Fantastic Generalization Measures and Where to Find Them

Anonymous authors
Paper under double-blind review

Abstract

Generalization of deep networks has been of great interest in recent years, resulting in a number of theoretically and empirically motivated complexity measures. However, most papers proposing such measures only study a small set of models, leaving open the question of whether the conclusion drawn from those experiments would remain valid in other settings. We present the first large scale study of generalization bounds and measures in deep networks. We train over two thousand convolutional networks with systematic changes in commonly used hyperparameters. Hoping to uncover potentially causal relationships between each measure and generalization, we analyze carefully controlled experiments and show surprising failures of some measures as well as promising measures for further research.

1 Introduction

Deep neural networks have seen tremendous success in a number of applications, but why (and how well) these models generalize is still a mystery (Neyshabur et al., 2014; Zhang et al., 2016; Recht et al., 2019). It is crucial to better understand the reason behind the generalization of modern deep learning models; such an understanding has multiple benefits, including providing guarantees for safety-critical scenarios and the design of better models.

A number of papers have attempted to understand the generalization phenomenon in deep learning models from a theoretical perspective e.g. (Neyshabur et al., 2015b; Bartlett et al., 2017; Neyshabur et al., 2018a; Golowich et al., 2017; Arora et al., 2018; Nagarajan & Kolter, 2019a; Wei & Ma, 2019a; Long & Sedghi, 2019). The most direct and principled approach for studying generalization in deep learning is to prove a ***generalization bound*** which is typically an upper bound on the test error based on some quantity that can be calculated on the training set. Unfortunately, finding tight bounds has proven to be an arduous undertaking. While encouragingly Dziugaite & Roy (2017) showed that PAC-Bayesian bounds can be optimized to achieve a reasonably tight generalization bound, current bounds are still not tight enough to accurately capture the generalization behavior. Others have proposed more direct empirical ways to characterize generalization of deep networks without attempting to deriving bounds (Keskar et al., 2016; Liang et al., 2017). However, as pointed by Dziugaite & Roy (2017), empirical correlation does not necessarily translate to a casual relationship between a measure and generalization.

A core component in (theoretical or empirical) analysis of generalization is the notion of ***complexity measure***; a quantity that monotonically relates to some aspect of generalization. More specifically, lower complexity should often imply smaller generalization gap. A complexity measure may depend on the properties of the trained model, optimizer, and possibly training data, but should not have access to a validation set.

Theoretically motivated complexity measures such as VC-dimension, norm of parameters, etc., are often featured as the major components of generalization bounds, where the monotonic relationship between the measures and generalization is mathematically established. In contrast, empirically motivated complexity measures such as sharpness (Keskar et al., 2016) are justified by experimentation and observation. In this work, we do not need to distinguish between theoretically vs empirically motivated measures, and simply refer to both as complexity measures.

Despite the prominent role of complexity measures in studying generalization, the empirical evaluation of these measures is usually limited to a few models, often on toy problems. A measure can only be considered reliable as a predictor of generalization gap if it is tested extensively on many models at a realistic problem size. To this end, we carefully selected a wide range of complexity measures from the literature. Some of the measures are motivated by generalization bounds such as those related to VC-dimension, norm or margin based bounds, and PAC-Bayesian bounds. We further selected a variety of empirical measures such as sharpness (Keskar et al., 2016), Fisher-Rao norm (Liang et al., 2017) and path norms (Neyshabur et al., 2017).

In this study, we trained a large number of models on the CIFAR-10 dataset (Krizhevsky et al., 2014). In order to create a wide range of generalization behaviors, we carefully varied hyperparameters that are believed to influence generalization. We also selected multiple optimization algorithms and looked at different stopping criteria for training convergence. Details of all our measures and hyperparameter selections are provided in Appendix D. Training under all combination of hyperparameters and optimization resulted in a large pool of models. For any such model, we considered 40 complexity measures. The key findings that arise from our large scale study are summarized below:

1. It is easy for some complexity measures to capture spurious correlations that do not reflect more causal insights about generalization; to mitigate this problem, we propose a more rigorous approach for studying them.
2. Many norm-based measures not only perform poorly, but *negatively* correlate with generalization specifically when the optimization procedure injects some stochasticity. In particular, the generalization bound based on the product of spectral norms of the layers (similar to that of Bartlett et al. (2017)) has very strong negative correlation with generalization.
3. Sharpness-based measures such as PAC-Bayesian bounds (McAllester, 1999) bounds and sharpness measure proposed by Keskar et al. (2016) perform the best overall and seem to be promising candidates for further research.
4. Measures related to the optimization procedures such as the gradient noise and the speed of the optimization can be predictive of generalization.

Our findings on the relative success of sharpness-based and optimization-based complexity measures for predicting the generalization gap can provoke further study of these measures.

1.1 Related Work

The theoretically motivated measures that we consider in this work belong to a few different families: PAC-Bayes (McAllester, 1999; Dziugaite & Roy, 2017; Neyshabur et al., 2017); VC-dimension (Vapnik & Chervonenkis, 1971); and norm-based bounds (Neyshabur et al., 2015b; Bartlett et al., 2017; Neyshabur et al., 2018a). The empirically motivated measures from prior literature that we consider are based on sharpness measure (Keskar et al., 2016); Fisher-Rao measure (Liang et al., 2017); distance of trained weights from initialization (Nagarajan & Kolter, 2019b) and path norm (Neyshabur et al., 2015a). Finally, we consider some optimization based measures based on the speed of the optimization method as motivated by (Hardt et al., 2015) and (Wilson et al., 2017a), and the magnitude of the gradient noise as motivated by the work of (Chaudhari & Soatto, 2018) and (Smith & Le, 2017).

A few papers have explored a large scale study of generalization in deep networks. Neyshabur et al. (2017) perform a small scale study of the generalization of PAC-Bayes, sharpness and a few different norms, and the generalization analysis is restricted to correlation. Jiang et al. (2018) studied the role of margin as a predictor of the generalization gap. However, they used a significantly more restricted set of models (e.g. no depth variations), the experiments were not controlled for potential undesired correlation (e.g. the models can have vastly different training errors) and some measures contained parameters that must be learned from the set of models. Novak et al. (2018) conducted large scale study but they only looked at correlation of a few measures to generalization. In contrast, we study thousands of models, and perform controlled experiments to avoid undesired artificial correlations. Our work

focuses on measures that can be computed on a single model and compares a large number of measures across a much wider range of models in a carefully controlled fashion.

1.2 Notation

We denote a probability distribution as $\mathscr{A}$, set as $\mathcal{A}$, tensor as $\mathbf{A}$, vector as $\mathbf{a}$, and scalar as a or α. Let $\mathscr{D}$ denote the data distributions over inputs and their labels, and let κ denote number of classes. We use $\triangleq$ for equality by definition. We denote by $\mathcal{S}$ a given dataset, consisting of m i.i.d tuples $\{(\mathbf{X}_1, y_1), \ldots, (\mathbf{X}_m, y_m)\}$ drawn from $\mathscr{D}$ where $\mathbf{X}_i \in \mathcal{X}$ is the input data and $y_i \in \{1, \ldots, \kappa\}$ the corresponding class label. We denote a feedforward neural network by $f_{\mathbf{w}} : \mathcal{X} \to \mathbb{R}^{\kappa}$, its weight parameters by $\mathbf{w}$, and the the number of weights by $\omega \triangleq \dim(\mathbf{w})$. No activation function is applied at the output (i.e. logits). Denote the weight tensor of the ith layer of the network by $\mathbf{W}_i$, so that $\mathbf{w} = \mathrm{vec}(\mathbf{W}_1, \ldots, \mathbf{W}_d)$, where d is the depth of the network, and vec represents the vectorization operator. Furthermore, denote by $f_{\mathbf{w}}(\mathbf{X})[j]$ the j-th output of the function $f_{\mathbf{w}}(\mathbf{X})$.

Let $\mathcal{R}$ be the set of binary relations, and $I : \mathcal{R} \to \{0, 1\}$ be the indicator function that is 1 if its input is true and zero otherwise. Let L be the 1-0 classification loss over the data distribution $\mathscr{D}$: $L(f_{\mathbf{w}}) \triangleq \mathbb{E}_{(\mathbf{X},y)\sim\mathscr{D}} \left[I\big(f_{\mathbf{w}}(\mathbf{X})[y] \leq \max_{j \neq y} f_{\mathbf{w}}(\mathbf{X})[j]\big) \right]$ and let $\hat{L}$ be the empirical estimate of 1-0 loss over $\mathcal{S}$: $\hat{L}(f_{\mathbf{w}}) \triangleq \frac{1}{m} \sum_{i=1}^{m} I\big(f_{\mathbf{w}}(\mathbf{X})[y_i] \leq \max_{j \neq y_i} f_{\mathbf{w}}(\mathbf{X})[j]\big)$. We refer to $L(f_{\mathbf{w}}) - \hat{L}(f_{\mathbf{w}})$ as the ***generalization gap***. For any input $\mathbf{X}$, we define the sample dependent margin[1] as $\gamma(\mathbf{X}) \triangleq \big(f_{\mathbf{w}}(\mathbf{X})\big)[y] - \max_{i \neq y} f_{\mathbf{w}}(\mathbf{X})_i$. Moreover, we define the overall margin γ as the 10th percentile (a robust surrogate for the minimum) of $\gamma(\mathbf{X})$ over the entire training set $\mathcal{S}$. More notation used for derivation is located in Appendix C.

2 Generalization: What is the goal and how to evaluate?

Generalization is arguably the most fundamental and yet mysterious aspect of machine learning. The core question in generalization is what causes the pair of a ***model*** and an ***optimization*** algorithm to generalize well beyond the training set. There are many hypotheses concerning this question, but what is the right way to compare these hypotheses? Here we briefly discuss some potential approaches to compare different hypotheses:

- Tightness of Generalization Bounds: Proving generalization bounds is very useful to establish the causal relationship between a complexity measure and the generalization error. However, almost all existing bounds are vacuous on current deep models and datasets and therefore, one cannot rely on their proof as an evidence on the causal relationship between a complexity measure and generalization currently[2].
- Regularizing the Complexity Measure: One may evaluate a measure by adding it as a regularizer and directly optimizing it, but this could fail due to two reasons. The complexity measure could change the loss landscape in non-trivial ways and make the optimization more difficult. In such cases, if the optimization fails to optimize the measure, no conclusion can be made about the causality. Another, and perhaps more critical, problem is the existence of implicit regularization of the optimization algorithm. This makes it hard to run a controlled experiment since one cannot simply turn off the implicit regularization; therefore, if optimizing a measure does not improve generalization it could be simply due to the fact that it is regularizing the model in the same way as the optimization is regularizing it implicitly.
- Correlation with Generalization: Evaluating measures based on correlation with generalization is interesting but potentially dangerous. To check the correlation, we should vary architectures and optimization algorithms to produce a set of models. If the set is generated in an artificial way and is not representative of the typical setting, the conclusions might be deceiving and might not generalize to typical

[1] This work only concerns with the output margins, but margin can also be defined at any layer. Please see (Elsayed et al., 2018) and (Wei & Ma, 2019b) for more discussions.

[2] See Dziugaite & Roy (2017) for a non-vacuous generalization bound and related discussions.

cases. One such example is training with different portions of random labels which artificially changes the dataset. Another pitfall is drawing conclusion from changing one or two hyper-parameters (e.g changing the width or batch-size and checking if a measure would correlate with generalization). In these cases, the hyper-parameter could be the true cause of both change in the measure and change in the generalization but the measure itself has no causal relationship with generalization.

