# Visualizing the Loss Landscape of Winning Lottery Tickets

Robert Bain<sup>1</sup>

## Abstract

The underlying loss landscapes of deep neural networks have a great impact on their training, but they have mainly been studied theoretically due to computational constraints. This work vastly reduces the time required to compute such loss landscapes, and uses them to study winning lottery tickets found via iterative magnitude pruning. We also share results that contradict previously claimed correlations between certain loss landscape projection methods and model trainability and generalization error.

# 1. Introduction

Deep neural networks (DNNs) are frequently trained using one of the many variants of stochastic gradient descent (SGD). These methods update a network's parameters using the gradients of the loss w.r.t. said parameters. DNNs have many degrees of freedom (i.e. weights), and their objective functions are thus very-high dimensional. For example, ResNet50 has over 23 million trainable parameters. The "loss landscape" is the same number of dimensions as the weight space plus 1, as each possible configuration of the DNN is evaluated for its loss over some number of test examples (i.e. examples not seen during training).

The first few sections of this paper cover the theoretical background surrounding the loss landscape and introduce the specific visualization method used in this work. The loss landscape (also referred to as "loss surface" and "objective landscape") is constructed by calculating the loss of multiple points in the weight space (i.e. different configurations) of a DNN. Later in the paper we introduce the lottery ticket hypothesis (LTH) and iterative magnitude pruning (IMP) by (Frankle & Carbin, 2019), and apply the same loss visualizations to winning lottery tickets (WLTs) created using IMP.

All figures in this paper contain a hyperlink in the caption to view the same data with *LossPlot*, an in-browser application built specifically to visualize these types of surface plots. The projected contours, radius based clippings, and other settings can be manually controlled from *LossPlot*. It is designed to be used with a mouse and keyboard.



*Figure 1.* The loss landscape of a ResNet56, with skip connections removed, trained on CIFAR10. Skip connections are a useful inductive bias that makes DNNs both easier to train and achieve lower generalization error. The loss surface is extremely chaotic and non-convex, which hints at some of the difficulty in training these types of deep networks. Note that the z-component is on a natural log scale. View and handle the data yourself here.

# 2. Contributions

- A method of computing the loss landscape that is 100x faster than previous methods.
- Demonstrate the effect of varying batch size on the loss surface of DNNs. Our results contradict previously noted correlations of certain visualization features with trainability and generalization error.
- Contrast the loss landscape of randomly pruned DNNs and winning lottery tickets produced by iterative magnitude pruning.

## 3. Loss Visualization Background

The ability to train DNNs can often be surprising given that they have highly non-convex structures in their objective landscapes. (Neyshabur et al., 2017) explore heuristics for characterizing a DNN's ability to generalize and their best

Loss Landscape and the Lottery Ticket Hypothesis



*Figure 2.* The loss landscape of a ResNet56 trained on CIFAR10. The only difference between each subplot is the number of test examples used to calculate each configuration's generalization loss. The same random directions were used across all subplots to ensure a fair comparison. You can handle the data yourself by clicking this link.

finding was a combination of sharpness and weight magnitude. Flatness is the size of the connected region with similar loss, and sharpness is the antithesis. Flatter minimizers will have larger connected regions of similar loss around the trained minimizer at the center of each surface plot (i.e. there will be little relative curvature around the trained solution). Sharp weight vectors on the other hand have deep crevices immediately surrounding the solution in the loss space.

Modern DNNs often use rectified activation functions (e.g. ReLU) which introduce redundancies in the weight space of models. Scaling one layer by a constant and then scaling the next layer by the inverse results in the same function. Combining batch norm with rectified activation functions has a similar effect, because everything is scaled to a common norm, so a relative comparison between different trained solutions becomes difficult (Dinh et al., 2017). Both of these introduce what are referred to as "scaling invariance".

All visualizations of the loss landscape are generated using the methods from (Li et al., 2017). Their open-sourced pytorch code evaluates a 2d grid slice of the weight space centered around the trained minimizer. There are many 2d grid slices to choose from given the dimensionality of modern DNNs. (Li et al., 2017) uses a dimensionality reduction technique to choose the slice of the weight space to plot its loss. It begins with creating two random weight vectors by sampling a Gaussian distribution  $\mathcal{N}(0, 1)$ .

