

---

# Deconstructing the Structure of Sparse Neural Networks

---

Maxwell Van Gelder\*  
University of Washington

Mitchell Wortsman  
University of Washington

Kiana Ehsani†  
University of Washington

## Abstract

Although sparse neural networks have been studied extensively, the focus has been primarily on accuracy. In this work, we focus instead on *network structure*, and analyze three popular algorithms. We first measure performance when structure persists and weights are reset to a different random initialization, thereby extending Zhou *et al.* [20]. This experiment reveals that accuracy can be derived from structure alone. Second, to measure structural robustness we investigate the *sensitivity* of sparse neural networks to further pruning after training, finding a stark contrast between algorithms. Finally, for a recent dynamic sparsity algorithm we investigate how early in training the structure emerges. We find that even after one epoch the structure is mostly determined, allowing us to propose a more efficient algorithm which does not require dense gradients throughout training. In looking back at algorithms for sparse neural networks and analyzing their performance from a different lens, we uncover several interesting properties and promising directions for future research.

## 1 Introduction

Sparse networks achieve impressive accuracy while using only a fraction of the parameters [13, 9, 7, 5, 4]. Though earlier work on sparse networks focused primarily on pruning *after* training, researchers have recently shown interest in pruning early in training [5, 14, 17, 18] or dynamically as training progresses [2, 19, 3].

Sparse neural networks continue to intrigue researchers for two main reasons. First, pruning neural networks can vastly reduce storage and computational costs, important steps towards reducing energy consumption and GreenAI [16]. Additionally, sparsifying neural networks provides a useful tool for understanding learning dynamics and other phenomena. For instance, the *lottery ticket hypothesis* [5] conjectures that overparameterized neural networks contain subnetworks which can be trained to competitive accuracy when their weights are reset to their initialization. These findings may suggest that SGD seeks out and trains a lucky subnetwork early in training. Moreover, in [20, 15] high accuracy is achieved by sparsifying a randomly initialized neural network with fixed weights.

In this paper we investigate sparse neural networks *structures*. We are inspired by the analysis of [20], wherein Zhou *et al.* carefully analyze which aspects of the Iterative Magnitude Pruning (IMP) [5] contribute to the accuracy. We bring a similar tool and perspective to the analysis of network structure across different algorithms. In addition to IMP we analyze two dynamic sparsity methods: Discovering Neural Wirings (DNW) [19] and Rigged Lottery Tickets (RigL) [3], algorithms where the connectivity changes throughout training. We then ask the following three questions:

1. What role does *structure* play in a sparse neural network’s performance? As in [20], we investigate this by using the sparse neural network structure found at the end of training

---

\*Correspondence to maxdvg@uw.edu. †Also affiliated with the Allen Institute for AI.

and resetting the weights to a different random initialization before retraining. Our findings match [20] for IMP subnetworks, which perform similarly to random subnetworks when retrained with new randomly initialized parameters. Surprisingly, we find that this isn't true for DNW and RigL subnetworks, which outperform random subnetworks across various sparsity values.

2. How robust are different network structures? Specifically, we investigate how quickly accuracy degrades when further pruning networks after training. We conclude that IMP, DNW, and RigL subnetworks are less sensitive than a random subnetwork when further pruning by magnitude.
3. In DNW, a dynamic sparsity method, how early in training does the final structure emerge? We find that surprisingly the structure emerges fairly early in training and leverage these results to design a more efficient algorithm.

## 2 Preliminaries and Related Work

We investigate four classes of sparse graphs. For the first class of graphs, which we refer to as a random graph, we choose a random set of weights to remain zero throughout training. We also consider Lottery Ticket (LT) [5] graphs, DNW [19] graphs, and RigL graphs [3] which are respectively produced by the following three algorithms:

**1. Iterative Magnitude Pruning (IMP).** IMP [5, 6] iteratively sparsifies a neural network in successive training runs. First a dense network is trained until completion, at which point the bottom  $p\%$  of the weights by magnitude are set to 0. The nonzero weights are then reset to their initialization and the process continues until the requisite sparsity is attained. Although the process is expensive, IMP produces highly effective sparse neural networks.

**2. Discovering Neural Wirings (DNW).** As in [2, 3], DNW [19] maintains sparsity throughout training. In the forwards pass, DNW selects the top  $p\%$  of weights by layer to be nonzero. In the backwards pass, all weights are updated via the straight-through estimator allowing the re-entry of dead weights.