While acknowledging all limitations of a correlation analysis, we try to improve the procedure and capture some of the causal effects as much as possible through careful design of experiments and metrics. To this end, we can try to partially capture this causal effect by changing some hyper-parameters and keeping the rest fixed and see if the measure can still correlate with generalization; we also adopt tools from causal inference to uncover the underlying causal structure. Further, to evaluate the effectiveness of a measure, it is crucial to test it on models with a wide range of variations, and each model must be sufficiently trained, if not to complete convergence. For practical reasons, they must also reach this stage within a reasonable time budget.

2.1 Training Models across Hyperparameter Space

In order to create models with different generalization behavior, we consider various hyper-parameter types, which are known or believed to influence generalization (e.g. batch size, dropout rate, etc.). Formally, denote each hyperparameter by θ_i taking values from the set Θ_i, for $i = 1, \ldots, n$ and n denoting the total number of hyperparameter types[3]. For each value of hyperparameters $\boldsymbol{\theta} \triangleq (\theta_1, \theta_2, \ldots, \theta_n) \in \Theta$, where $\Theta \triangleq \Theta_1 \times \Theta_2 \times \cdots \times \Theta_n$, we train the architecture until the training loss (cross entropy value) reaches a given threshold ϵ. See the Appendix A.2 for a discussion on the choice of the stopping criterion. Doing this for each hyper-parameter configuration $\boldsymbol{\theta} \in \Theta$, we obtain a ***total*** of $|\Theta|$ models. The space Θ reflects our prior knowledge about a reasonable hyperparameter space, both in terms of their types and values. Regarding the latter, one could, for example, create Θ_i by grid sampling of a reasonable number of points within a reasonable range of values for θ_i.

2.2 Kendall's Rank-Correlation Coefficient

One way to evaluate the quality of a complexity measure μ is through ***ranking***. Given a set of models resulted by training with hyperparameters in the set Θ, their associated generalization gap $\{g(\boldsymbol{\theta}) \,|\, \boldsymbol{\theta} \in \Theta\}$ and their respective values of the measure, $\{\mu(\boldsymbol{\theta}) \,|\, \boldsymbol{\theta} \in \Theta\}$, our goal is to analyze how consistent a measure (e.g. ℓ_2 `norm` of network weights) is with the empirically observed generalization. To this end, we construct a set $\mathcal{T}$, where each element of the set is associated with one of the trained models. Each element has the form of a pair: complexity measure μ versus generalization gap g.

$$\mathcal{T} \triangleq \cup_{\boldsymbol{\theta} \in \Theta} \big\{ \big(\mu(\boldsymbol{\theta}), g(\boldsymbol{\theta}) \big) \big\} \,. \tag{1}$$

An ideal complexity measure must be such that, for any pair of trained models, if $\mu(\boldsymbol{\theta}_1) > \mu(\boldsymbol{\theta}_2)$, then so is $g(\boldsymbol{\theta}_1) > g(\boldsymbol{\theta}_2)$. We use Kendall's rank coefficient τ to capture to what degree such consistency holds among the elements of $\mathcal{T}$.

$$\tau(\mathcal{T}) \triangleq \frac{1}{|\mathcal{T}|(|\mathcal{T}|-1)} \sum_{(\mu_1, g_1) \in \mathcal{T}} \sum_{(\mu_2, g_2) \in \mathcal{T} \setminus (\mu_1, g_1)} \operatorname{sign}(\mu_1 - \mu_2) \operatorname{sign}(g_1 - g_2) \tag{2}$$

Note that τ can vary between 1 and -1 and attains these extreme values at perfect agreement (two rankings are the same) and perfect disagreement (one ranking is the reverse of the other) respectively. If complexity and generalization are independent, the coefficient becomes zero.

2.3 An evaluation criterion based on controlled experiments

While Kendall's correlation coefficient is an effective tool widely used to capture relationship between 2 rankings of a set of objects, we found that certain measures can achieve high τ

[3]In our analysis we use $n = 7$ hyperparameters: batchsize, dropout probability, learning rate, network depth, weight decay coefficient, network width, optimizer.

values in a trivial manner – i.e. the measure may strongly correlate with the generalization performance without necessarily capturing the cause of generalization. We will analyze this phenomenon in greater details in subsequent sections. To mitigate the effect of spurious correlations, we propose a new quantity for reflecting the correlation between measures and generalization based on a more controlled setting.

None of the existing complexity measures is perfect. However, they might have different sensitivity and accuracy w.r.t. different hyperparameters. For example, `sharpness` may do better than other measures when only a certain hyperparameter (say batch size) changes. To understand such details, in addition to $\tau(\mathcal{T})$, we compute τ for consistency within each hyperparameter axis Θ_i, and then average the coefficient across the remaining hyperparameter space. Formally, we define:

$$m_i \triangleq |\Theta_1 \times \cdots \times \Theta_{i-1} \times \Theta_{i+1} \times \cdots \times \Theta_n| \tag{3}$$

$$\psi_i \triangleq \frac{1}{m_i} \sum_{\theta_1 \in \Theta_1} \cdots \sum_{\theta_{i-1} \in \Theta_{i-1}} \sum_{\theta_{i+1} \in \Theta_{i+1}} \cdots \sum_{\theta_n \in \Theta_n} \tau\left(\cup_{\theta_i \in \Theta_i} \{ (\mu(\boldsymbol{\theta}), g(\boldsymbol{\theta})) \} \right) \tag{4}$$

The inner τ reflects the ranking correlation between the generalization and the complexity measure for a small group of models where the only difference among them is the variation along a single hyperparameter θ_i. Then average the value across all combinations of the other hyperparameter axis. Intuitively, if a measure is good at predicting the effect of hyperparameter θ_i over the model distribution, then its corresponding ψ_i should be high. Then, we compute the ψ_i's average across all hyperparamter axis:

$$\Psi \triangleq \frac{1}{n} \sum_{i=1}^{n} \psi_i \tag{5}$$

If a measure achieves a high Ψ on a given model distribution Θ, then it should achieve high individual ψ across all hyperparameters. A measure that excels at predicting changes in a single hyperparameter (high ψ_i) but fails at the other hyperparameters (low ψ_j for all $j \neq i$) will not do well on Ψ. On the other hand, if the measure performs well on Ψ, it means that the measure can reliably rank the generalization for each of the hyper-parameter changes.

A ***thought experiment*** to illustrate why Ψ captures a *better* causal nature of the generalization than Kendall's τ is as follows. Suppose there exists a measure that perfectly captures the depth of the network while producing random prediction if 2 networks have the same depth, this measure would do reasonably well in terms of τ but much worse in terms of Ψ. In the experiments we consider in the following sections, we found that such a measure would achieve overall $\tau = 0.362$ but $\Psi = 0.11$.

Finally, we acknowledge this this measure is only a small step towards the difficult problem of capturing the causal relationship between complexity measures and generalization in empirical settings, and we hope this encourages future work in this direction.

3 Generating a family of trained models

We chose 7 common hyperparameters related to optimization and architecture design, with 3 choices for each hyperparameter, to generate $3^7 = 2187$ models. We trained these models to convergence. Convergence criterion is chosen as when cross entropy loss reaches the value 0.01. Any model that was not able to achieve this value of cross-entropy[4] was discarded from further analysis. The latter is different from the DEMOGEN dataset (Jiang et al., 2018) where the models are not trained to the same cross-entropy. Putting the stopping criterion on *the training loss* rather than *the number of epochs* is crucial since otherwise one can simply use cross-entropy loss value to predict generalization. Please see Appendix Section A.2 for a discussion on the choice of stopping criterion.

To construct a pool of trained models with vastly different generalization behaviors while being able to fit the training set, we covered a wide range of hyperparameters for training. Our base model is inspired by the Network-in-Network (Gao et al., 2011). The hyperparameter categories we test on are: weight decay coefficient (`weight_decay`), width of the layer

[4]In our analysis, less than 5 percent of the models do not reach this threshold.

(`width`), mini-batch size (`batchsize`), learning rate (`learning_rate`), dropout probability (`dropout`), depth of the architecture (`depth`) and the choice of the optimization algorithms (`opt`). We select 3 choices for each hyperparameter (i.e. $|\Theta_1| = 3\,\forall i$), but different optimization algorithms may use different hyperparameter sets such as learning rate or weight regularization. Please refer to Appendix A.3 for the details on the models, and Appendix A.1 for the reasoning behind the design choices.

Figure 1 shows some summarizing statistics of the models in this study. On the left we show the number of models that achieve above 99% training accuracy for every individual hyperparameter choice. Since we have $3^7 = 2187$ models in total, the maximum number of model for each hyperparameter type is $3^{7-1} = 718$; the majority of the models in our pool were able to reach this threshold. In the middle we show the distribution of the cross-entropy value over the entire training set. While we want the models to be at exactly 0.01 cross-entropy, in practice it is computationally prohibitive to constantly evaluate the loss over the entire training set; further, to enable reasonable temporal granularity, we estimate the training loss with 100 randomly sampled minibatch. These computational compromises result in long-tailed distribution of training loss centered at 0.01. As shown in Table 1, even such minuscule range of cross-entropy difference could lead to positive correlation with generalization, highlighting the importance of training loss as a stopping criterion. On the right, we show the distribution of the generalization gap. We see that while all the models' training accuracy is above 0.99, there is a wide range of generalization gap, which is ideal for evaluating complexity measures.

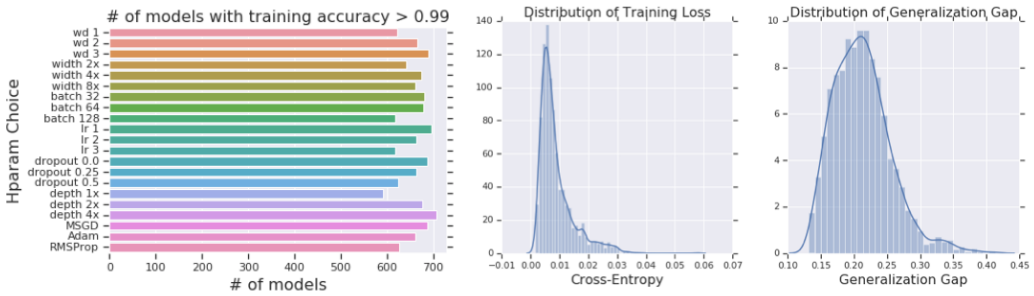

Figure 1: **Left**: Number of models with training accuracy above 0.99 for each hyperparameter type. **Middle**: Distribution of training cross-entropy; distribution of training error can be found in Fig. 2. **Right**: Distribution of generalization gap.

4 Baselines

The first baseline we consider is performance of a measure against an oracle who observes the noisy generalization gap. Concretely, we rank the models based on the true generalization gap with some additive noise. The resulting ranking correlation indicates how close the performances of all models are. As the scale of the noise approaches 0, the oracle's prediction tends towards perfect (i.e. 1). This baseline accounts for the potential noise in the training procedure and gives an anchor for gauging the difficulty of each hyperparameter type. Formally, given a set of hyper-parameters Θ', we define oracle ϵ to be the expectation of τ or Ψ where the measure is $\{g(\boldsymbol{\theta}) + \mathcal{N}(0, \epsilon^2) \,|\, \boldsymbol{\theta} \in \Theta'\}$. We report the performance of the noisy oracle in Table 1 for $\epsilon \in \{0.01, 0.02, 0.05, 0.1\}$.

Second, to understand how our hyperparameter choices affect the optimization, we give each hyperparameter type a *canonical order* which is believed to have correlation with generalization (e.g. larger learning rate generalizes better) and measure their τ. The exact canonical ordering can be found in Appendix A.4. Note that unlike other measures, each canonical ordering can only predict generalization for its own hyperparameter type, since its corresponding hyperparameter remains fixed in any other hyperparameter type; consequently, each column actually represents different measure for the *canonical measure* row. Assuming that each canonical measure is uninformative of any other canonical measures, the Ψ metric for each canonical measure is $\frac{1}{7}$ of its performance on it corresponding hyperparameter type.

