



CENTER FOR
**Brains
Minds+
Machines**

CBMM Memo No. 91

July 24, 2018

Classical Generalization Bounds Are Surprisingly Tight For Deep Networks

Qianli Liao¹, Brando Miranda¹, Andrzej Banburski¹, Jack Hidary² and Tomaso Poggio¹

¹Center for Brains, Minds, and Machines, MIT

²Alphabet (Google) X

Abstract

Deep networks are usually trained and tested in a regime in which the training classification error is not a good predictor of the test error. Thus the consensus has been that generalization, defined as convergence of the empirical to the expected error, does not hold for deep networks. Here we show that, when normalized appropriately after training, deep networks trained on exponential type losses show a good linear dependence of test loss on training loss. The observation, motivated by a previous theoretical analysis of overparametrization and overfitting, not only demonstrates the validity of classical generalization bounds for deep learning but suggests that they are tight. In addition, we also show that the bound of the classification error by the normalized cross entropy loss is empirically rather tight on the data sets we studied.



This material is based upon work supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216.



CENTER FOR
Brains
Minds+
Machines

CBMM Memo No. 91

July 26, 2018

A Surprising Linear Relationship Predicts Test Performance in Deep Networks

Qianli Liao¹, Brando Miranda¹, Andrzej Banburski¹, Jack Hidary² and Tomaso Poggio¹

¹Center for Brains, Minds, and Machines, MIT

²Alphabet (Google) X

Abstract: ¹ Given two networks with the same training loss on a dataset, when would they have drastically different test losses and errors? Better understanding of this question of generalization may improve practical applications of deep networks. In this paper we show that with cross-entropy loss it is surprisingly simple to induce significantly different generalization performances for two networks that have the same architecture, the same meta parameters and the same training error: one can either pretrain the networks with different levels of "corrupted" data or simply initialize the networks with weights of different Gaussian standard deviations. A corollary of recent theoretical results on overfitting shows that these effects are due to an intrinsic problem of measuring test performance with a cross-entropy/exponential-type loss, which can be decomposed into two components both minimized by SGD — one of which is not related to expected classification performance. However, if we factor out this component of the loss, a linear relationship emerges between training and test losses. Under this transformation, classical generalization bounds are surprisingly tight: the empirical/training loss is very close to the expected/test loss. Furthermore, the empirical relation between classification error and normalized cross-entropy loss seem to be approximately monotonic.



This work was supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF - 1231216.

¹A previous version of this paper is [9].

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 4 |
| 2 | Observation: Networks with the Same Training Performance Show Different Test Performance | 4 |
| 3 | Theory: Measuring Generalization Correctly | 6 |
| 3.1 | The Normalized Cross-Entropy Loss | 6 |
| 3.2 | Bounds on the Cross-Entropy Loss imply Bounds on the Classification Error | 8 |
| 4 | Experiments: Normalization Leads to Surprising Tight Generalization | 8 |
| 4.1 | Tight Linearity | 9 |
| 4.2 | Randomly Labeled Training Data | 9 |
| 4.3 | Higher Capacity leads to Higher Test Error | 9 |
| 5 | Discussion | 12 |
| | Appendix A Results on MNIST | 14 |
| | Appendix B Results on CIFAR-100 | 16 |
| | Appendix C L_p norms for normalization | 18 |
| | Appendix D Additional Evidence for a Linear Relationship | 19 |
| | Appendix E Higher Capacity leads to higher Test Error | 22 |
| | Appendix F Numerical values of normalized loss | 23 |
| | Appendix G Deep Neural Network Architecture | 25 |
| G.1 | Three layer network | 25 |
| G.1.1 | Network with 24 filters | 25 |

| | |
|---|-----------|
| G.1.2 Network with 34 filters | 26 |
| G.2 Five layer network | 26 |
| Appendix H Intuition: Shallow Linear Network | 26 |

1 Introduction

Despite many successes of deep networks and a growing amount of research, several fundamental questions remain unanswered, among which is a key puzzle is about the apparent lack of generalization, defined as convergence of the training performance to the test performance with increasing size of the training set. How can a network predict well without generalization? What is the relationship between training and test performances?

In this paper, we investigate the question of why two deep networks with the same training loss have different testing performances. This question is valuable in theory and in practice since training loss is an important clue for choosing among different deep learning architectures. It is therefore worth studying how much training loss can tell us about generalization. In the rest of the paper, we will use the term “error” to mean “classification error” and the term “loss” to mean “cross-entropy loss”, the latter being the objective minimized by stochastic gradient descent (SGD).

In addition to training error and loss, there are many factors (such as choices of network architecture) that can affect generalization performance and we cannot exhaustively study them in this paper. Therefore, we restrict our models to have the same architecture and training settings within each experiment. We tried different architectures in different experiments and observed consistent results.

2 Observation: Networks with the Same Training Performance Show Different Test Performance

First we start with a common observation: even when two networks have the same architecture, same optimization meta parameters and same training loss, they usually have different test performances (i.e. error and loss), presumably because the stochastic nature of the minimization process converge to different minima among the many existing in the loss landscape [11, 13, 12].

With standard settings the differences are usually small (though significant, as shown later). We propose therefore two approaches to magnify the effect:

- Initialize networks with different levels of “random pretraining”: the network is pretrained for a specified number of epochs on “corrupted” training data — the labels of a portion of the examples are swapped with each other in a random fashion.
- Initialize the weights of the networks with different standard deviations of a diagonal Gaussian distribution. As it turns out, different standard deviations yield different test performance.