**3. Rigged Lottery (RigL).** RigL [3] maintains sparsity throughout training and does not require dense gradients during most iterations of training. RigL graph structure is static for the majority of training iterations, and only changes every  $n$  iterations up to some stopping iteration  $T_{end}$ . On those iterations where graph structure *is* changed, the bottom  $k$  parameters are removed from the network by absolute magnitude. To replace these  $k$  lost parameters, the  $k$  parameters with the greatest magnitude gradient are introduced into the network.  $k$  decreases as the number of iterations increases, until  $T_{end}$  is reached.

## 3 How Does Structure Affect Accuracy?

The four algorithms we investigate find sparse neural graph structures using different methods. These differences between algorithms might be expected to lead each algorithm to generate graph structures with unique characteristics. Understanding how much graph *structure* is responsible for the accuracy of each of these algorithms, as opposed to the importance of a specific initialization, is important for understanding the generalizability of the structures they create.

To determine how the four different classes of graphs perform under varying levels of sparsity, we compare their accuracy at each degree of sparsity when their parameters are reinitialized with multiple new random seeds. This experiment extends [20], where authors conduct this analysis for LT graphs.

Using ResNet-18 [10] on CIFAR-10 [12], we find that LT graphs perform similarly to random graphs when their weights are reset randomly. This corroborates the findings of [20]. Surprisingly, however, we find that DNW and RigL graphs outperform random graphs, suggesting that their performance may in part due to a characteristic of their structure which is absent in LT and random graphs. As illustrated in Figure 1, this performance gap is especially pronounced in the sparse regimes.

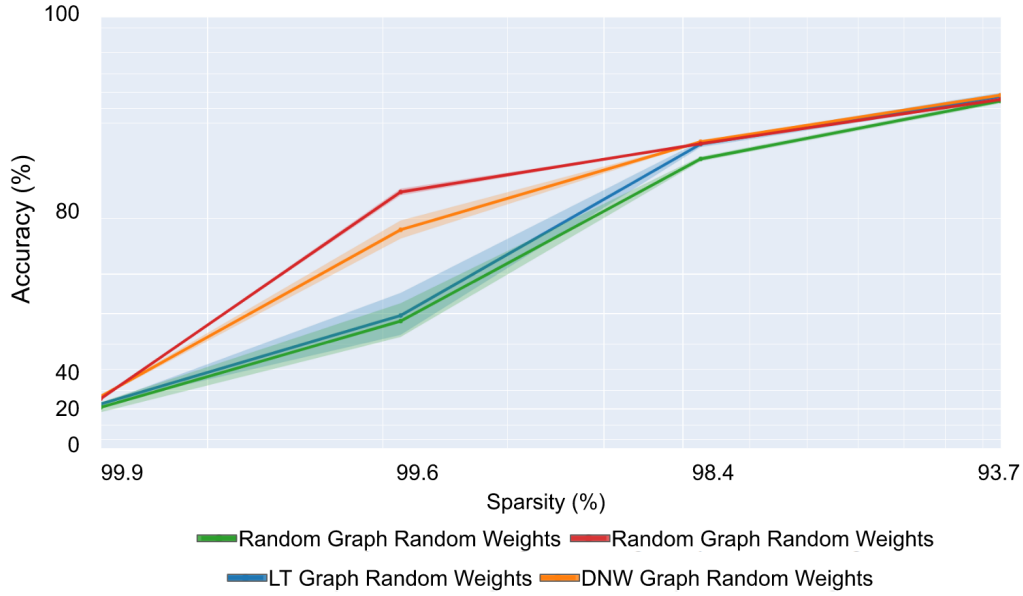


Figure 1: **Accuracy Vs. Sparsity.** Performance of ResNet-18 on CIFAR-10 when trained with the structure found by sparsity algorithms but with weights reset randomly. Error bars show standard deviation of accuracy when graph structures of each type are reinitialized with three random seeds.

#### 4 Sensitivity to Further Pruning

In linear optimization, *sensitivity* analysis can be used to identify the stability of a solution for maximizing or minimizing a linear program. When small perturbations to the inequalities defining a linear program leave the optimal solution unchanged, the optimal solution is said to be stable. In order to test the robustness of a graph’s structure, we experiment with an analogous metric which measures the stability of a sparse neural network.