We next look at of the most well-known complexity measures in machine learning is the VC-Dimension. Bartlett et al. (2019) proves bounds on the VC dimension of piece-wise

linear networks with potential weight sharing. In Appendix D.1, we extend their result to include pooling layers and multi-class classification. We report two measure based on VC-dimension bounds and parameter counting. These measures could only be predicative when the architecture changes which only happens in depth and width hyperparameter type. We observe that in both types VC-dimension and number of parameters are negatively correlated with generalization gap which confirms the widely known empirical observation that overparametrization improves generalization in deep learning.

Finally, we report the measures that only look at the output of the network. In particular, we look at the cross-entropy loss, margin γ and the entropy of the output. All three measures are closely related to each other. In fact, the results in Table 1 also reflects this similarity. The results confirm the general understanding that larger margin, lower cross-entropy and higher entropy would lead to better generalization. Please see Appendix D.1.1 for definitions and more discussions on these measures.

	batch size	dropout	learning rate	depth	optimizer	weight decay	width	overall τ	Ψ
vc dim 19	0.000	0.000	0.000	-0.909	0.000	0.000	-0.171	-0.251	-0.154
# params 20	0.000	0.000	0.000	-0.909	0.000	0.000	-0.171	-0.175	-0.154
$1/\gamma$ 22	0.312	-0.593	0.234	0.758	0.223	-0.211	0.125	0.124	0.121
-entropy 23	0.346	-0.529	0.251	0.632	0.220	-0.157	0.104	0.148	0.124
cross-entropy 21	0.440	-0.402	0.140	0.390	0.149	0.232	0.080	0.149	0.147
oracle 0.01	0.590	0.850	0.706	0.902	0.550	0.630	0.509	0.850	0.677
oracle 0.02	0.380	0.657	0.536	0.717	0.374	0.388	0.360	0.714	0.487
oracle 0.05	0.172	0.375	0.305	0.384	0.165	0.184	0.204	0.438	0.256
oracle 0.1	0.095	0.183	0.212	0.245	0.136	0.117	0.104	0.262	0.156
canonical ordering	0.652	0.969	0.733	0.909	-0.055	0.735	0.171	N/A	N/A

Table 1: Numerical results for Baselines and Oracular measures. Columns under each hyperparameter contain the values of ψ_i (4) for that hyperparameter; overall τ contains the value of (2) for all models; and Ψ contains value of (5). Red numbers next to each measure's name lead to the actual definition of the measure.

5 Surprising Failure of some Norm/margin-based Measures

In machine learning, a long standing measure for quantifying the complexity of a function, and therefore generalization, is using various norms of the given function. Some of the norms can be directly optimized to improve the generalization. For example, ℓ_2 regularization on the parameters of a model can be seen as imposing an isotropic Gaussian prior over the parameters in maximum a posteriori estimation. We choose several representative norms or measures based on norms and compute our correlation coefficient between the measures and the generalization gap of the model. We study the following measures and their variants (Table 2): *spectral bound*, *Frobenius distance from initialization*, ℓ_2 *Frobenius norm of the parameters*, *Fisher-Rao metric* and *path norm*.

	batch size	dropout	learning rate	depth	optimizer	weight decay	width	overall τ	Ψ
Frob distance 40	-0.317	-0.833	-0.718	0.526	-0.214	-0.669	-0.166	-0.263	-0.341
Spectral orig 26	-0.355	-0.745	-0.683	-0.909	-0.150	-0.047	-0.256	-0.534	-0.449
Parameter norm 42	0.236	-0.516	0.174	0.330	0.187	0.124	-0.170	0.073	0.052
Path norm 44	0.252	**0.270**	0.049	**0.934**	0.153	0.338	**0.178**	0.373	0.311
Fisher-Rao 45	0.396	0.147	0.240	-0.553	0.120	**0.551**	0.177	0.078	0.154
sum-of-fro/margin 38	**0.520**	-0.369	**0.598**	0.882	**0.380**	0.738	-0.080	**0.391**	**0.381**
oracle 0.02	0.380	0.657	0.536	0.717	0.374	0.388	0.360	0.714	0.487

Table 2: Numerical results for selected Norm/Margin-based Measures

Spectral bound: The most surprising observation here is that the spectral complexity is strongly negatively correlated with generalization, and negatively correlated with changes within every hyperparameter type. Most notably, it has strong negative correlation with the depth of the network, which may suggest that the largest singular values are not sufficient to capture the capacity of the model. To better understand the reason behind this observation, we investigate using different components of the spectral complexity as the measure. An interesting observation is that the Frobenius distance to initialization is negatively correlated, but the Frobenius norm of the parameters is slightly positively correlated with generalization, which contradicts some theories suggesting solutions closer to initialization should generalize better. A tempting hypothesis is that weight decay favors solution closer to the origin, but we did an ablation study on only models with 0 weight decay and found that the distance from initialization still correlates negatively with generalization.

These observations correspond to choosing different reference matrices $\mathbf{W}_i^0$ for the bound: the distance corresponds to using the initialization as the reference matrices while the Frobenius norm of the parameters corresponds to using the origin as the reference. Since the Frobenius norm of the parameters shows better correlation, we use zero reference matrices in the spectral bound. This improved both τ and Ψ, albeit still negative. In addition, we extensively investigate the effect of different terms of the Spectral bound to isolate the effect; however, the results do not improve. We also found that the measure *sum-of-fro/margin* performs particularly well. This measure is invariant to rebalancing of layers but does not correspond to any existing bound. We believe investigating this quantity is an interesting direction for future work. These experiments can be found in the Appendix D.2.

Path norm: While path-norm is a proper norm in the function space but not in parameter space, we observe that it is positively correlated with generalization in all hyper-parameter types and achieves comparable τ (0.373) and Ψ (0.311).

Fisher-Rao metric: The Fisher-Rao metric is a lower bound (Liang et al., 2017) on the path norm that has been recently shown to capture generalization. We observed that it overall shows worse correlation than the path norm; in particular, it is negatively correlated ($\tau = -0.553$) with the depth of the network while contrast with path norm which properly captures the effect of depth on generalization. A more interesting observation is that the Fisher-Rao metric achieves a positive $\Psi = 0.154$ but its $\tau = 0.078$ is essentially at chance. This may suggest that the metric can capture a single hyper-parameter change but is not able to capture the interactions between different hyperparameter type.

Effect of Randomness: Dropout and batch size (first 2 columns of Table 2) directly introduce randomness into the training dynamic. For batch size, we observed that the Frobenius displacement and spectral complexity both correlate negatively with the changes in batch size while the Frobenius norm of the parameters correlates positively with generalization. On the other hand, when changes happen to the magnitude dropout probability, we observed that all of the proper norms are negatively correlated with the generalization changes. Since increasing dropout usually reduces the generalization gap, this implies that increasing the dropout probability may be at least partially responsible for the growth in these norms. This is unexpected since increasing norm in principle implies higher model capacity which is usually more prone to overfitting.

6 Success of Sharpness-based Measures

A natural category of generalization measures is centered around the concept of "***sharpness***" of the local minima, capturing the sensitivity of the empirical risk (i.e. the loss over the entire training set) to perturbations in model parameters. Such notion of stability under perturbation is captured elegantly by the PAC-Bayesian framework (McAllester, 1999) which has provided promising insights for studying generalization of deep neural networks (Dziugaite & Roy, 2017; Neyshabur et al., 2017; 2018a). In this sections, we investigate PAC-Bayesian generalization bounds and several of its variants which rely on different priors and different notions of sharpness (Table 3).

In order to evaluate a PAC-Bayesian bound, one needs to come up with a prior distribution over the parameters that is chosen in advance before observing the training set. Then, given any posterior distribution on the parameters which could depend on the training set, a PAC-Bayesian bound (Theorem 46) states that the expected generalization error of the parameters generated from the posterior can be bounded by the KL-divergence of the prior and posterior. The posterior distribution can be seen as adding perturbation on final parameters. Dziugaite & Roy (2017) shows contrary to other generalization bounds, it is possible to calculate non-vacuous PAC-Bayesian bounds by optimizing the bound over a large set of Gaussian posteriors. Neyshabur et al. (2017) demonstrates that PAC-Bayesian bounds where prior and posterior are isotropic Gaussian distributions are good measure of generalization on small scale experiments (e.q. 47).

PAC-Bayesian framework captures sharpness in the expected sense since we add randomly generated perturbations to the parameters. Another possible notion of sharpness is the worst-case sharpness where we search for the direction that changes the loss the most. This is motivated by (Keskar et al., 2016) where they observe that this notion would correlate to

generalization in the case of different batch sizes. We can use PAC-Bayesian framework to construct generalization bounds for this worst-case perturbations as well. We refer to this worst case bound as the sharpness bound (eq. 50). The main component in both PAC-Bayes and worst-case sharpness bounds is the ratio of norm of parameters to the magnitude of the perturbation where the magnitude is chosen to be the largest number so that the training error of the perturbed model is at most 0.1. While mathematically, the sharpness bound should always yield higher complexity than the Pac-Bayes bound, we observed that the former has higher correlation both in terms of τ and Ψ. In addition, we measured inverse of perturbation magnitude as a measure removing the norm in the numerator to compare it with the bound and did not observe a significant difference.

	batch size	dropout	learning rate	depth	optimizer	weight decay	width	overall τ	Ψ
sharpness 52	0.542	-0.359	0.716	0.816	0.297	0.591	0.185	0.400	0.398
pacbayes 49	0.526	-0.076	0.705	0.546	0.341	0.564	-0.086	0.293	0.360
sharpness mag 60	0.401	-0.514	0.321	-0.909	0.181	0.281	-0.171	-0.158	-0.059
pacbayes mag 57	0.532	-0.480	0.508	**0.902**	0.188	0.155	0.186	0.410	0.284
$1/\sigma$ sharpness 54	0.532	-0.326	0.711	0.776	0.296	0.592	0.263	0.399	0.406
$1/\sigma$ pacbayes 53	0.501	-0.033	0.744	0.200	**0.346**	0.609	0.056	0.303	0.346
$1/\alpha'$ sharpness mag 62	**0.570**	**0.148**	**0.762**	0.824	0.297	**0.741**	**0.269**	**0.484**	**0.516**
$1/\sigma'$ pacbayes mag 61	0.490	-0.215	0.505	0.896	0.186	0.147	0.195	0.365	0.315
oracle 0.02	0.380	0.657	0.536	0.717	0.374	0.388	0.360	0.714	0.487

Table 3: Numerical results for selected sharpness-based Measures; all the measure use the origin as the reference and `mag` refers to magnitude-aware version of the measure.

Magnitude-aware Perturbation Bounds Perturbing the parameters without taking their magnitude into account can cause many of them to switch signs. Therefore, one cannot apply large perturbations to the model without changing the loss significantly. One possible modification to improve the perturbations is to choose the perturbation magnitude based on the magnitude of the parameter. In that case, it is guaranteed that if the magnitude of perturbation is less than that of the parameter, then the sign of the parameter does not change. Following Keskar et al. (2016), we pick the magnitude of the perturbation with respect to the magnitude of parameters. We formalize this notion of importance based on magnitude and derive two alternative generalization bounds for expected sharpness (eq. 55) and worst case sharpness (eq. 58) that include the magnitude of the parameters into the prior. Formally, we design α' and σ' to be the ratio of parameter magnitude to the perturbation magnitude in sharpness and PAC-Bayes bounds respectively. While this change did not improve upon the original PAC-Bayesian measures, we observed that looking at $1/\alpha'$ has surprising predictive power in terms of the generalization which surpasses the performance of oracle 0.02. This measure is close the one suggested in Keskar et al. (2016).