We show the results of “random pretraining” with networks on CIFAR-10 (Figure 1) and CIFAR-100 (Figure 9) and initialization with different standard deviations on CIFAR-10 (Figure 2) and CIFAR-100 (Figure 10).




Figure 1: *Random Pretraining vs Generalization Performance on CIFAR-10: a 5-layer ConvNet (described in section G.2) is pretrained on training data with partially “corrupted” labels for 30 epochs. It is then trained on normal data for 80 epochs. Among the network snapshots saved from all the epochs we pick a network that is closest to an arbitrarily (but low enough) chosen reference training loss (0.006 here). The number on the x axis indicates the percentage of labels that are swapped randomly. As pretraining data gets increasingly “corrupted”, the generalization performance of the resultant model becomes increasingly worse, even though they have similar training losses and the same zero classification error in training. Batch normalization (BN) is used. After training, the means and standard deviations of BN are “absorbed” into the network’s weights and biases. No data augmentation is performed.*




Figure 2: *Standard Deviation in weight initialization vs generalization performance on CIFAR-10: the network is initialized with weights of different standard deviations. The other settings are the same as in Figure 1. As the norm of the initial weights becomes larger, the generalization performance of the resulting model is worse, even though all models have the same classification error in training.*

3 Theory: Measuring Generalization Correctly

In the previous section, we observed that it is possible to obtain networks that have the same training loss (and error) but very different test performances. This indicates that training loss is not a good proxy of test loss. In general, it is quite common to get zero training error (and a very small training loss) when using overparametrized networks. Therefore, comparing models based on training statistics becomes very difficult. In an extreme case of this situation, recent work [15] shows that an overparametrized deep network can fit a randomly labeled training set and of course fail to generalize at all. In some sense, deep nets seem to “fool” training loss: extremely low training loss can be achieved without any generalization.

Previous theoretical work[12] has provided an answer to the above puzzle. In this paper, we show a corollary of the theory and motivate it in a simple way showing why this “super-fitting” phenomenon happens and under which conditions generalization holds for deep networks.

3.1 The Normalized Cross-Entropy Loss

Notation: We define (as in [12]) a deep network with K layers with the usual elementwise scalar activation functions $\sigma(z) : \mathbf{R} \rightarrow \mathbf{R}$ as the set of functions $f(W; x) = \sigma(W^K \sigma(W^{K-1} \dots \sigma(W^1 x)))$, where the input is $x \in \mathbf{R}^d$, the weights are given by the matrices W^k , one per layer, with matching dimensions. We use the symbol W as a shorthand for the set of W^k matrices $k = 1, \dots, K$. For simplicity we consider here the case of binary classification in which f takes scalar values, implying that the last layer matrix W^K is $W^K \in \mathbf{R}^{1, K_l}$. There are no biases apart from the input layer where the bias is instantiated by one of the

input dimensions being a constant. The activation function in this paper is the ReLU activation.

Consider different zero minima of the empirical risk obtained with the same network on the same training set. Can we predict their expected error from empirical properties only? A natural way to approach the problem of ranking two different minimizers of the empirical risk starts with the “positive homogeneity” property of ReLU networks (see [12]) :

$$f(W^1, \dots, W_k; x) = \rho_1, \dots, \rho_K f(\tilde{W}^1, \dots, \tilde{W}_k; x) \quad (1)$$

where $W_k = \rho_k \tilde{W}_k$ and $\|\tilde{W}_k\| = 1$.

This property is valid for layerwise normalization under any norm. Note that $f(W^1, \dots, W_k; x)$ and $f(\tilde{W}^1, \dots, \tilde{W}_k; x)$ have the same classification performance on any given (test) set. It follows that different empirical minimizers should be compared in terms of their normalized form: the ρ factors affect an exponential type loss – driving it to zero by increasing the ρ s to infinity – but do not change the classification performance which only depends on the sign of $y_n f(x_n)$. Consider the cross-entropy for two classes, given by

$$L = \sum_{n=1}^N \ell(f(x_n), y_n) = \sum_{n=1}^N \ln(1 + e^{-y_n f(x_n)}) = \sum_{n=1}^N \ln(1 + e^{-y_n (\rho_1, \dots, \rho_K) f(\tilde{W}^1, \dots, \tilde{W}_k; x_n)}). \quad (2)$$

What is the right norm to be used to normalize the capacity of deep networks? The “right” normalization should make different minimizers equivalent from the point of view of their intrinsic capacity – for instance equivalent in terms of their Rademacher complexity. The positive homogeneity property implies that the correct norm for a multilayer network should be based on a product of norms. Since different norms are equivalent in \mathbb{R}^n it is enough to use L_2 .

The argument can be seen more directly looking at a typical generalization bound (out of many) of the form [4], where now l is the bounded loss function $l(y, f(W; x))$:

With probability $\geq (1 - \delta) \forall g$

$$|\mathbf{E}(l) - \mathbf{E}_S(l)| \leq c_1 \mathcal{C}_N(\mathbf{L}) + c_2 \sqrt{\frac{\ln(\frac{1}{\delta})}{2N}} \quad (3)$$

where $\mathbf{E}(l)$ is the expected loss, $\mathbf{E}_S(l)$ is the empirical loss, N is the size of the training set, $\mathcal{C}_N(\mathbf{G})$ is an empirical complexity of the class of functions \mathbf{G} on the unit sphere (with respect to each weight per layer) to which l belongs and c_1 and c_2 are constant depending on properties of the loss function. The prototypical example of such a complexity measure is the empirical Rademacher average, typically used for classification in a bound of the form 3. Another example involves covering numbers. We call the bound “tight” if the right hand side – that we call “offset” – is small. Notice that a very small offset, assuming that the bound holds, implies a linear relation with slope 1 between $\mathbf{E}(l)$ and $\mathbf{E}_S(l)$. Thus a very tight bound implies linearity with slope 1. Many learning algorithms have the closely related property of “generalization”: $\lim_{N \rightarrow \infty} |\mathbf{E}(l) - \mathbf{E}_S(l)| \rightarrow 0$ (a notable exception is the nearest neighbor classifier), since in general we hope that $\mathcal{C}_N(\mathbf{G})$ is such to decrease as N increases.

