqueezeNet incorporates techniques such as replacing 3×3 filters with 1×1 filters and delaying pooling until the final stages of the network to increase activation magnitude in the convolution layers. MobileNetV1 leverages the insight that a convolution can be expressed as a composition of *depth-wise separable* and *point-wise* convolutions, and introduces two hyper-parameters to control the *width* of the network and the *resolution* of the input images [Howard et al. \(2017\)](#). MobileNetV2 extends its predecessor by incorporating residual blocks (such as in [He et al. \(2015\)](#)) along with depth-wise separable convolutions, and achieves state-of-the-art performance for a number of inference tasks while approaching the accuracy of very deep, dense networks such as ResNet and VGG [Sandler et al. \(2018\)](#).

Network Pruning A third approach to network sparsity (and the one that we will predominantly focus on in this thesis) is known as *network pruning*. In this setting, the full model is first trained

end-to-end, and subsequently parameters are removed via a heuristic. The model is then re-trained after pruning to recover from the loss in accuracy.

Pruning and related concepts were introduced as early as 1990 by [Le Cun et al. \(1990\)](#) through the concept of optimal brain damage. [Han et al. \(2015\)](#) and [Li et al. \(2016\)](#) provide two modern approaches to network pruning. [Han et al. \(2015\)](#) introduce a methodology to prune networks with one round of fine-tuning (i.e. train, prune, then fine-tune) that achieved significant parameter savings on VGG (about a 50% parameter reduction). They also extend this to *iterative* pruning, which performs multiple rounds of fine-tuning. Broadly, these methods are referred to as *unstructured* pruning, since all layers are pruned simultaneously by the heuristic.

Similarly, [Li et al. \(2016\)](#) introduce *structured* pruning, in which entire layers/structured subsets of layers are pruned prior to fine-tuning; they describe a method for pruning CNNs in which whole filters are pruned. The filter-based pruning reduces the parameter count by about 38% for ResNet-110 without an appreciable drop in accuracy. It is important to note here that while unstructured, magnitude-based pruning can find much smaller networks that reach the same accuracy, structured pruning yields computational speedups as well by making the network shallower.

[Frankle and Carbin \(2019\)](#) use the two unstructured pruning methods introduced by [Han et al. \(2015\)](#). In particular, they use both *iterative* and *one-shot* pruning to find sparse, trainable sub-networks. All results are found by using simple magnitude-based pruning, in which the smallest $p\%$ of weights by magnitude are pruned at each iteration. For a summary of current state-of-the-art approaches to pruning, see [Gale et al. \(2019\)](#). Pruning has also been investigated in the context of transfer learning by [Molchanov et al. \(2017\)](#), who describe a method for pruning entire filters based on the first-order gradient information of the layers to improve the generalization performance of ResNet-50 for transfer learning tasks.

Pruning as Network Architecture Search Another interpretation of network pruning frames it as a form of *network architecture search*. Conventional wisdom holds that widely-used network architectures are massively over-parameterized; this overparameterization aids in generalization performance, and has been shown to also positively affect optimization in certain circumstances ([Arora et al., 2018](#)). Accordingly, several authors take the view that pruning uncovers hidden networks within a massively over-parameterized deep network. [Gordon et al. \(2017\)](#) propose MorphNets, which learn a sparse architecture throughout the training process via a series of interleaved *prune* and *expand* phases. Similarly, [He et al. \(2018\)](#) develop an automatic approach to finding sparse architectures which they call “AutoML for Model Compression” (AMC). They frame the problem of coarse, structured pruning as a reinforcement learning problem and describe an automatic pipeline for pruning channels and fine-tuning.

B.2 The Lottery Ticket Hypothesis

[Frankle and Carbin \(2019\)](#) introduced the Lottery Ticket Hypothesis, and provided a wide set of experimental evidence including experiments on basic fully-connected networks as well as convolutional networks trained on CIFAR-10. They also posed the question of whether the sparse representation produced by pruning can be transferred between tasks. Subsequently, [Frankle et al. \(2019\)](#) showed that a slight modification of the above method finds winning lottery tickets for ResNet50 trained on ImageNet. The authors achieve test accuracy within a percentage point of the un-pruned model with as many as 79% of original parameters removed. In particular, they introduce *late resetting*, in which the sparse sub-network is re-initialized not to $m \odot \theta_0$, but rather $m \odot \theta_t$, for some $t \ll j$ (recall that j is the iteration at which the original network $f(\cdot)$ reaches its minimum validation accuracy). By resetting the network to a checkpoint several epochs *after* initialization, they find strong evidence of winning tickets for large-scale visual recognition problems. In a recent follow-up work, [Zhou et al. \(2019\)](#) investigate some of the underlying dynamics of winning tickets, and find that masking can be viewed as a form of training, insofar that it biases the retained weights in a particular direction. [Lee et al. \(2019\)](#) introduce a distinct but related approach called SNIP, which performs one-shot pruning on the network based on the sensitivity of individual connections. However, while exceeding the performance of randomly-reinitialized lottery tickets, the “winning ticket initialization” discovered by [Frankle and Carbin \(2019\)](#) outperforms SNIP by an appreciable margin.⁴

⁴See <https://openreview.net/forum?id=rJl-b3RcF7> for details.