Finding σ In case of models with extremely small loss, the perturbed loss should roughly increase monotonically w.r.t the perturbation scale. Leveraging this observation, we design algorithms for computing the perturbation scale σ such that the first term on the RHS is as close to a fixed value as possible for all models. In our experiments, we choose the deviation to be 0.1 which translates to 10% training error. These search algorithms are paramount to compare measures between different models. We provide the detailed algorithms in the Appendix E. To improve upon our algorithms, one could try a computational approach similar to Dziugaite & Roy (2017) to obtain a numerically better bound which may result in stronger correlation; however, due to practical computational constraints, we could not do so for the large number of models we consider.

7 Potential of Optimization-based Measures

Optimization is an indispensable component of deep learning. Numerous optimizers have been proposed for more stable training and faster convergence. How the optimization scheme and speed of optimization influence generalization of a model has been a topic of contention among the deep learning community (Merity et al., 2017; Hardt et al., 2015). We study 3 representative optimizers `Momentum SGD`, `Adam`, and `RMSProp` with different initial learning rates in our experiments to thoroughly evaluate this phenomenon. We also consider other optimization related measures that are believed to correlate with generalization. These include (Table 4):

1. Number of iterations required to reach cross-entropy equals 0.1

2. Number of iterations required going from cross-entropy equals 0.1 to cross-entropy equals 0.01
3. Variance of the gradients after only seeing the entire dataset once (1 epoch)
4. Variance of the gradients when the cross-entropy is approximately 0.01

	batch size	dropout	learning rate	depth	optimizer	weight decay	width	overall τ	Ψ
step to 0.1 63	-0.664	-0.861	-0.255	**0.440**	-0.030	-0.628	0.043	-0.264	-0.279
step 0.1 to 0.01 64	-0.151	-0.069	-0.014	0.114	0.072	-0.046	-0.021	-0.088	-0.016
grad noise 1 epoch 65	0.071	**0.378**	0.376	-0.517	0.121	0.221	0.037	0.070	0.098
grad noise final 66	**0.452**	0.119	**0.427**	0.141	**0.245**	**0.432**	**0.230**	**0.311**	**0.292**
oracle 0.02	0.380	0.657	0.536	0.717	0.374	0.388	0.360	0.714	0.487

Table 4: Optimization-Based Measures

The number of iterations roughly characterizes the speed of optimization, which has been argued to correlate with generalization. For the models we consider, we observed that the initial phase (step to reach cross-entropy value of 0.1) of the optimization is negatively correlated with the speed of optimization for both τ and Ψ. This would suggest that the difficulty of optimization during the initial phase optimization benefits the final generalization. On the other hand, the speed of optimization going from cross-entropy 0.1 to cross-entropy 0.01 does not seem to be correlated with the generalization of the final solution. On the other hand, the speed of optimization is not an explicit capacity measure so either positive or negative correlation could potentially be informative.

On the other hand, the variance of the gradients characterizes the smoothness of the loss surface. Towards the end of the training, the variance of the gradients also captures a particular type of "flatness" of the local minima. This measure is surprisingly predictive of the generalization both in terms of τ and Ψ, and more importantly, is positively correlated across every type of hyperparameter. To the best of our knowledge, this is the first time this phenomenon has been observed. The connection between variance of the gradient and generalization is perhaps natural since much of the recent advancement in deep learning such as residual network (He et al., 2016) or batch normalization have enabled using larger learning rates to train neural networks. Stability with higher learning rates implies smaller noises in the minibatch gradient. We hope that our work encourages future works in other possible measures based on optimization and during training.

8 Conclusion

We conducted large scale experiments to test the correlation of different measures with the generalization of deep models and propose a framework to better disentangle the cause of correlation from spurious correlation. We confirmed the effectiveness of the PAC-Bayesian bounds through our experiments and corroborate it as a promising direction for cracking the generalization puzzle. Further, we provide an extension to existing PAC-Bayesian bounds that consider the importance of each parameter. We also found that several measures related to optimization are surprisingly predictive of generalization and worthy of further investigation. On the other hand, several surprising failures about the norm-based measures were uncovered. In particular, we found that regularization that introduces randomness into the optimization can increase various norm of the models and spectral complexity related norm-based measures are unable to capture generalization – in fact, most of them are negatively correlated. Our experiments demonstrate that the study of generalization measure can be misleading when the number of models studied is small and the metric of quantifying the relationship is not carefully chosen. We hope this work will incentivize more rigorous treatment of generalization measures in future work.

To the best of our knowledge, this work is one of the most comprehensive study of generalization to date, but there are a few short-comings. Due to computational constraints, we were only able to study 7 most common hyperparameter types and relatively smaller architectures, which do not reflect the models used in production. Indeed, if more hyperparameters are considered, one could expect to better capture the causal relationship. We also only studied classification models on convolutional networks. Future work will aim to address these limitations.

References

Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. Stronger generalization bounds for deep nets via a compression approach. *arXiv preprint arXiv:1802.05296*, 2018.

Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.

Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, pp. 6240–6249, 2017.

Peter L. Bartlett, Nick Harvey, Christopher Liaw, and Abbas Mehrabian. Nearly-tight vc-dimension and pseudodimension bounds for piecewise linear neural networks. *Journal of Machine Learning Research*, 20(63):1–17, 2019. URL http://jmlr.org/papers/v20/17-612.html.

Pratik Chaudhari and Stefano Soatto. Stochastic gradient descent performs variational inference, converges to limit cycles for deep networks. In *2018 Information Theory and Applications Workshop (ITA)*, pp. 1–10. IEEE, 2018.

Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. Sharp minima can generalize for deep nets. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1019–1028. JMLR. org, 2017.

Gintare Karolina Dziugaite and Daniel M Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *arXiv preprint arXiv:1703.11008*, 2017.

Gamaleldin Elsayed, Dilip Krishnan, Hossein Mobahi, Kevin Regan, and Samy Bengio. Large margin deep networks for classification. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 31*, pp. 842–852. Curran Associates, Inc., 2018. URL http://papers.nips.cc/paper/7364-large-margin-deep-networks-for-classification.pdf.

Jianxi Gao, Sergey V Buldyrev, Shlomo Havlin, and H Eugene Stanley. Robustness of a network of networks. *Physical Review Letters*, 107(19):195701, 2011.

Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. *arXiv preprint arXiv:1712.06541*, 2017.

Moritz Hardt, Benjamin Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. *arXiv preprint arXiv:1509.01240*, 2015.

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*, pp. 770–778, 2016.

Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015. URL http://arxiv.org/abs/1502.03167.

Yiding Jiang, Dilip Krishnan, Hossein Mobahi, and Samy Bengio. Predicting the generalization gap in deep networks with margin distributions. *arXiv preprint arXiv:1810.00113*, 2018.

Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.

Aryeh Kontorovich. Dudley-pollard packing theorem. http://aiweb.techfak.uni-bielefeld.de/content/bworld-robot-control-software/, 2016.

Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. The cifar-10 dataset. *online: http://www. cs. toronto. edu/kriz/cifar. html*, 55, 2014.

Tengyuan Liang, Tomaso Poggio, Alexander Rakhlin, and James Stokes. Fisher-rao metric, geometry, and complexity of neural networks. *arXiv preprint arXiv:1711.01530*, 2017.

Philip M Long and Hanie Sedghi. Size-free generalization bounds for convolutional neural networks. *arXiv preprint arXiv:1905.12600*, 2019.

David A McAllester. Pac-bayesian model averaging. In *COLT*, volume 99, pp. 164–170. Citeseer, 1999.

Stephen Merity, Nitish Shirish Keskar, and Richard Socher. Regularizing and optimizing lstm language models. *arXiv preprint arXiv:1708.02182*, 2017.

Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. Foundations of machine learning. adaptive computation and machine learning. *MIT Press*, 31:32, 2012.

Vaishnavh Nagarajan and J Zico Kolter. Deterministic pac-bayesian generalization bounds for deep networks via generalizing noise-resilience. *arXiv preprint arXiv:1905.13344*, 2019a.

Vaishnavh Nagarajan and J Zico Kolter. Generalization in deep networks: The role of distance from initialization. *arXiv preprint arXiv:1901.01672*, 2019b.

Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *arXiv preprint arXiv:1412.6614*, 2014.

Behnam Neyshabur, Ruslan R Salakhutdinov, and Nati Srebro. Path-sgd: Path-normalized optimization in deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 2422–2430, 2015a.

Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pp. 1376–1401, 2015b.

Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, pp. 5947–5956, 2017.

Behnam Neyshabur, Srinadh Bhojanapalli, and Nathan Srebro. A pac-bayesian approach to spectrally-normalized margin bounds for neural networks. *International Conference on Learning Representations*, 2018a.

Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. Towards understanding the role of over-parametrization in generalization of neural networks. *arXiv preprint arXiv:1805.12076*, 2018b.

Roman Novak, Yasaman Bahri, Daniel A Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. Sensitivity and generalization in neural networks: an empirical study. *arXiv preprint arXiv:1802.08760*, 2018.

Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548*, 2017.

Konstantinos Pitas, Mike Davies, and Pierre Vandergheynst. Pac-bayesian margin bounds for convolutional neural networks. *arXiv preprint arXiv:1801.00171*, 2017.

Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? *arXiv preprint arXiv:1902.10811*, 2019.

J. K. Rowling. Fantastic beasts and where to find them. In David Yates (ed.), *Harry Potter film series*. WarnerBros., 2016. URL http://www.fantasticbeasts.com.

Hanie Sedghi, Vineet Gupta, and Philip M. Long. The singular values of convolutional layers. *CoRR*, abs/1805.10408, 2018. URL http://arxiv.org/abs/1805.10408.

Samuel L Smith and Quoc V Le. A bayesian perspective on generalization and stochastic gradient descent. *arXiv preprint arXiv:1710.06451*, 2017.

Vladimir N Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Theory of probability and its applications*, pp. 11–30. Springer, 1971.

Thomas Verma and Judea Pearl. Equivalence and synthesis of causal models. In *Proceedings of the Sixth Annual Conference on Uncertainty in Artificial Intelligence*, UAI '90, pp. 255–270, New York, NY, USA, 1991. Elsevier Science Inc. ISBN 0-444-89264-8. URL http://dl.acm.org/citation.cfm?id=647233.719736.

Colin Wei and Tengyu Ma. Data-dependent sample complexity of deep neural networks via lipschitz augmentation. *arXiv preprint arXiv:1905.03684*, 2019a.

Colin Wei and Tengyu Ma. Improved sample complexities for deep networks and robust classification via an all-layer margin, 2019b.

Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. In *Advances in Neural Information Processing Systems*, pp. 4148–4158, 2017a.

Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. In *Advances in Neural Information Processing Systems*, pp. 4148–4158, 2017b.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.

A Experiments

A.1 More training details

During our experiments, we found that Batch Normalization (Ioffe & Szegedy, 2015) is crucial to reliably reach a low cross-entropy value for all models; since normalization is a indispensable components of modern neural networks, we decide to use batch normalization in all of our models. We *remove* batch normalization before computing any measure by fusing the γ, β and moving statistics with the convolution operator that precedes the normalization. This is important as Dinh et al. (2017) showed that common generalization measures such as sharpness can be easily manipulated with re-parameterization. We also discovered that the models trained with data augmentation often cannot fit the data (i.e. reach cross-entropy 0.01) completely. Since a model with data augmentation tends to consistently generalize better than the models without data augmentation, measure that reflects the training error (i.e. value of cross-entropy) will easily predict the ranking between two models even though it has only learned that one model uses data augmentation (see the thought experiments from the previous section). While certain hyperparameter configuration can reach cross-entropy of 0.01 even with data augmentation, it greatly limits the space of models that we can study. Hence, we make the design choice to not include data augmentation in the models of this study. Note that from a theoretical perspective, data augmentation is also challenging to analyze since the training samples generated from the procedure are no longer identical and independently distributed. All values for all the measures we computed over these models can be found in Table 5 in Appendix A.5.