We expect layerwise normalization to yield the same complexity for the different minimizers. The case of linear functions and binary classification is an example. The Rademacher complexity of a set \mathbf{F} of linear functions $f(x) = w^T x$ can be bounded as $\mathcal{R}_N(\mathbf{F}) \leq c \sqrt{\frac{1}{n}} XW$ where X is an L_p bound on the vectors x

and W is an L_q bound on the vectors w with L_p and L_q being dual norms, that is $\frac{1}{p} + \frac{1}{q} = 1$. This includes L_2 for both w and x but also L_1 and L_∞ . Since different L_p norms are equivalent in finite dimensional spaces (in fact $\|x\|_p \geq \|x\|_q$ when $p \leq q$, with opposite relations holding with appropriate constants, e.g. $\|x\|_p \leq n^{1/p-1/q} \|x\|_q$ for $x \in \mathbb{R}^n$), Equation 3 holds under normalization with different L_p norms for w (in the case of linear networks). Notice that the other term in the bound (the last term on the right-hand side of Equation 3) is the same for all networks trained on the same dataset.

The results of [14] – which, as mentioned earlier, was the original motivation behind the arguments above and the experiments below – imply that the weight matrices at each layer converge to the minimum Frobenius norm for each minimizer.

3.2 Bounds on the Cross-Entropy Loss imply Bounds on the Classification Error

Running SGD on the cross-entropy loss in deep networks, as usually done, is a typical approach in machine learning: minimize a convex and differentiable *surrogate* of the 0 – 1 loss function (in the case of binary classification). The approach is based on the fact that the logistic loss (for two classes, cross-entropy becomes the logistic loss) is an upper bound to the binary classification error: thus minimization of the loss implies minimization of the error. One way to formalize these upper bounds is to consider the excess classification risk $R(f) - R^*$, where $R(f)$ is the classification loss associated with f and R^* is the Bayes risk [3]. Let us call $R_\ell(f) = \mathbf{E}(\ell)$ the cross-entropy risk and R_ℓ^* the optimal cross entropy risk. Then the following bound holds in terms of the so-called ψ -transform of the logistic loss ℓ :

$$\psi(R(f) - R^*) \leq R_\ell(f) - R_\ell^* \tag{4}$$

where the ψ function for the exponential loss – a good proxy for the logistic loss – can be computed analytically as $\psi(x) = 1 - \sqrt{1 - x^2}$. Interestingly the shape of $\psi(x)$ for $0 \leq x \leq 1$ is quite similar to the empirical dependence of classification error on normalized cross-entropy loss in figures such as the bottom right of Figure 4.

Unfortunately lower bounds are not available. As a consequence, we cannot prove that among two networks the one with better normalized training cross-entropy (which implies a similar loss at test) guarantees a lower classification error at test. This difficulty is not specific to deep networks or to our main result. It is common to all optimization techniques using a surrogate loss function.

4 Experiments: Normalization Leads to Surprising Tight Generalization

In this section we discuss the experiment results after normalizing the capacity of the networks, as discussed in the previous section. What we observe is a linear relationship between the train loss and the test loss, implying that the expected loss is very close to the empirical loss.

4.1 Tight Linearity

Our observations are that the linear relationship between train loss and test loss hold in a robust way under several conditions:

- *Independence from Initialization:* The linear relationship is independent of whether the initialization is via pretraining on randomly labeled natural images or whether it is via larger initialization, as shown by Figures 4 and 5. In particular, notice that the left of Figure 14 and Figure 15 clearly shows a linear relation for minimizers obtained with the default initialization in Pytorch [10, 7]. This point is important because it shows that the linear relationship is not an artifact of pretraining on random labels.
- *Independence from Network Architecture:* The linear relationship of the test loss and train loss is independent from the network architectures we tried. Figure 3, 16 show the linear relationship for a 3 layer network without batch normalization while Figures 4, 5 show the linear relationship for a 5 layer network with batch normalization on CIFAR10. Additional evidence can be found in the appendix in Figures 13, 14 and 17.
- *Independence from Data Set:* Figures 3, 4, 5, 16, 13, 14, 15, 17 show the linear relationship on CIFAR10 while Figures 11 and 12 show the linear relationship on CIFAR100.
- *Norm independence:* Figures 13 show that the L_p norm used for normalization does not matter – as expected.
- *Normalization is independent of training loss:* Figure 3 shows that networks with different cross-entropy training losses (which are sufficiently large to guarantee zero classification error), once normalized, show the same linear relationship between train loss and test loss.

4.2 Randomly Labeled Training Data

Since the empirical complexity in Equation 3 does not depend on the labels of the training set, Equation 3 should also hold in predicting expected error when the training is on randomly labeled data [16] (though different stricter bounds may hold separately for each of the two cases). The experiments in Figures 17, 4 and 11 show exactly this result. The new datapoint (corresponding to the randomly trained network) is still a tight bound: *the empirical loss is very similar to the test loss*. This is quite *surprising*: the corresponding classification error is zero for training and at chance level for testing!