To determine the robustness of each graph type, we introduce an algorithm for sensitivity analysis on neural graph structures. This sensitivity analysis is performed by further eliminating non-zero parameters *after training* and evaluating the performance of the reduced network. We select which weights to remove both randomly and by order of ascending magnitude, as illustrated in Figure 2. Our findings indicate that when removing weight by magnitude the performance of random graphs decreases faster than DNW, LT, and RigL graphs. This suggests that lower magnitude parameters in LT, DNW and RigL graphs are less critical. This is partially expected for DNW and RigL, which continuously swap out the lower magnitude weights. When removing weights randomly, DNW and LT slightly outperform RigL and random graphs, suggesting that there may be more redundancy within LT and DNW graphs.

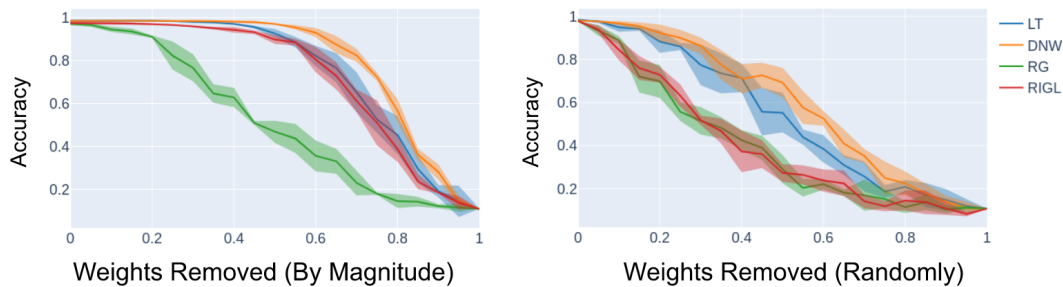


Figure 2: Sensitivity to parameter zeroing for a 90.5% sparse fully connected network on MNIST where weights are removed in order of ascending magnitude (**left**) and randomly (**right**). Error bars show standard deviation with three random seeds.

## 5 How Early in Dynamic Training Does Structure Emerge?

In order to improve dynamic training algorithms, it is helpful to have tools which monitor how they modify graph structure as they train. One simple metric which allows us to roughly track graph evolution during training is to track node in-degrees. By tracking the in-degree of each node in a layer we are able to gain insight into how networks evolve during training with dynamic algorithms. Surprisingly, when analysing DNW, we find that the majority of the structural changes happen in the first epoch, and that the remainder of the structural change is mainly an amplification of the structure created in the first epoch. Nodes which have higher in-degrees after the first epoch of training continue to increase the number of connections they have, while those nodes with lower in-degrees after the first epoch see declines in their importance over the rest of training.

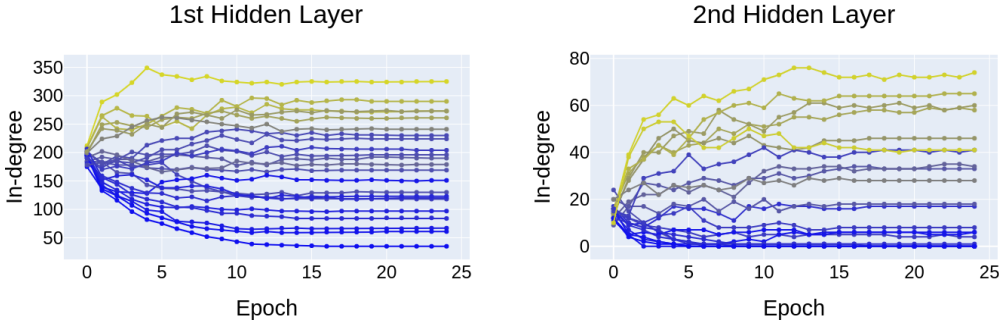


Figure 3: In-degree evolution throughout training with DNW. The in-degree of a random subsampling of neurons is shown for a small sparsified fully connected network on MNIST with 90.5% sparsity. Color is used to show in-degree after one epoch of training, where yellow is a higher in-degree after the first epoch and blue is a lower in-degree after the first epoch. These results showcase a strong correlation between in-degree after the first epoch and final in-degree.

For normal DNW training, the gradient is propagated to all parameters with the straight through estimator [1], allowing dead weights to resurface. In a very sparse graph, this means that a high percentage of the computational power is consumed looking for new edges to add to the graph. However, we have shown in Figure 3 that the final DNW structure is primarily an amplification of the structure after the first epoch, and therefore much of this searching is superfluous.