A.2 The choice of stopping criterion

The choice of stopping criterion is very essential and could completely change the evaluation and the resulting conclusions. In our experiments we noticed that if we pick the stopping criterion based on number of iterations or number of epochs, then since some models optimize faster than others, they end up fitting the training data more and in that case the cross-entropy itself can be very predictive of generalization. To make it harder to distinguish models based on their training performance, it makes more sense to choose the stopping criterion based on the training error or training loss. We noticed that as expected, models with the same cross-entropy usually have very similar training error so that suggests that this choice is not very important. However, during the optimization the training error behavior is noisier than cross-entropy and moreover, after the training error reaches zero, it cannot distinguish models while the cross-entropy is still meaningful after fitting the data. Therefore, we decided to use cross-entropy as the stopping criterion.

A.3 All Model Specification

As mentioned in the main text, the models we use resemble Network-in-Network (Gao et al., 2011) which is a class of more parameter efficient convolution neural networks that achieve reasonably competitive performance on modern image classification benchmarks. The model consists blocks of modules that have 1 3×3 convolution with stride 2 followed by 2 1×1 convolution with stride 1. We refer to this single module as a `NiN-block` and construct models of different size by stacking NiN-block. For simplicity, all NiN-block have the same number of output channels c_{out}. Dropout is applied at the end of every NiN-block. At the end of the model, there is a 1×1 convolution reducing the channel number to the class number (i.e. 10 for CIFAR-10) followed by a global average pooling to produce the output logits.

For width, we choose from c_{out} from 3 options: $\{2 \times 96, 4 \times 96, 8 \times 96\}$.

For depth, we choose from 3 options: $\{2 \times \text{NiNblock}, 4 \times \text{NiNblock}, 8 \times \text{NiNblock}\}$

For dropout, we choose from 3 options: $\{0.0, 0.25, 0.5\}$

For batch size, we choose from: $\{32, 64, 128\}$

Since each optimizer may require different learning rate and in some cases, different regularization, we fine-tuned the hyper-parameters for each optimizer while keeping 3 options for every hyper-parameter choices[5].

Momentum SGD: We choose momentum of 0.9 and choose the initial learning rate η from $\{0.1, 0.032, 0.01\}$ and regularization coefficient λ from $\{0.0, 0.0001, 0.0005\}$. The learning rate decay schedule is $\times 0.1$ at iterations $[60000, 90000]$.

Adam: We choose initial learning rate η from $\{0.001, 3.2e-4, 1e-4\}$, $\epsilon = 1e-3$ and regularization coefficient λ from $\{0.0, 0.0001, 0.0005\}$. The learning rate decay schedule is $\times 0.1$ at iterations $[60000, 90000]$.

RMSProp: We choose initial learning rate η from $\{0.001, 3.2e{-}4, 1e{-}4\}$ and regularization coefficient λ from $\{0.0, 0.0001, 0.0003\}$. The learning rate decay schedule is $\times 0.1$ at iterations $[60000, 90000]$.

A.4 Canonical Measures

Based on empirical observations made by the community as a whole, the canonical ordering we give to each of the hyper-parameter categories are as follows:

1. Batchsize: smaller batchsize leads to smaller generalization gap
2. Depth: deeper network leads to smaller generalization gap
3. Width: wider network leads to smaller generalization gap
4. Dropout: The higher the dropout (≤ 0.5) the smaller the generalization gap
5. Weight decay: The higher the weight decay (smaller than the maximum for each optimizer) the smaller the generalization gap
6. Learning rate: The higher the learning rate (smaller than the maximum for each optimizer) the smaller the generalization gap
7. Optimizer: Generalization gap of `Momentum SGD` $<$ Generalization gap of `Adam` $<$ Generalization gap of `RMSProp`

A.5 All results

Below we present all of the measures we computed and their respective τ and Ψ on the 2000+ models we trained and additional plots.

[5]While methods with adaptive methods generally require less tuning, in practice researchers have observed performance gains from tuning the initial learning rate and learning rate decay.

	ref	batchsize	dropout	learning rate	depth	optimizer	weight decay	width	overall τ	Ψ
vc dim	19	0.000	0.000	0.000	−0.909	0.000	0.000	−0.171	−0.251	−0.154
# params	20	0.000	0.000	0.000	−0.909	0.000	0.000	−0.171	−0.175	−0.154
sharpness	51	0.537	−0.523	0.449	0.826	0.221	0.233	−0.004	0.282	0.248
pacbayes	48	0.372	−0.457	0.042	0.644	0.179	−0.179	−0.142	0.064	0.066
sharpness-orig	52	0.542	−0.359	0.716	0.816	0.297	0.591	0.185	0.400	0.398
pacbayes-orig	49	0.526	−0.076	0.705	0.546	0.341	0.564	−0.086	0.293	0.360
frob-distance	40	−0.317	−0.833	−0.718	0.526	−0.214	−0.669	−0.166	−0.263	−0.341
spectral-init	25	−0.330	−0.845	−0.721	−0.908	−0.208	−0.313	−0.231	−0.576	−0.508
spectral-orig	26	−0.262	−0.762	−0.665	−0.908	−0.131	−0.073	−0.240	−0.537	−0.434
spectral-orig-main	28	−0.262	−0.762	−0.665	−0.908	−0.131	−0.073	−0.240	−0.537	−0.434
fro/spec	33	0.563	0.351	0.744	−0.898	0.326	0.665	−0.053	−0.008	0.243
prod-of-spec	32	−0.464	−0.724	−0.722	−0.909	−0.197	−0.142	−0.218	−0.559	−0.482
prod-of-spec/margin	31	−0.308	−0.782	−0.702	−0.907	−0.166	−0.148	−0.179	−0.570	−0.456
sum-of-spec	35	−0.464	−0.724	−0.722	0.909	−0.197	−0.142	−0.218	0.102	−0.223
sum-of-spec/margin	34	−0.308	−0.782	−0.702	0.909	−0.166	−0.148	−0.179	0.064	−0.197
spec-dist	41	−0.458	−0.838	−0.568	0.738	−0.319	−0.182	−0.171	−0.110	−0.257
prod-of-fro	37	0.440	−0.199	0.538	−0.909	0.321	0.731	−0.101	−0.297	0.117
prod-of-fro/margin	36	0.513	−0.291	0.579	−0.907	0.364	0.739	−0.088	−0.295	0.130
sum-of-fro	39	0.440	−0.199	0.538	0.913	0.321	0.731	−0.101	0.418	0.378
sum-of-fro/margin	38	0.520	−0.369	0.598	0.882	0.380	0.738	−0.080	0.391	0.381
1/margin	22	−0.312	0.593	−0.234	−0.758	−0.223	0.211	−0.125	−0.124	−0.121
neg-entropy	23	0.346	−0.529	0.251	0.632	0.220	−0.157	0.104	0.148	0.124
path-norm	44	0.363	−0.190	0.216	0.925	0.272	0.195	0.178	0.370	0.280
path-norm/margin	43	0.363	0.017	0.148	0.922	0.230	0.280	0.173	0.374	0.305
param-norm	42	0.236	−0.516	0.174	0.330	0.187	0.124	−0.170	0.073	0.052
fisher-rao	45	0.396	0.147	0.240	−0.516	0.120	0.551	0.177	0.090	0.160
cross-entropy	21	0.440	−0.402	0.140	0.390	0.149	0.232	0.080	0.149	0.147
$1/\sigma$ pacbayes	53	0.501	−0.033	0.744	0.200	0.346	0.609	0.056	0.303	0.346
$1/\sigma$ sharpness	54	0.532	−0.326	0.711	0.776	0.296	0.592	0.263	0.399	0.406
num-step-0.1-to-0.01-loss	64	−0.151	−0.069	−0.014	0.114	0.072	−0.046	−0.021	−0.088	−0.016
num-step-to-0.1-loss	63	−0.664	−0.861	−0.255	0.440	−0.030	−0.628	0.043	−0.264	−0.279
$1/\alpha'$ sharpness mag	62	0.570	0.148	0.762	0.824	0.297	0.741	0.269	0.484	0.516
$1/\sigma'$ pacbayes mag	61	0.490	−0.215	0.505	0.896	0.186	0.147	0.195	0.365	0.315
pac-sharpness-mag-init	59	−0.293	−0.841	−0.698	−0.909	−0.240	−0.631	−0.171	−0.225	−0.541
pac-sharpness-mag-orig	60	0.401	−0.514	0.321	−0.909	0.181	0.281	−0.171	−0.158	−0.059
pacbayes-mag-init	56	0.425	−0.658	−0.035	0.874	0.099	−0.407	0.069	0.175	0.052
pacbayes-mag-orig	57	0.532	−0.480	0.508	0.902	0.188	0.155	0.186	0.410	0.284
grad-noise-final	66	0.452	0.119	0.427	0.141	0.245	0.432	0.230	0.311	0.292
grad-noise-epoch-1	65	0.071	0.378	0.376	−0.517	0.121	0.221	0.037	0.070	0.098
oracle 0.01		0.579	0.885	0.736	0.920	0.529	0.622	0.502	0.851	0.682
oracle 0.02		0.414	0.673	0.548	0.742	0.346	0.447	0.316	0.726	0.498
oracle 0.05		0.123	0.350	0.305	0.401	0.132	0.201	0.142	0.456	0.236
oracle 0.1		0.069	0.227	0.132	0.223	0.086	0.121	0.093	0.241	0.136
canonical ordering		−0.652	0.969	0.733	0.909	−0.055	0.735	0.171	0.005	0.402

Table 5: All complexity measures (rows), hyperparameters (columns) and the rank-correlation coefficients that we investigated.

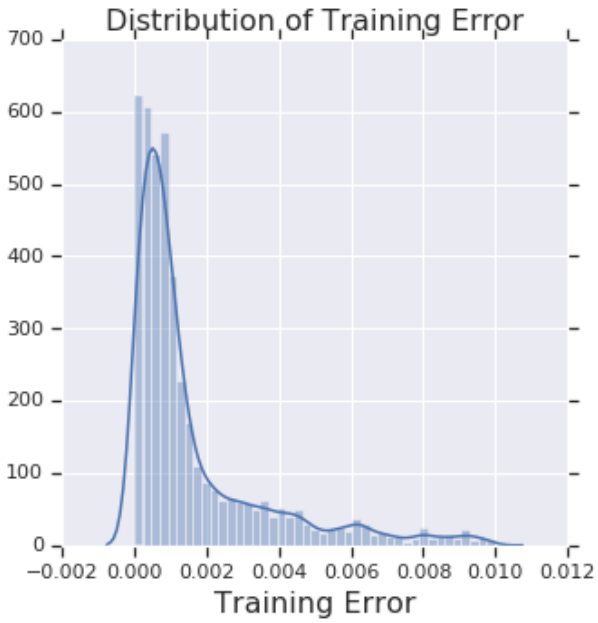

Figure 2: Distribution of training error on the trained models.