4.3 Higher Capacity leads to Higher Test Error

Our arguments so far imply that among two unnormalized minimizers of the exponential loss that achieve the same given small loss $L = \epsilon$, the minimizer with higher product of the norms ρ_1, \dots, ρ_K has the higher capacity and thus the highest expected loss. Experiments support this claim, see Figure 18. Notice the linear relationship of test loss with increasing capacity on the top right panels of Figure 4, 5, 11, 12.

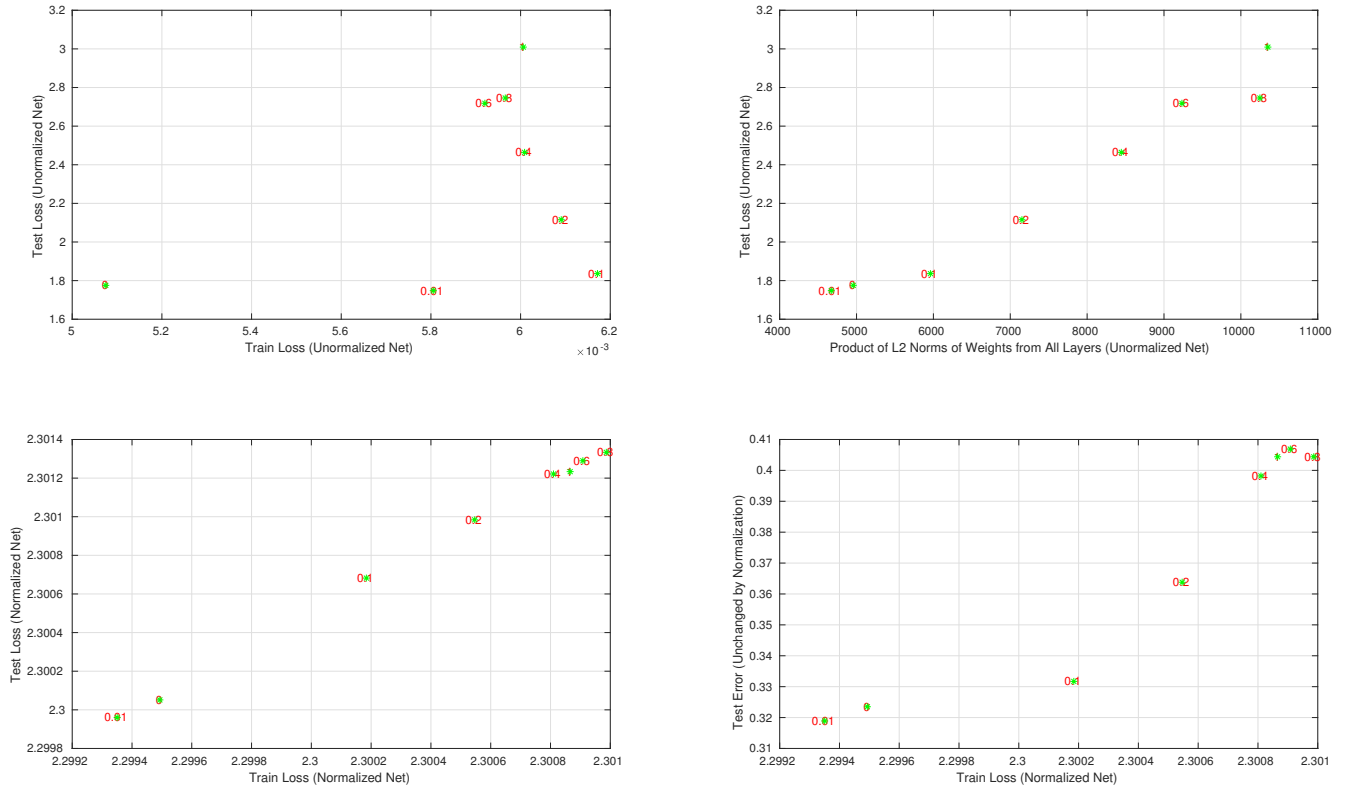


Figure 4: *Random pretraining experiment with batch normalization on CIFAR-10 using a 5 layer neural network as described in section G.2. The red numbers in the figures indicate the percentages of “corrupted labels” used in pretraining. The green stars in the figures indicate the precise locations of the points. **Top left:** Training and test losses of unnormalized networks: there is no apparent relationship. **Top right:** the product of L2 norms from all layers of the network. We observe a positive correlation between the norm of the weights and the testing loss. **Bottom:** under layerwise normalization of the weights (using the Frobenius norm), the classification error does not change (bottom right) while the cross-entropy loss changes (bottom left). There is a surprisingly good linear relationship between training and testing losses, implying tight generalization.*