To fix the problems caused by scaling invariance, (Li et al., 2017) use "filter-wise normalized directions". There Gaussian sampled direction vectors of length 1 are scaled along the components relating to convolutional neural network (CNN) filters and that subset of the weights norm. I.e. they scale the component subsets of the random directions using the corresponding magnitude of filters:

$$d_{i,j} \leftarrow \frac{d_{i,j}}{\|d_{i,j}\|} \|\theta_{i,j}\|$$

*d* represents 1 of the 2 random directions and  $\theta$  are the DNN's weights.  $d_{i,j}$  represents the j-th filter of the i-th layer of *d*.  $\|\cdot\|$  is the Frobenius norm (Li et al., 2017). Filters are not just limited to CNNs. A fully connected layer's filters are the weights connected to one neuron in the resultant layer.

Filter-wise normalization with random directions captures the weighted average of the principle curvatures (i.e. eigenvalues of the Hessian matrix) and enables visual comparison

Loss Landscape and the Lottery Ticket Hypothesis



*Figure 3.* The same model architecture as in Figure 2 but trained without skip connections. The number of examples required to get a general notion of the overall shape is similar, but the landscape has become far more chaotic & non-convex.

of loss surfaces around different trained solutions, even among different architectures (Li et al., 2017). (Li et al., 2017) claim that it also causes the flatness of minimizers to negatively correlate with generalization error. This is contested later in this paper.

(Li et al., 2017) reported that it took hours to produce some of their plots using multi-GPU machines. They did not report on the resolution of their plots, but we were able to produce all of this work's 22 individual surface plots with a resolution of 125 x 125 in 5 hours and 30 minutes using a single K40 GPU.

## 4. Creating Loss Visualizations Faster

The full test set was used previously to produce these types of loss surface plots, yet the general shape of the graph does not change much after a couple hundred validation examples.

Figure 2 shows the effect of varying the number of test examples to use during evaluation of each specific weight vector. The network being evaluated is a ResNet50 (He et al., 2015) trained on CIFAR10. The overall shape changes a lot for the few first examples. By 10 examples the shape has already taken form, and going from 1k to 3k test examples

has little effect on the plot.

This novel observation can realistically lead to a 100x speedup over contemporary methods, even more for large datasets. For example, all surface plots in this paper used a copious random sample of 250 test set images, where CI-FAR10 has 10k examples in its test set. This leads to a 40x reduction in computation for this dataset.

Figure 3 shows the same evaluation example sweep but using the released ResNet56 with no skip connections (i.e. noshorts). From the first example the chaotic and non-convex nature of the loss landscape is already apparent. The dominant curvatures of the objective landscape surrounding the trained minimizers is vastly different with and without skip connections.

Both ResNet56 models were trained by (Li et al., 2017). They can be obtained from their open-source code here. The preceding experiments used different random directions than the original authors in order to verify that the resulting shape is generally consistent across pairs of randomly sampled directions.

Loss Landscape and the Lottery Ticket Hypothesis



*Figure 4.* Low batch sizes produce shallow and flat minimizers. The solutions' landscapes become progressively deeper and sharper as the batch size increases. Eventually the trend reverses and the minimizers become shallow and flat again. This experiment's data can be viewed with more detail and control with this link.

## 5. Effects of Increasing Batch Size

(Keskar et al., 2017) has hypothesized that larger batch sizes create worse DNNs in part be due to reduced noise in minibatches, making their weight trajectories more susceptible to falling into and getting stuck in exotic basin structures that result in sharp minimizers. This hypothesis was explored by (Li et al., 2017) and the experiment that follows is a natural continuation of theirs.

A LeNet5 (Lecun et al., 1998) architecture with dropout and batch-normalization is trained on FMNIST (Xiao et al., 2017) using Adam (Kingma & Ba, 2017) and a variable number of mini-batch sizes. Figure 4 highlights the effects on the underlying loss surface. Given the faster method of computing these plots we were able to generate 5 highresolution figures, each with a batch size between 2 and 16,000. The best performance lies between a batch size of 160 and 1,600. It would be interesting to continue increasing the batch size to see if it continually gets worse or if there is a cyclic trend in performance.