	batchsize	dropout	learning rate	depth	optimizer	weight decay	width	overall τ	Ψ
vc dim	0.0000	0.0000	0.0000	−1.0000	0.0000	0.0000	−0.0478	−0.3074	−0.1497
# params	0.0000	0.0000	0.0000	−1.0000	0.0000	0.0000	−0.0478	−0.1934	−0.1497
sharpness	0.1898	−0.4092	0.4569	0.9752	0.1708	0.2444	0.1202	0.5438	0.2497
pacbayes	0.0606	−0.5806	0.0503	0.9447	0.0831	−0.2123	0.0034	0.3688	0.0499
sharpness-orig	0.2324	−0.1807	0.6329	0.9595	0.2196	0.5018	0.1923	0.5175	0.3654
pacbayes-orig	0.1983	−0.2055	0.5979	0.8863	0.2286	0.4583	0.0655	0.3708	0.3185
frob-distance	−0.1071	−0.8603	−0.6270	0.8874	−0.1677	−0.6319	−0.0302	0.1765	−0.2196
spec-init	−0.2854	−0.7928	−0.6423	−0.9989	−0.1063	−0.2913	−0.0799	−0.6284	−0.4567
spec-orig	−0.1362	−0.6110	−0.4688	−0.9932	−0.0513	0.0671	−0.1096	−0.6163	−0.3290
spec-prig-main	−0.1362	−0.6110	−0.4688	−0.9628	−0.0513	0.0671	−0.2797	−0.5870	−0.3490
fro/spec	0.6285	0.3961	0.6646	−0.8274	0.2317	0.6047	0.0525	−0.1264	0.2501
prod-of-spec	−0.2603	−0.5835	−0.6095	−0.9628	−0.1063	−0.0343	−0.2705	−0.5615	−0.4039
prod-of-spec/margin	−0.2582	−0.6419	−0.5852	−0.9289	−0.0918	−0.0681	−0.2477	−0.5404	−0.4031
sum-of-spec	−0.2603	−0.5835	−0.6095	0.9989	−0.1063	−0.0343	−0.2705	0.4627	−0.1237
sum-of-spec/margin	−0.2582	−0.6419	−0.5852	0.9921	−0.0918	−0.0681	−0.2477	0.4421	−0.1287
spec-dist	−0.2734	−0.7752	−0.3386	0.9616	−0.0669	−0.2637	−0.0434	0.3542	−0.1142
prod-of-fro	0.5098	−0.0369	0.5439	−1.0000	0.1861	0.6508	0.0126	−0.4983	0.1238
prod-of-fro/margin	0.4673	−0.1262	0.5534	−1.0000	0.2079	0.6375	0.0091	−0.5001	0.1070
sum-of-fro	0.5098	−0.0369	0.5439	0.9853	0.1861	0.6508	0.0126	0.5928	0.4074
sum-of-fro/margin	0.4673	−0.1262	0.5534	0.9492	0.2079	0.6375	0.0091	0.5638	0.3855
1/margin	0.0684	0.3861	−0.1519	−0.9314	−0.1018	0.3211	0.0216	−0.3829	−0.0554
neg-entropy	−0.0320	−0.4506	0.1481	0.7999	0.1360	−0.2460	−0.0106	0.3001	0.0492
path-norm	0.2150	0.2565	0.0464	0.9854	0.1018	0.3885	0.0614	0.5626	0.2936
path norm/margin	0.1107	0.0291	0.1340	0.9978	0.1504	0.2098	0.0683	0.5798	0.2429
parameter-norm	0.3246	−0.4794	0.1730	0.6639	0.0780	0.1383	−0.0398	0.3747	0.1227
fisher-rao	0.2313	0.0500	0.0222	−0.6189	0.1008	0.3190	0.0546	−0.2844	0.0227
cross-entropy	0.0923	−0.4091	−0.0042	−0.0096	0.0811	0.1118	−0.0432	−0.0693	−0.0258
$1/\sigma$ pacbayes	0.1867	−0.1862	0.6164	0.6665	0.2280	0.4985	0.1512	0.3148	0.3087
$1/\sigma$ sharpness	0.2321	−0.1549	0.6330	0.9363	0.2253	0.5163	0.2179	0.4930	0.3723
num-step-0.1-to-0.01-loss	−0.1224	−0.1610	−0.0061	0.1556	0.0737	−0.0415	−0.0154	−0.0720	−0.0167
num-step-to-0.1-loss	−0.6656	−0.9120	−0.3613	0.9556	−0.1450	−0.5974	−0.0414	0.0944	−0.2524
$1/\alpha'$ sharpness mag	0.4546	0.3254	0.6650	0.9831	0.2753	0.6495	0.2680	0.5676	0.5173
$1/\alpha'$ pacbayes mag	0.2525	0.1250	0.4758	0.9805	0.1629	0.2698	0.0871	0.5674	0.3362
pac-sharpness-mag-init	−0.0787	−0.7181	−0.4883	−1.0000	−0.0640	−0.4720	−0.0502	−0.2254	−0.4102
pac-sharpness-mag-orig	0.5005	−0.3831	0.3153	−1.0000	0.1648	0.2440	−0.0502	−0.1687	−0.0298
pacbayes-mag-init	0.2289	−0.3322	0.2298	−0.9786	0.1625	−0.0429	−0.0484	−0.1309	−0.1116
pacbayes-mag-orig	0.0984	−0.6821	0.2351	−0.9842	0.1304	0.0542	−0.0484	−0.1682	−0.1709
grad-noise-final	0.2128	−0.1862	0.2458	0.0343	0.1711	0.3211	0.1149	0.0594	0.1305
grad-noise-epoch-1	0.1590	0.1912	−0.0159	0.0118	−0.0534	0.2760	−0.0046	0.1222	0.0806
oracle 0.01	0.3811	0.6463	0.4293	0.9517	0.3478	0.3946	0.3572	0.8070	0.5012
oracle 0.02	0.2410	0.4102	0.2964	0.8730	0.1886	0.2190	0.1741	0.6854	0.3432
oracle 0.05	0.1238	0.2235	0.1530	0.6706	0.0522	0.1057	0.0785	0.5162	0.2010
oracle 0.1	−0.0239	0.0708	0.0844	0.4356	0.0408	0.0526	0.0512	0.3322	0.1017
canonical ordering	−0.6732	0.9539	0.6424	1.0000	−0.1028	0.6662	0.0478	0.0123	0.3620

Table 6: **SVHN** results: All complexity measures (rows), hyperparameters (columns) and the rank-correlation coefficients that we investigated.

	batchsize	dropout	learning rate	depth	optimizer	weight decay	width	overall τ	Ψ
vc dim	0.0000	0.0000	0.0000	−0.7520	0.0000	0.0000	−0.0392	−0.1770	−0.1130
# params	0.0000	0.0000	0.0000	−0.7520	0.0000	0.0000	−0.0392	−0.1194	−0.1130
sharpness	0.2059	−0.1966	0.1336	0.6358	−0.0532	−0.0127	−0.0317	0.2325	0.0973
pacbayes	0.1480	−0.0488	−0.0611	0.5493	−0.0570	−0.2340	−0.0563	0.1477	0.0343
sharpness-orig	0.2271	0.0167	0.4462	0.6262	0.0600	0.1563	−0.0058	0.2995	0.2181
pacbayes-orig	0.2587	0.1655	0.5282	0.5238	0.1102	0.1318	−0.0174	0.3104	0.2430
frob-distance	−0.1814	−0.7677	−0.6504	0.3767	−0.2403	−0.3831	−0.0392	−0.2652	−0.2693
spec-init	−0.1495	−0.5752	−0.6208	−0.7407	−0.2650	−0.2885	−0.0945	−0.4333	−0.3906
spec-orig	−0.0837	−0.4196	−0.4747	−0.7379	−0.1776	−0.1468	−0.1085	−0.3860	−0.3070
spec-orig-main	−0.0837	−0.4196	−0.4747	−0.7284	−0.1776	−0.1468	−0.1857	−0.3940	−0.3166
fro/spec	0.2606	0.3893	0.7221	−0.7435	0.4169	0.4404	0.0615	0.0477	0.2210
prod-of-spec	−0.2034	−0.5619	−0.6199	−0.7520	−0.2184	−0.1269	−0.0691	−0.4176	−0.3645
prod-of-spec/margin	−0.1257	−0.4727	−0.5549	−0.7181	−0.2260	−0.2113	−0.1707	−0.4238	−0.3542
sum-of-spec	−0.2034	−0.5619	−0.6199	0.7520	−0.2184	−0.1269	−0.0691	0.0547	−0.1496
sum-of-spec/margin	−0.1257	−0.4727	−0.5549	0.7501	−0.2260	−0.2113	−0.1707	0.0868	−0.1445
spec-dist	−0.2005	−0.8378	−0.5692	0.5832	−0.3751	−0.0899	−0.0392	−0.1517	−0.2184
prod-of-fro	0.2854	−0.1532	0.4967	−0.7520	0.3609	0.4656	0.0054	−0.2162	0.1013
prod-of-fro/margin	0.2816	−0.1987	0.4613	−0.7520	0.2365	0.3729	0.0130	−0.2113	0.0592
sum-of-fro	0.2854	−0.1532	0.4967	0.7652	0.3609	0.4656	0.0054	0.3407	0.3180
sum-of-fro/margin	0.2816	−0.1987	0.4613	0.7643	0.2365	0.3729	0.0130	0.3356	0.2758
1/margin	−0.1652	0.3153	−0.0850	−0.5474	0.1263	0.1738	0.2142	−0.1295	0.0046
neg-entropy	0.2200	−0.3496	0.3906	0.5584	0.1614	−0.2819	−0.2095	0.1378	0.0699
path-norm	0.2549	0.5258	0.2951	0.8161	0.2593	0.2223	0.0420	0.3892	0.3451
path-norm/margin	0.2510	0.0441	0.3314	0.7718	0.1206	0.0571	−0.0558	0.3580	0.2172
param-norm	0.2472	−0.0090	0.3754	0.1287	0.2716	0.1569	−0.0458	0.0865	0.1607
fisher-rao	0.0727	0.3722	0.0162	−0.5314	−0.1595	0.0355	0.0231	0.0246	−0.0245
cross-entropy	0.1452	−0.1095	−0.0630	0.3484	−0.2080	−0.1140	−0.2236	0.1410	−0.0321
$1/\sigma$ pacbayes	0.2525	0.1905	0.4993	0.3698	0.0822	0.1298	0.0660	0.3213	0.2272
$1/\sigma$ sharpness	0.2120	0.0008	0.3879	0.6097	0.0161	0.1191	−0.0073	0.3005	0.1912
num-step-0.1-to-0.01-loss	−0.0053	−0.0747	−0.0792	0.1688	0.0318	0.0621	−0.0168	−0.0210	0.0124
num-step-to-0.1-loss	−0.3219	−0.8336	−0.2626	0.2859	−0.0699	−0.4231	−0.0062	−0.2350	−0.2331
$1/\alpha'$ sharpness mag	0.2127	0.2602	0.4458	0.6430	0.0354	0.1846	0.0071	0.3613	0.2555
$1/\alpha'$ pacbayes mag	0.1660	0.0065	0.4001	0.6820	0.0319	−0.0879	−0.1308	0.2878	0.1525
pac-sharpness-mag-init	−0.1776	−0.7743	−0.6476	−0.7520	−0.2498	−0.3803	−0.0392	−0.1602	−0.4315
pac-sharpness-mag-orig	0.3789	−0.0109	0.5033	−0.7520	0.3067	0.2688	−0.0392	−0.0867	0.0937
pacbayes-mag-init	0.1404	−0.2537	0.1203	−0.7501	0.0446	−0.2183	−0.0392	−0.1123	−0.1366
pacbayes-mag-orig	0.1309	−0.4026	0.2961	−0.7520	0.0389	−0.1434	−0.0392	−0.1075	−0.1245
grad-noise-final	0.1013	0.3514	0.3706	0.2730	0.1035	−0.0652	0.0250	0.3538	0.1656
grad-noise-epoch-1	0.0801	0.4045	0.3792	−0.3701	0.1349	0.1328	0.0814	0.1279	0.1204
oracle 0.01	0.5789	0.8862	0.7507	0.8274	0.5878	0.5464	0.5123	0.8470	0.6700
oracle 0.02	0.3588	0.7288	0.5922	0.5804	0.3970	0.3440	0.3927	0.7032	0.4848
oracle 0.05	0.1114	0.4149	0.3066	0.2937	0.1918	0.1473	0.1697	0.4267	0.2336
oracle 0.1	0.1037	0.2281	0.1738	0.1957	0.1225	0.0692	0.0876	0.2423	0.1401
canonical ordering	−0.3254	0.9459	0.7125	0.7520	−0.0598	0.4628	0.0392	−0.0151	0.3610