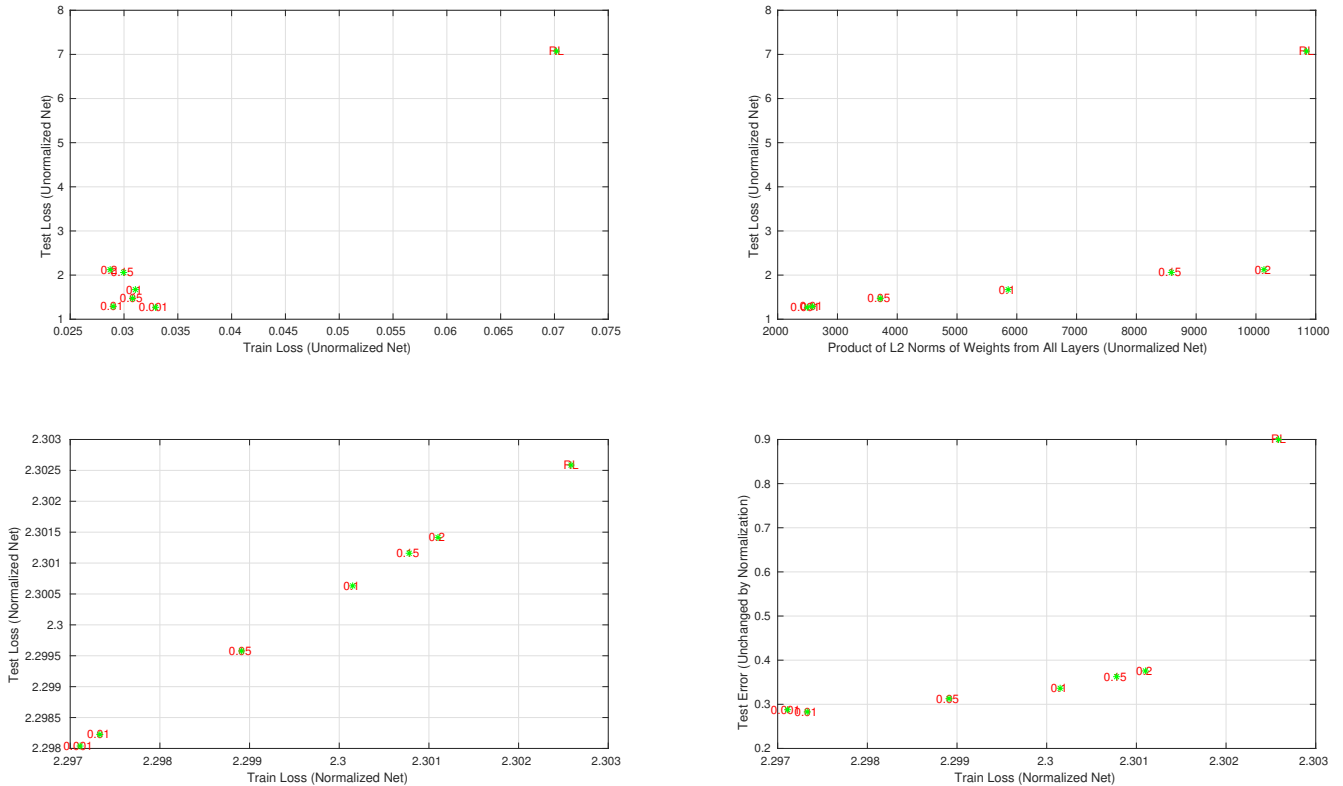


Figure 5: Same as Figure 4 but using different standard deviations for initialization of the weights instead of “random pretraining”. The red numbers in the figures indicate the standard deviations used in initializing weights. The “RL” point (initialized with standard deviation 0.05) refers to training and testing on completely random labels.

5 Discussion

Our results support the theoretical arguments that an appropriate measure of complexity for a deep network should be based on a product norm[1], [2], [8]. The most relevant recent results are by [5] where the generalization bounds depend on the products of layerwise norms without the dependence on the number of layers present in previous bounds. It is also interesting to notice that layerwise Frobenius norm equal to 1 as in our normalized networks avoids the implicit dependence of the network complexity on the number of layers. The results from [12], which triggered this paper, imply products of layerwise Frobenius norms. As we discussed earlier, all L_p norms are equivalent in our setup.

The linear relation we found is quite surprising since it implies that the classical generalization bound of Equation 3 is not only valid but is *tight*: the *offset* in the right-hand side of the Equation is quite small, as shown by Figures 6 and 16 (this also supported by Figures 3, 13 and 15). For instance, the offset in Figure 6 is only 0.0844. This implies of course linearity in our plots of training vs testing normalized cross-entropy loss. It shows that there is indeed generalization - defined as expected loss converging to training loss for large data sets - in deep neural networks. This result contradicts the claims in the title of the recent and influential paper “Understanding deep learning requires rethinking generalization” [16] though those claims *do apply* to the unnormalized loss.

Though it is impossible to claim that better expected cross-entropy loss implies better test error (because the latter is only an upper bound to the classification error), our empirical results in various of the figures

show an approximately monotonic relation between normalized test (or training) cross-entropy loss and test classification error with roughly the shape of the ψ transform of the logistic loss function[3]. Notice, in particular, that the normalized cross-entropy loss in training for the randomly labeled set is close to log 10 which is the cross-entropy loss for chance performance found at test.

Though more experiments are necessary the linear relationship we found seems to hold in a robust way across different types of networks, different data sets and different initializations.

Our results, which are mostly relevant for theory, yield a recommendation for practitioners: it is better to monitor during training the empirical cross-entropy loss of the normalized network instead of the unnormalized cross-entropy loss. The former is the one that matters in terms of stopping time and test performance (see Figures 13 and 14).

More significantly for the theory of Deep Learning, the observations of this paper clearly demand a critical discussion of several commonly held ideas related to generalization such as dropout, SGD and flat minima.




Figure 6: *The left part of the figure shows the test cross-entropy loss plotted against the training loss for the normal, unnormalized network. Notice that all points have **zero classification error at training**. The **RL** points have zero classification error at training and **chance test error**. Once the trained networks are normalized layerwise, the training cross-entropy loss is a very good proxy for the test cross-entropy loss (right plot). The line – regressed just on the 0 and RL points – with slope 0.9642, offset 0.0844, both ordinary and adjusted $R^2 = 0.9999$, root mean square (RMSE) was 6.9797×10^{-5} seems to mirror well Equation 3.*

Acknowledgments

We thank Lorenzo Rosasco, Yuan Yao, Misha Belkin, Youssef Mroueh, Amnon Sashua and especially Sasha Rakhlin for illuminating discussions. Xavier Boix suggested the idea of random pretraining to obtain get different test performances. NSF funding provided by CBMM.