B.3 Evidence against the Lottery Ticket Hypothesis

While there is a sufficient body of evidence in favor of the Lottery Ticket Hypothesis, other works have focused on a holistic comparison of various pruning heuristics and methods. In particular, [Gale et al. \(2019\)](#) conduct an extensive investigation into various methods for inducing sparsity, including variational dropout due to [Molchanov et al. \(2017\)](#), ℓ_0 regularization due to [Louizos et al. \(2017\)](#), and unstructured magnitude pruning as originally described by [Han et al. \(2015\)](#). While they demonstrate that naive magnitude pruning outperforms more complex methods, they also find that naive magnitude pruning (without late-resetting) fails to find winning tickets for both ResNet50 on ImageNet, as well as for Transformer-based models on the WMT 2014 English-to-German translation task.

In addition, [Liu et al. \(2019\)](#) consider network pruning from the perspective of network architecture search and show that, among other results, a careful choice of learning rate can alleviate the negative effects of random re-initialization. In particular, they show that with a carefully-chosen learning rate, the “winning ticket initialization” provides no appreciable benefit over random initialization for the pruned sub-network.

C Limitations & Future Work

Our results in Sections 3.1 and 4 rely on small-scale datasets such as CIFAR-10, SmallNORB, and FashionMNIST. In a similar manner to [Frankle et al. \(2019\)](#), we would like to scale up our existing experiments to deeper models such as ResNet-50/ResNet-152 trained on ImageNet. We also believe that larger-scale tasks will show a far greater accuracy drop-off when using the winning ticket or random re-initialization, since fully-connected networks cannot begin to approach the accuracy of convolutional networks on larger-scale tasks. In addition, we also believe that the Ticket Transfer Hypothesis offers a number of interesting directions for future research in network sparsity.

Transferring winning tickets between tasks In all our experiments, we considered a transfer learning setting in which we fine-tune the pre-trained network on a different dataset but on the same task (object recognition). It would be interesting to verify if winning tickets found for a source task such as object detection could transfer to distinct down-stream tasks such as semantic segmentation or image captioning.

Winning tickets via structured pruning [Frankle and Carbin \(2019\)](#) posed the original Lottery Ticket Hypothesis in the context of unstructured magnitude pruning. While we are able to obtain massive parameter savings, pruned models cannot take advantage of this increased sparsity to improve prediction speed, unless specialized hardware is used ([Han et al., 2015](#)). Another approach to pruning is given by [Li et al. \(2016\)](#), which proposed *structured layer-wise* pruning which eliminates whole convolutional filters. It is interesting to ask if pruning ResNet and VGG in this manner will produce winning tickets that are more computationally efficient for transfer learning.

Connection to ensemble learning In addition to scaling up our experiments to larger tasks in different contexts, we also propose an avenue towards determining the underlying dynamics of the lottery ticket hypothesis. In particular, both [Frankle and Carbin \(2019\)](#) and [Frankle et al. \(2019\)](#) determine that winning tickets found for ResNet-type architectures outperform VGG for vision tasks. Moreover, our results in Section 4 corroborate this finding, as winning tickets found on VGG fail to reach commensurate accuracy at all sparsity levels. We hypothesize that the improved performance of ResNet in these tasks can be accounted for by its implicit training dynamics as an ensemble; in particular, [Veit et al. \(2016\)](#) show that ResNet behaves like an ensemble of shallow residual networks during training. This view of ResNet as an *overparameterized ensemble* (with weight-sharing) leads to a natural interpretation of our unstructured pruning; we prune all the weak learners in the overparameterized ensemble simultaneously. This also might suggest why winning tickets found on ResNet transfer well to other tasks. To this end, we ask if a lottery ticket-style experiment can be devised on a traditional ensemble model (such as a random forest) to investigate this connection.

D Network Configurations

<i>Network</i>	Fully-Connected		ResNet18	VGG19
<i>Convolutions</i>			16 3x[16,16] 3x[32,32] 3x[64,64]	2x64, pool 2x128 pool, 4x256, pool 4x512, pool, 4x512
<i>FC Layers</i>	1024, 10	1024, 100, 10	avg-pool, 10	avg-pool, 10
<i>All/Conv Weights</i>	10,240/0	1.02M/0	11.6M/10.9M	20.03M/20.02M
<i>Iterations/Batch</i>	30k/128	30k/128	30k/128 CIFAR-10 25k/128 NORB 30k/128 FMNIST	112k/64 CIFAR-10 50k/64 NORB 112k/64 FMNIST
<i>Optimizer</i>	SGD 0.01,0.005, Momentum 0.9		SGD 5e-3,1e-3,1e-4, Momentum 0.9, Weight Decay 1e-4	SGD 1e-2-1e-3,1e-4, Momentum 0.9, Weight Decay 1e-4
<i>Pruning Rate</i>	N/A		Conv: 20%, FC: 0%	

Table 1: Model configurations for vision experiments

<i>Domain</i>	Source	Target	
<i>Dataset</i>	CIFAR-10	SmallNORB	FMNIST
<i>Train/Test</i>	50k/10k	40k/10k	60k/10k
<i>Num. class</i>	10	5	10
<i>Shape</i>	32x32x3	28x28x1	28x28x1
<i>Task</i>	Multi-Class Classification		

Table 2: Summary of source and target datasets and tasks for vision experiments.