Table 7: **CIFAR-10 Loss=0.1** results: All complexity measures (rows), hyperparameters (columns) and the rank-correlation coefficients that we investigated.

	batchsize	dropout	learning rate	depth	optimizer	weight decay	width	overall τ	Ψ
vc dim	0	0	0	−0.9073	0	0	−0.1487	−0.2509	−0.1509
# params	0	0	0	−0.9073	0	0	−0.1487	−0.1751	−0.1509
sharpness	0.5492	−0.5155	0.4636	0.8247	0.2134	0.2025	0.0083	0.2848	0.2495
pacbayes	0.3896	−0.4459	0.0427	0.6289	0.1721	−0.1757	−0.1266	0.0647	0.0693
sharpness-orig	0.5493	−0.3492	0.7147	0.8101	0.3006	0.5655	0.1976	0.3996	0.3984
pacbayes-orig	0.5399	−0.0847	0.7237	0.5377	0.3561	0.5597	−0.0693	0.2895	0.3662
frob-distance	−0.3048	−0.8366	−0.7253	0.5301	−0.2437	−0.6701	−0.1499	−0.2606	−0.3429
spec-init	−0.3414	−0.8436	−0.7326	−0.9068	−0.2422	−0.3134	−0.2133	−0.5743	−0.5133
spec-orig	−0.2633	−0.7593	−0.678	−0.9068	−0.1611	−0.0683	−0.2273	−0.5354	−0.4377
spec-orig-main	−0.2633	−0.7593	−0.678	−0.9064	−0.1611	−0.0683	−0.2662	−0.5451	−0.4432
fro / spec	0.5884	0.3703	0.7501	−0.9014	0.3661	0.658	−0.0219	−0.0086	0.2585
prod-of-spec	−0.4718	−0.7237	−0.7302	−0.9072	−0.2385	−0.1409	−0.2126	−0.5598	−0.4893
prod-of-spec/margin	−0.3222	−0.7803	−0.716	−0.9066	−0.2066	−0.1614	−0.1727	−0.5698	−0.4665
sum-of-spec	−0.4718	−0.7237	−0.7302	0.9072	−0.2385	−0.1409	−0.2126	0.1023	−0.2301
sum-of-spec/margin	−0.3222	−0.7803	−0.716	0.9066	−0.2066	−0.1614	−0.1727	0.0662	−0.2075
spec-dist	−0.4506	−0.8263	−0.5791	0.7297	−0.3413	−0.2027	−0.1485	−0.1044	−0.2598
prod-of-fro	0.4659	−0.1885	0.5283	−0.9072	0.3342	0.7255	−0.0835	−0.2972	0.1250
prod-of-fro/margin	0.5377	−0.372	0.5888	−0.9072	0.4024	0.7329	−0.0673	−0.2957	0.1308
sum-of-fro	0.4659	−0.1885	0.5283	0.9099	0.3342	0.7255	−0.0835	0.4157	0.3845
sum-of-fro/margin	0.5377	−0.372	0.5888	0.8832	0.4024	0.7329	−0.0673	0.3894	0.3865
1/margin	−0.3334	0.5914	−0.2543	−0.7539	−0.2257	0.2097	−0.0988	−0.1257	−0.1236
input grad norm	0.5235	0.263	0.0544	0.6239	0.0888	0.5969	0.2054	0.3836	0.3366
neg-entropy	0.3686	−0.5443	0.2609	0.6326	0.2296	−0.1567	0.0973	0.1472	0.1269
path-norm	0.2457	0.262	0.0397	0.9296	0.1271	0.3291	0.1558	0.3718	0.2984
path norm/margin	0.3692	−0.2022	0.2159	0.9189	0.2523	0.2103	0.1582	0.3724	0.2747
param-norm	0.2414	−0.5194	0.1611	0.3346	0.1866	0.1198	−0.1509	0.0729	0.0533
fisher-rao	0.4327	0.1625	0.2494	−0.5317	0.1322	0.5559	0.1484	0.1028	0.1642
one epoch loss	0.3939	−0.4362	0.0477	0.1573	0.1149	−0.0475	0.0128	−0.0147	0.0347
cross-entropy	0.4443	−0.4015	0.1518	0.3821	0.1367	0.2322	0.0676	0.1515	0.1447
$1/\sigma$ pacbayes	0.5109	−0.0349	0.7551	0.2032	0.3738	0.6048	0.0686	0.2993	0.3545
$1/\sigma$ sharpness	0.536	−0.3169	0.7154	0.7529	0.3021	0.5726	0.2615	0.3976	0.4034
min(norm distance)	0.2414	−0.5194	0.1611	0.3346	0.1866	0.1198	−0.1509	0.0729	0.0533
num-step-0.1-to-0.01-loss	−0.1458	−0.0816	−0.0166	0.1318	0.0949	−0.0348	−0.0387	−0.086	−0.0130
step to	−0.6798	−0.5418	−0.4441	0.3493	−0.0578	−0.6909	0.0102	−0.2812	−0.2936
num-step-to-0.1-loss	−0.68	−0.8526	−0.2662	0.4545	−0.0291	−0.6484	0.0291	−0.2626	−0.2847
$1/\alpha'$ sharpness mag	0.5802	0.1381	0.7537	0.8181	0.3163	0.7371	0.2416	0.481	0.5122
$1/\alpha'$ pacbayes mag	0.5089	−0.2388	0.5203	0.8959	0.1907	0.1628	0.1738	0.3649	0.3162
pac-sharpness-mag-init	−0.2967	−0.8451	−0.7165	−0.9072	−0.2637	−0.6387	−0.1488	−0.2256	−0.5452
pac-sharpness-mag-orig	0.4145	−0.5227	0.3102	−0.9072	0.1916	0.2586	−0.1488	−0.159	−0.0577
pacbayes-mag-init	0.4783	−0.6438	0.2402	−0.9072	0.1446	−0.1006	−0.1488	−0.1669	−0.1339
pacbayes-mag-orig	0.4694	−0.7749	0.317	−0.9072	0.1343	0.0315	−0.1488	−0.1682	−0.1255
grad-noise-final	0.3663	0.0039	0.3066	0.0813	0.1773	0.4492	0.1615	0.2521	0.2209
grad-noise-epoch-1	−0.0376	0.3618	0.2691	−0.5688	−0.0342	0.2535	−0.0616	−0.0252	0.0260
oracle 0.01	0.588	0.8718	0.7047	0.9094	0.5191	0.6117	0.5107	0.852	0.6736
oracle 0.02	0.3904	0.6862	0.5405	0.7226	0.35	0.3969	0.336	0.7197	0.4889
oracle 0.05	0.1827	0.3694	0.3099	0.3893	0.1478	0.1676	0.1665	0.4518	0.2476
oracle 0.1	0.106	0.2132	0.1694	0.2084	0.0922	0.0859	0.082	0.259	0.1367
canonical ordering	−0.668	0.9753	0.7421	0.9073	−0.0511	0.7268	0.1487	−0.0039	0.3973

Table 8: **Average CIFAR-10 Loss=0.01**: All complexity measures (rows), hyperparameters (columns) and the average rank-correlation coefficients over 5 runs. The numerical values are consistent of that of Table 5.

	batchsize	dropout	learning rate	depth	optimizer	weight decay	width	overall τ	Ψ
vc dim	0	0	0	0.0038	0	0	0.0179	0.0006	0.0026
# params	0	0	0	0.0038	0	0	0.0179	0.0009	0.0026
sharpness	0.0124	0.0129	0.0153	0.0036	0.0196	0.0154	0.0181	0.0026	0.0056
pacbayes	0.0171	0.0159	0.0108	0.0086	0.0074	0.0078	0.0169	0.0008	0.0048
sharpness-orig	0.0082	0.0106	0.0062	0.0073	0.0192	0.0151	0.0164	0.0034	0.0048
pacbayes-orig	0.011	0.0062	0.0111	0.0083	0.0162	0.013	0.0173	0.0025	0.0047
frob-distance	0.0102	0.0049	0.0067	0.0058	0.017	0.0102	0.0176	0.0035	0.0043
spec-init	0.0061	0.0029	0.0072	0.004	0.0192	0.0191	0.0127	0.001	0.0045
spec-orig	0.0015	0.0096	0.0072	0.004	0.0166	0.0234	0.0136	0.0009	0.0049
spec-orig-main	0.0015	0.0096	0.0072	0.0037	0.0166	0.0234	0.0083	0.0004	0.0046
fro / spec	0.0164	0.0105	0.0034	0.0048	0.0205	0.0151	0.0203	0.0024	0.0055
prod-of-spec	0.0053	0.0109	0.0048	0.0037	0.0237	0.0249	0.0101	0.0008	0.0055
prod-of-spec/margin	0.0075	0.0078	0.0082	0.0039	0.0225	0.0232	0.0054	0.0006	0.0051
sum-of-spec	0.0053	0.0109	0.0048	0.0037	0.0237	0.0249	0.0101	0.0014	0.0055
sum-of-spec/margin	0.0075	0.0078	0.0082	0.0035	0.0225	0.0232	0.0054	0.0015	0.0051
spec-dist	0.012	0.0095	0.0081	0.0084	0.0221	0.0122	0.0177	0.0036	0.0052
prod-of-fro	0.016	0.0096	0.0117	0.0037	0.0191	0.0121	0.0174	0.0014	0.0052
prod-of-fro/margin	0.0112	0.0126	0.0083	0.0037	0.0224	0.0093	0.0141	0.0014	0.0049
sum-of-fro	0.016	0.0096	0.0117	0.0034	0.0191	0.0121	0.0174	0.0024	0.0052
sum-of-fro/margin	0.0112	0.0126	0.0083	0.0054	0.0224	0.0093	0.0141	0.002	0.0049
1/margin	0.0191	0.0059	0.0154	0.0068	0.0221	0.0079	0.0224	0.0026	0.0060
input grad norm	0.0147	0.0186	0.019	0.0018	0.0222	0.0161	0.011	0.0043	0.0061
neg-entropy	0.0163	0.0169	0.012	0.0093	0.022	0.0184	0.0204	0.0025	0.0064
path-norm	0.0103	0.006	0.0079	0.0034	0.0174	0.0115	0.0178	0.0014	0.0044
path norm/margin	0.0095	0.0172	0.0054	0.0056	0.0157	0.0224	0.0192	0.0019	0.0056
param-norm	0.0125	0.0061	0.0071	0.0077	0.0083	0.0051	0.0175	0.0016	0.0038
fisher-rao	0.0192	0.0153	0.0084	0.0083	0.0311	0.01	0.0158	0.0069	0.0065
one epoch loss	0.0169	0.0128	0.0146	0.0066	0.0223	0.0126	0.0173	0.005	0.0058
cross-entropy	0.0221	0.0128	0.0174	0.0138	0.0151	0.014	0.0183	0.0023	0.0062
$1/\sigma$ pacbayes	0.0095	0.0031	0.0081	0.0066	0.0173	0.0132	0.0162	0.0035	0.0044
$1/\sigma$ sharpness	0.0084	0.009	0.0077	0.0126	0.0185	0.0119	0.0121	0.0039	0.0045
min(norm distance)	0.0125	0.0061	0.0071	0.0077	0.0083	0.0051	0.0175	0.0016	0.0038
num-step-0.1-to-0.01-loss	0.0049	0.0094	0.0071	0.0182	0.0147	0.0081	0.0222	0.0023	0.0051
step to	0.0118	0.011	0.0162	0.0169	0.0135	0.0101	0.012	0.002	0.0050
num-step-to-0.1-loss	0.0119	0.0059	0.0101	0.0236	0.0191	0.0148	0.0152	0.002	0.0058
$1/\alpha'$ sharpness mag	0.0108	0.0224	0.0048	0.0082	0.0262	0.0097	0.0201	0.0031	0.0062
$1/\alpha'$ pacbayes mag	0.0198	0.0166	0.0084	0.0037	0.0228	0.015	0.0237	0.0044	0.0065
pac-sharpness-mag-init	0.0113	0.0039	0.0139	0.0037	0.0186	0.0155	0.0179	0.0011	0.0051
pac-sharpness-mag-orig	0.016	0.0061	0.0127	0.0037	0.0188	0.0139	0.0179	0.0008	0.0052
pacbayes-mag-init	0.022	0.0059	0.0171	0.0037	0.0173	0.0131	0.0179	0.001	0.0057
pacbayes-mag-orig	0.0221	0.0077	0.0083	0.0037	0.0213	0.0134	0.0179	0.0009	0.0057
grad-noise-final	0.0447	0.0598	0.0628	0.0337	0.0394	0.0243	0.0363	0.0309	0.0170
grad-noise-epoch-1	0.0547	0.0165	0.0542	0.0316	0.082	0.0173	0.0514	0.0478	0.0186
oracle 0.01	0.0178	0.0078	0.0153	0.0108	0.0189	0.0086	0.026	0.0026	0.0061
oracle 0.02	0.0133	0.0135	0.0081	0.0138	0.0272	0.0167	0.0058	0.0033	0.0058
oracle 0.05	0.0091	0.0249	0.0133	0.0136	0.0171	0.015	0.0239	0.0076	0.0066
oracle 0.1	0.0188	0.0333	0.0292	0.0341	0.0145	0.0185	0.0321	0.0107	0.0102
canonical ordering	0.0111	0.004	0.0073	0.0038	0.0185	0.0108	0.0179	0.0027	0.0045