These results contradict the results of (Li et al., 2017), which had previously suggested that these visualizations created correlations between model generalization and the sharpness or flatness of the landscape around the trained minimizer.

# 6. Lottery Ticket Hypothesis Background

The lottery ticket hypothesis (LTH) (Frankle & Carbin, 2019) states that densely connected DNNs contain sparse sub-network that can exceed the test accuracy of the original network after training on at most the same number of instances. These sparse networks are called winning lottery tickets (WLTs) and do much better than the average random sub-network.

In their work (Frankle & Carbin, 2019) used Algorithm 1 to create winning lottery tickets. The dense network has a mask m applied element-wise to its weights:  $m \odot \theta$ . The mask starts out as all ones. After training converges the smallest p% of the remaining weights are pruned and their equivalent elements in the mask are set to 0. All remaining weights are re-initialized to their original random values and training begins again. This pattern of training, masking, pruning, and re-setting the weights is repeated until the specified sparsity is reached.

The following notation is adopted from (Frankle & Carbin, 2019):

A	lgorithm	<b>1</b> Iterative	Magnitude	Pruning A	Algorithm
	0		0	0	0

- 1. Train network until convergence.
- 2. Remove the smallest 10% of remaining weights per layer & set them to 0.
- 3. Re-initialize the remaining weights to their original values.
- 4. Repeat steps 2 & 3 until the desired sparsity is met.

 $P_m = \frac{\|m\|_0}{|\theta|}$  is the sparsity of m, e.g.,  $P_m = 25\%$  when 75% of weights are pruned.

#### 7. Loss Landscape of Lottery Tickets

Figure 6 details how the loss landscape evolves as more and more weights are pruned via Algorithm 1. The LeNet5 architectures are trained on CIFAR10 using an Adam optimizer (Kingma & Ba, 2017) and mini-batch size of 9.6k. The best models were obtained using early stopping. After every 35 epochs 10% of the remaining weights were pruned. A total of 35 pruning iterations occurred, ending with only 3% of the weights intact. Figure 7 shows the same progression of  $P_m$  but using randomly created masks.

Figure 5 shows the learning trends for both WLTs and random tickets. The results become interesting around  $P_m$ = 50%. It is here that the performance gap between IMP and random pruning begins to occur. It continues to grow all the way through the pruning iterations. By the end of training the gap in test accuracy is 16%. Note the scale on the bottom-right subplot of Figure 7, where  $P_m = 4.2\%$ . At first glance the slope might seem non-trivial, but the z-component's range is much lesser on this plot. Relative to the others, this loss surface is very flat. This might be an exotic structure that prevents training. These "gradient deserts" seem to be dominated by near 0 curvature, meaning small gradients and weight updates, essentially stalling training. Independently confirming that this gradient desert exists is suggested by the author of this research.

Flatness and shallowness in randomly pruned ticket's loss landscape projections is associated with poor generalization. The IMP method consistently produced more convex and sharper minimizers relative to random masking.

#### 8. Discussion

WLTs have been demonstrated to exist in many architecture (Zhou et al., 2020; Frankle & Carbin, 2019; Yu et al., 2020; Zhang et al., 2021) and in different task domains like reinforcement learning (RL) and natural language processing (Yu et al., 2020). Finding WLTs when less than 50% of the network is pruned seems trivial, but what pragmatic benefit do they offer to practitioners? Finding these winning tickets still requires a lot of compute and fine-tuning. If the sample



Figure 5. Training curves of both randomly pruned tickets and tickets derived from Algorithm 1. After around  $P_m = 50\%$  the gap in performance shows. When  $P_m = 3.2\%$  the winning ticket is still performing as well as the original dense network and generalizes twice as well as the random ticket of the same sparsity.

efficiency gains of WLTs can be had from less compute and data, this could impact fields like Deep RL where sample efficiency is essentially non-existent.