Based on these results, we propose a hybrid dynamic/static training method where the structure is allowed to change in the first  $k$  epochs using the normal DNW algorithm, and structure is fixed for the remainder of training. Accordingly we refer to the  $k$ th epoch by the *freezing* epoch. We illustrate the performance of our hybrid algorithm in Figure 4.




Figure 4: The effect of the *freezing* epoch on performance in our hybrid dynamic/static training method for ResNet-18 on CIFAR-10. Error bars give standard deviation from three random seeds at each freezing epoch.

We find that, after the first few epochs, further training with DNW yields diminishing returns. This corroborates our finding that structure is primarily determined within the first few epochs of DNW. However, if compute is unbounded it is still best to train with DNW to completion. In comparison to newer dynamic training algorithms including RigL and Top-KAST [3, 11], our modified DNW algorithm is likely still less computationally efficient for most applications. We believe the more important finding, however, is that simple introspection into graph evolution during dynamic training can make improvements to existing algorithms easier to identify.

## 6 Conclusion

We have taken a retrospective look at sparse neural networks through a different lens. Our findings offer interesting insight into a growing field which is becoming increasingly practical [8]. We have shown that in contrast with IMP, structures found by DNW and RigL perform well independently of initialization. This suggests that there are properties in dynamically generated structures that are absent in LT and random graphs which are conducive to sparse learning. Moreover we have shown that random graphs are much more sensitive to changes in connectivity than the other graph types. Finally, using a simple tool we find that the majority of structural changes in DNW occur in the first few epochs, leading us towards a more efficient algorithm. More generally, we believe that analyzing existing algorithms from a variety perspectives is often very useful and underexplored.

## References

- [1] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- [2] Tim Dettmers and Luke Zettlemoyer. Sparse networks from scratch: Faster training without losing performance. *arXiv preprint arXiv:1907.04840*, 2019.
- [3] Utku Evci, Trevor Gale, Jacob Menick, Pablo Samuel Castro, and Erich Elsen. Rigging the lottery: Making all tickets winners. *arXiv preprint arXiv:1911.11134*, 2019.
- [4] Utku Evci, Fabian Pedregosa, Aidan Gomez, and Erich Elsen. The difficulty of training sparse neural networks. *arXiv preprint arXiv:1906.10732*, 2019.
- [5] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018.
- [6] Jonathan Frankle, David J Schwab, and Ari S Morcos. Training batchnorm and only batchnorm: On the expressive power of random features in cnns. *arXiv preprint arXiv:2003.00152*, 2020.
- [7] Trevor Gale, Erich Elsen, and Sara Hooker. The state of sparsity in deep neural networks. *arXiv preprint arXiv:1902.09574*, 2019.
- [8] Trevor Gale, Matei Zaharia, Cliff Young, and Erich Elsen. Sparse gpu kernels for deep learning. *arXiv preprint arXiv:2006.10901*, 2020.
- [9] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [11] Siddhant Jayakumar, Razvan Pascanu, Jack Rae, Simon Osindero, and Erich Elsen. Top-kast: Top-k always sparse training. *Advances in Neural Information Processing Systems*, 33, 2020.
- [12] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [13] Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In *Advances in neural information processing systems*, pages 598–605, 1990.

- [14] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip Torr. Snip: Single-shot network pruning based on connection sensitivity. In *International Conference on Learning Representations*, 2018.
- [15] Vivek Ramanujan, Mitchell Wortsman, Aniruddha Kembhavi, Ali Farhadi, and Mohammad Rastegari. What’s hidden in a randomly weighted neural network? *arXiv preprint arXiv:1911.13299*, 2019.
- [16] Roy Schwartz, Jesse Dodge, Noah A Smith, and Oren Etzioni. Green ai. corr abs/1907.10597 (2019). *arXiv preprint arXiv:1907.10597*, 2019.
- [17] Hidenori Tanaka, Daniel Kunin, Daniel LK Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. *arXiv preprint arXiv:2006.05467*, 2020.
- [18] Chaoqi Wang, Guodong Zhang, and Roger Grosse. Picking winning tickets before training by preserving gradient flow. In *International Conference on Learning Representations*, 2019.
- [19] Mitchell Wortsman, Ali Farhadi, and Mohammad Rastegari. Discovering neural wirings. In *Advances in Neural Information Processing Systems*, pages 2684–2694, 2019.
- [20] Hattie Zhou, Janice Lan, Rosanne Liu, and Jason Yosinski. Deconstructing lottery tickets: Zeros, signs, and the supermask. In *Advances in Neural Information Processing Systems*, pages 3592–3602, 2019.