References

- [1] M. Anthony and P. Bartlett. *Neural Network Learning - Theoretical Foundations*. Cambridge University Press, 2002.
- [2] Peter L. Bartlett, Dylan J. Foster, and Matus Telgarsky. Spectrally-normalized margin bounds for neural networks. *CoRR*, abs/1706.08498, 2017.
- [3] Peter L. Bartlett, Michael I. Jordan, and Jon D. McAuliffe. Convexity, classification, and risk bounds. Technical report, JOURNAL OF THE AMERICAN STATISTICAL ASSOCIATION, 2003.
- [4] O. Bousquet, S. Boucheron, and G. Lugosi. Introduction to statistical learning theory. pages 169–207, 2003.
- [5] Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. *CoRR*, abs/1712.06541, 2017.
- [6] Sham M Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Advances in neural information processing systems*, pages 793–800, 2009.
- [7] Yann LeCun, Leo Bottou, G. Orr, and Glaus-Robert Muller. Efficient backprop. 1999.
- [8] Tengyuan Liang, Tomaso Poggio, Alexander Rakhlin, and James Stokes. Fisher-rao metric, geometry, and complexity of neural networks. *CoRR*, abs/1711.01530, 2017.
- [9] Qianli Liao, Brando Miranda, Jack Hidary, and Tomaso Poggio. Classical generalization bounds are surprisingly tight for deep networks. Technical report, Center for Brains, Minds and Machines (CBMM), 2018.
- [10] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [11] T. Poggio and Q. Liao. Theory II: Landscape of the empirical risk in deep learning. *arXiv:1703.09833*, *CBMM Memo No. 066*, 2017.
- [12] T. Poggio, Q. Liao, B. Miranda, A. Banburski, X. Boix, and J. Hidary. Theory IIIb: Generalization in deep networks. *arXiv:1703.09833*, *CBMM Memo No. 090*, 2018.
- [13] T. Poggio, Q. Liao, B. Miranda, L. Rosasco, X. Boix, J. Hidary, and H. Mhaskar. Theory of deep learning III: explaining the non-overfitting puzzle. *arXiv:1703.09833*, *CBMM Memo No. 073*, 2017.
- [14] C. Zhang, Q. Liao, A. Rakhlin, K. Sridharan, B. Miranda, N. Golowich, and T. Poggio. Theory of deep learning IIb: Optimization properties of SGD. *CBMM Memo 072*, 2017.
- [15] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
- [16] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *CoRR*, abs/1611.03530, 2016.

APPENDIX

A Results on MNIST

This section shows figures replicating the main results on the MNIST data set. Figures 7 and 8 show that the linear relationship holds after normalization on the MNIST data set. Figure 8 shows the linear relationship holds after adding the point trained only on random labels.

B Results on CIFAR-100

This section shows figures replicating the main results on CIFAR-100. Figure 9 shows how different test performance can be obtained with pretraining on random labels while Figure 10 shows that different increasing initializations are also effective.

More importantly, Figures 11 and 12 show that the linear relationship holds after normalization, regardless of whether the training was done with pretraining on random labels or with large initialization.




Figure 9: Same as Figure 1, but on CIFAR-100.




Figure 10: Same as Figure 2, but on CIFAR-100.