(Zhou et al., 2020) found that even better than IMP is keeping weights whose values change by the greatest magnitude, instead of just keeping the largest magnitude weights. (Frankle & Carbin, 2019) unintentionally foreshadowed this when noting in their Appendix F that winning tickets' weights move further than other weights. (Zhou et al., 2020) also showed that re-initializing the weights is not as important as retaining the original signs of the weights, lending even more evidence to the idea that re-initializing to the original values is not vital to finding WLTs.

(Liu et al., 2019) revised some of the experiments from (Frankle & Carbin, 2019) on pruning larger networks while training on ImageNet. This domain usually requires a more exhaustive hyper-parameter search to discover WLTs, but when discovered the results tend to more impressive (Frankle & Carbin, 2019). (Liu et al., 2019) demonstrate that using the more de facto training regime with a larger learning rate and momentum SGD instead of Adam produces better results than even the winning tickets, making LTH even less practical.

(Zhang et al., 2021) deduced that 1 hidden-layer neural networks (with some other assumptions about i.i.d. sampling) should have an enlarged convex region appear around a guaranteed optimal solution when pruned correctly. It is difficult to make sense of this work's results in terms of that hypothesis, but it might be an interesting future research direction.

Loss Landscape and the Lottery Ticket Hypothesis



Figure 6. Pictured are the WLTs of LeNet5 trained on CIFAR10. The tickets perform comparably yet the landscapes get flatter and shallower as pruning continues. The subplot titles are tuples of (mask method,  $P_m$ , test accuracy). You can view the data yourself here.

As originally noted by (Frankle & Carbin, 2019) their IMP method produces sparse networks in such a way that GPUs do not benefit. The sparsity allows more compression, but not faster inference. In part, this is due to the asynchronous checkpoints that occur while GPUs process matrix calculations in parallel. Multiplications by 0 in a certain batch might be faster (i.e. fewer clock cycles) but the slowest multiplication in that batch of operations has to finish before more calculations can be queued.

(Li et al., 2017) experiments using PCA and found that the majority of variance in weight updates lies in only a few dimensions. Perhaps these few very important dimensions are preserved with IMP. It could also be the case that larger learning rates could dislodge the weights from a gradient desert and find more amenable points in the weight space, allowing the recovery of useful training for even randomly masked networks.

The filter-wise normalized random directions method of visualizing the loss landscape has been shown twice not correlating flatness or sharpness with generalization error. However, the convexity of the visualization still seems to correlate with trainability and test performance when the minimizer is deep enough.



*Figure 7.* Objective landscapes from random lottery tickets of LeNet5 trained on CIFAR10. These tickets perform much worse relative to the WLTs obtained using IMP. They are also vastly more shallow. The data is accessible from the web using this link.

#### References

- Dinh, L., Pascanu, R., Bengio, S., and Bengio, Y. Sharp minima can generalize for deep nets, 2017.
- Frankle, J. and Carbin, M. The lottery ticket hypothesis: Finding sparse, trainable neural networks, 2019.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition, 2015.
- Keskar, N. S., Mudigere, D., Nocedal, J., Smelyanskiy, M., and Tang, P. T. P. On large-batch training for deep learning: Generalization gap and sharp minima, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization, 2017.
- Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradientbased learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.
- Li, H., Xu, Z., Taylor, G., Studer, C., and Goldstein, T. Visualizing the loss landscape of neural nets, 2017.
- Liu, Z., Sun, M., Zhou, T., Huang, G., and Darrell, T. Rethinking the value of network pruning, 2019.

- Neyshabur, B., Bhojanapalli, S., McAllester, D., and Srebro, N. Exploring generalization in deep learning, 2017.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.
- Yu, H., Edunov, S., Tian, Y., and Morcos, A. S. Playing the lottery with rewards and multiple languages: lottery tickets in rl and nlp, 2020.
- Zhang, S., Wang, M., Liu, S., Chen, P.-Y., and Xiong, J. Why lottery ticket wins? a theoretical perspective of sample complexity on sparse neural networks, 2021. URL https://openreview.net/forum? id=8pz6GXZ3YT.
- Zhou, H., Lan, J., Liu, R., and Yosinski, J. Deconstructing lottery tickets: Zeros, signs, and the supermask, 2020.