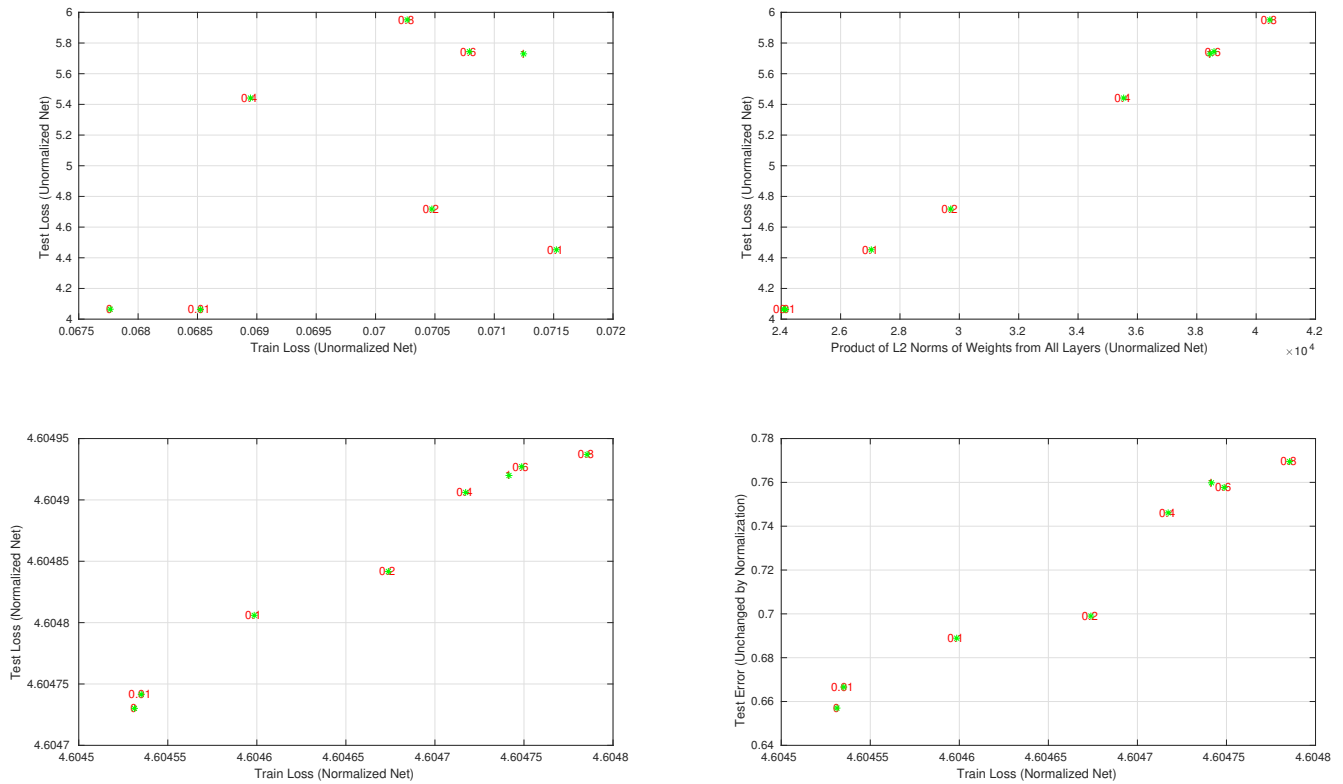


Figure 11: Same as Figure 4 but on CIFAR-100

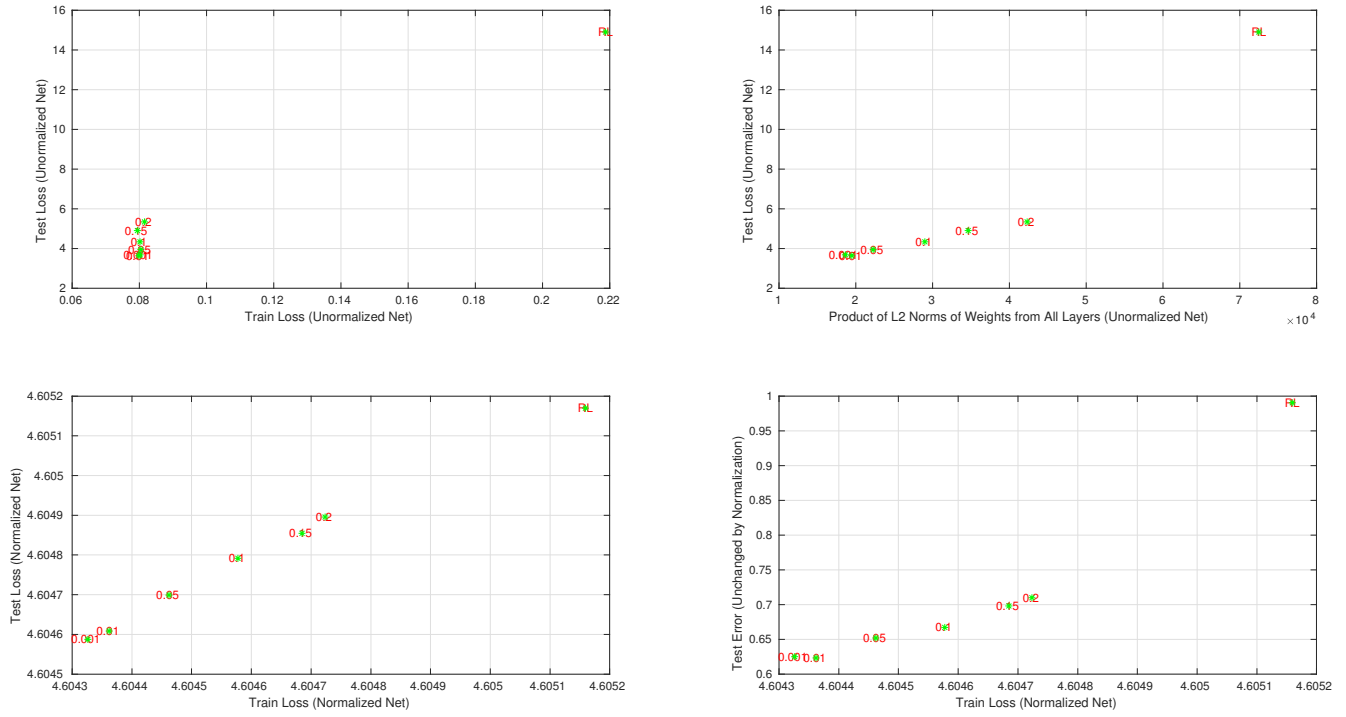


Figure 12: Same as Figure 5 but on CIFAR-100.

C L_p norms for normalization

This section verifies that the effectiveness of layerwise normalization does not depend on which L_p norm is used. Figure 13 provides evidence for this. The constants underlying the equivalence of different norms depend on the dimensionality of the vector spaces.




Figure 13: *Test loss/error vs training loss with all networks normalized layerwise by the L1 norm (divided by 100 to avoid numerical issues because the L1 norms are here very large). The model was a 3 layer neural network described in section G.1.1 and was trained with 50K examples on CIFAR10. The networks were normalized after training each up to epoch 300 and thus different points on the plot correspond to different training losses. Figure 3 shows that the normalized network does not depend on the value of the training loss before normalization. The slope and intercept of the line of best fit are 0.8358 and 0.3783 respectively. The ordinary and adjusted R^2 values are both 0.9998 while the root mean square (RMSE) was 5.6567×10^{-5} . The numbers in the figure indicate the amount of corruption of random labels used during the pretraining.*

D Additional Evidence for a Linear Relationship

This section provides additional experiments showing a linear relationship between the test loss and the training loss after layerwise normalization.

Figure 14 shows that the linear relationship holds if all models are stopped at approximately the same train loss.

Figure 15 is the same as Figure 6, except that the randomly pretrained networks were removed. The important thing to note is that the *offset* is 0.0844, which implies that the 3 is surprisingly tight.




Figure 14: Test loss/error vs training loss with all networks normalized layerwise by the Frobenius norm of the weights. The model was a 3 layer neural network described in section G.1.1 and was trained with 50K examples on CIFAR10. The models were obtained by pretraining on random labels and then by fine tuning on natural labels. SGD without batch normalization was run on all networks in this plot until each reached approximately 0.0044 ± 0.0001 cross-entropy loss on the training data. The numbers in the figure indicate the amount of corruption of the random labels used in pretraining.




Figure 15: Left: test loss vs training loss for networks normalized layerwise by the Frobenius norm. Right: test loss vs training loss for unnormalized networks. The model was a 3 layer neural network described in section G.1.1 and was trained with 50K examples on CIFAR10. In this experiments the networks converged (and had zero train error) but not to the same loss. The networks were trained for 300 epochs. The cross entropy training losses range approximately from 1.29×10^{-3} to 1.38×10^{-3} . The 0s in the figure indicate that there was no random pretraining. The slope and intercept of the line of best fit are 0.8117 and 0.4333 respectively. The ordinary R^2 is 0.9660 and the adjusted R^2 value is 0.9592 while the root mean square (RMSE) was 6.3624×10^{-5} .




Figure 16: *The figure shows the cross-entropy loss on the test set vs the training loss for networks normalized layerwise in terms of the Frobenius norm. The model was a 3 layer neural network described in section G.1.1 and was trained with 50K examples on CIFAR10. All networks were trained for 300 epochs. In this experiments the networks converged (and had zero train error) but not to the same loss. The slope and intercept of the line of best fit are 0.9642 and 0.0844 respectively. The ordinary and adjusted R^2 values are both 0.9999 while the root mean square (RMSE) was 6.9797×10^{-5} . The points labeled 1 were trained on random labels; the training loss was estimated on the same randomly labeled data set. The points marked with 0 were only trained on natural labels.*




Figure 17: The figure shows cross-entropy loss on the test set vs the training loss for networks normalized layerwise in terms of the Frobenius norm. The model was a 3 layer neural network described in section G.1.1 and was trained with 50K examples on CIFAR10. All networks were trained for 300 epochs. In this experiments the networks converged (and had zero train error) but not to the same loss. The slope and intercept of the line of best fit are 0.8789 and 0.2795 respectively. The ordinary and adjusted R^2 values are both 0.9721 while the root mean square (RMSE) was 5.8304×10^{-4} . See Figure 16 for other details.

E Higher Capacity leads to higher Test Error

Figure 18 shows that when the capacity of a network (as measured by the product norm of the layers) increases, so does the test error.

all very close to 1.

As a control experiment refer to Figure 20 for the histogram of the final layer for the most likely class for the unnormalized network. Note that not only are the differences between loss values far from machine precision $\sim 10^{-7.5}$, but also that the output values of the normalized neural networks are far from machine precision too as shown by Figure 19. Notice that normalization does not affect the classification loss. In summary, though the values of the points in the figures differ slightly in terms of loss from each other, those differences are highly significant and reproducible.




Figure 19: *Histogram of the values of $f(x)$ for the most likely class of the layerwise normalized neural network over the 50K images of the MNIST training set. The average value $f(x)$ of the most likely class according to the normalized neural network is 0.026683 with standard deviation 0.007144.*




Figure 20: *Histogram of the values of $f(x)$ for the most likely class of the unnormalized neural network over the 50K images of the MNIST training set. The average value $f(x)$ of the most likely class according to the unnormalized neural network is 60.564373 with standard deviation 16.214078.*

G Deep Neural Network Architecture

G.1 Three layer network

G.1.1 Network with 24 filters

The model is a 3-layer convolutional ReLU network with the first two layers containing 24 filters of size 5 by 5; the final layer is fully connected; only the first layer has biases. There is no pooling.

The network is overparametrized: it has 154,464 parameters (compared to 50,000 training examples).

G.1.2 Network with 34 filters

The model is the same 3-layer convolutional ReLU network as in section G.1.1 except it had 34 units.

The network was still overparametrized: it has 165,784 parameters (compared to 50,000 training examples).

G.2 Five layer network

The model is a 5-layer convolutional ReLU network with (with no pooling). It has in the five layers 32, 64, 64, 128 filters of size 3 by 3; the final layer is fully connected; batch-normalization is used during training.

The network is overparametrized with about 188,810 parameters (compared to 50,000 training examples).

H Intuition: Shallow Linear Network

Consider a shallow linear network, trained once with GD on a set of N data $\{x_i^a, y_i^a\}_{i=1}^N$ and separately on a different set of N data $\{x_i^b, y_i^b\}_{i=1}^N$. Assume that in both cases the problem is linearly separable and thus zero error is achieved on the training set. The question is which of the two solutions will generalize better. The natural approach is to look at the L_2 norm of the solutions: $\sum_n (y_n - w^T x_n)^2$

$$f(w^a, x^a) = \langle w^a, x^a \rangle \quad f(w^b, x^b) = \langle w^b, x^b \rangle \quad (5)$$

In both cases GD converges to zero loss with the minimum norm, maximum margin solution for w . Thus the solution with the larger margin (and the smaller norm) should have a lower expected error. In fact, generalization bounds ([6]) appropriate for this case depend on the product of the norm $\|w\|$ and of a bound on the norm of the data.