Table 9: **Standard deviation CIFAR-10 Loss=0.01** results: All complexity measures (rows), hyperparameters (columns) and the standard deviation of each entry measured over 5 runs. The standard deviation for Ψ is computed assuming that each hyperparamters are independent from each other. We see that all standard deviation are quite small, suggesting the results in of Table 5 are statistically significant.

B Causal Analysis

Relying on correlation is intuitive but perhaps unsatisfactory. In our experiments, we change several hyper-parameters and assess the correlation between a complexity measure and generalization. When we observe correlation between a complexity measure and generalization, we want to differentiate the following two scenarios:

- Changing a hyper-parameter causes the complexity measure to be low and lower value of the measure causes the generalization gap to be low.
- Changing a hyper-parameter causes the complexity measure to be low and changing the same hyper-parameter also causes the generalization to be low but the lower value of the complexity measure by itself has no effect on generalization.

The above two scenarios are demonstrated in Figure 3. In attempt to truly understand these relationships, we will rely on the tools from probabilistic causality. Our approach is inspired by the seminal work on Inductive Causation (IC) Algorithm by Verma & Pearl (1991), which provides a framework for learning a graphical model through conditional independence test. While the IC algorithm traditionally initiates the graph to be fully connected, we will take advantage of our knowledge about generalization and prune edges of the initialized graph to expedite the computations. Namely, we assume that the choice of hyperparameter does not directly explain generalization, but rather it induces changes in some measure μ which can be used to explain generalization.

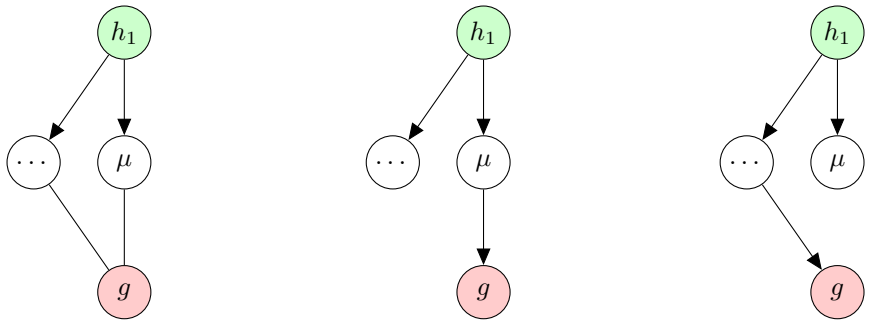

Figure 3: **Right**: Graph at initialization. **Middle**: The ideal graph where the measure μ can directly explain observed generalization. **Right**: Graph for correlation where μ cannot explain observed generalization.

Our primary interest is to establish the existence of an edge between μ and g. Suppose there exists a large family of complexity measures and among them there is a true complexity measure μ^* that can fully explain generalization. Then to verify the existence of the edge between μ and g, we can perform the ***conditional independent test*** by reading the conditional mutual information between μ and g given that a set of hyperparameter $\mathcal{S}$ is observed (e.g. $\mathcal{S}$ contains a single hyperparameter such as the learning rate). For any function $\phi : \Theta \to \mathbb{R}$, let $V_\phi(\theta_1, \theta_2) = \text{sign}(\phi(\theta_1) - \phi(\theta_2))$. Furthermore, let $U_\mathcal{S}$ be a random variable that correspond to the values of hyperparameters in $\mathcal{S}$. We calculate the conditional mutual information as follows:

$$\mathcal{I}(V_\mu, V_g \,|\, U_\mathcal{S}) = \sum_{U_\mathcal{S}} p(U_\mathcal{S}) \sum_{V_\mu \in \{\pm 1\}} \sum_{V_g \in \{\pm 1\}} p(V_\mu, V_g \,|\, U_\mathcal{S}) \log \left(\frac{p(V_\mu, V_g \,|\, U_\mathcal{S})}{p(V_\mu \,|\, U_\mathcal{S}) p(V_g \,|\, U_\mathcal{S})} \right) \tag{6}$$

Since in our case the conditional mutual information between a complexity measure and generalization is at most equal to the conditional entropy of generalization, we normalize it with the conditional entropy to arrive at a criterion ranging between 0 and 1:

$$\mathcal{H}(V_g \,|\, U_\mathcal{S}) = -\sum_{U_\mathcal{S}} p(U_\mathcal{S}) \sum_{V_g \in \{\pm 1\}} p(V_g \,|\, U_\mathcal{S}) \log(p(V_g \,|\, U_\mathcal{S})) \tag{7}$$

$$\hat{\mathcal{I}}(V_\mu, V_g \,|\, U_\mathcal{S}) = \frac{\mathcal{I}(V_\mu, V_g \,|\, U_\mathcal{S})}{\mathcal{H}(V_g \,|\, U_\mathcal{S})} \tag{8}$$

According to the IC algorithm, an edge is kept between two nodes if there exists no subset $\mathcal{S}$ of hyperparameters such that the two nodes are independent, i.e. $\hat{\mathcal{I}}(V_\mu, V_g \,|\, U_\mathcal{S}) = 0$. In our setting, letting $\mathcal{S}$ be the set of all hyperparameters is not possible as both the conditional entropy and conditional mutual information would become zero. Moreover, due to computational reasons, we only look at $|\mathcal{S}| \leq 2$:

$$\mathcal{K}(\mu) = \min_{U_\mathcal{S} \text{ s.t } |\mathcal{S}| \leq 2} \hat{\mathcal{I}}(V_\mu, V_g \,|\, U_\mathcal{S}) \tag{9}$$

On a high level, the larger $\mathcal{K}$ is for a measure μ, the more likely an edge exists between μ and g, and therefore the more likely μ can explain generalization. For details on the set-up, please refer to Appendix B.1 on how these quantities are estimated.

B.1 Definition of Random Variables

Since the measures are results of complicated interactions between the data, the model, and the training procedures, we cannot manipulate it to be any values that we want. Instead, we use the following definition of random variables: suppose $\mathcal{S}$ is a subset of all the components of $\boldsymbol{\theta}$ (e.g. $\mathcal{S} = \{\varnothing\}$ for $|\mathcal{S}| = 0$, $|\mathcal{S}| = \{\text{learning rate}\}$ for $|\mathcal{S}| = 1$ or $|\mathcal{S}| = \{\text{learning rate, dropout}\}$ for $|\mathcal{S}| = 2$). Specifically we denote $\mathcal{S}_{ab}$ as the collective condition $\{\theta_1^{(a)} = v_1, \theta_1^{(b)} = v_2, \ldots, \theta_{|\mathcal{S}|}^{(a)} = v_{2|\mathcal{S}|-1}, \theta_{|\mathcal{S}|}^{(b)} = v_{2|\mathcal{S}|}\}$. We can then define and empirical measure four probability $\Pr(\mu^{(a)} > \mu^{(b)}, g^{(a)} > g^{(b)} \,|\, \mathcal{S}_{ab})$, $\Pr(\mu^{(a)} > \mu^{(b)}, g^{(a)} < g^{(b)} \,|\, \mathcal{S}_{ab})$, $\Pr(\mu^{(a)} < \mu^{(b)}, g^{(a)} > g^{(b)} \,|\, \mathcal{S}_{ab})$ and $\Pr(\mu^{(a)} < \mu^{(b)}, g^{(a)} < g^{(b)} \,|\, \mathcal{S}_{ab})$.

	$\mu^{(a)} > \mu^{(b)}$	$\mu^{(a)} \leq \mu^{(b)}$
$g^{(a)} > g^{(b)}$	p_{00}	p_{01}
$g^{(a)} \leq g^{(b)}$	p_{10}	p_{11}

Figure 4: Joint Probability table for a single $\mathcal{S}_{ab}$

Together forms a 2 by 2 table that defines the joint distribution of the Bernoulli random variables $\Pr(g^{(a)} > g^{(b)} \,|\, \mathcal{S}_{ab})$ and $\Pr(\mu^{(a)} > \mu^{(b)} \,|\, \mathcal{S}_{ab})$. For notation convenience, we use $\Pr(\mu, g \,|\, \mathcal{S}_{ab})$, $\Pr(g \,|\, \mathcal{S}_{ab})$ and $\Pr(\mu \,|\, \mathcal{S}_{ab})$ to denote the joint and marginal. If there are $N = 3$ choices for each hyperparameter in $\mathcal{S}$ then there will be $N^{|\mathcal{S}|}$ such tables for each hyperparameter combination. Since each configuration occurs with equal probability, for that arbitrary $\boldsymbol{\theta}^{(a)}$ and $\boldsymbol{\theta}^{(b)}$ drawn from Θ conditioned on that the components of $\mathcal{S}$ are observed for both models, the joint distribution can be defined as $\Pr(\mu, g \,|\, \mathcal{S}) = \frac{1}{N^{|\mathcal{S}|}} \sum_{\mathcal{S}_{ab}} \Pr(\mu, g \,|\, \mathcal{S}_{ab})$ and likely the marginals can be defined as $\Pr(\mu \,|\, \mathcal{S}) = \frac{1}{N^{|\mathcal{S}|}} \sum_{\mathcal{S}_{ab}} \Pr(\mu \,|\, \mathcal{S}_{ab})$ and $\Pr(g \,|\, \mathcal{S}) = \frac{1}{N^{|\mathcal{S}|}} \sum_{\mathcal{S}_{ab}} \Pr(g \,|\, \mathcal{S}_{ab})$. With these notations established, all the relevant quantities can be computed by iterating over all pairs of models.

B.2 Results

Table 10 contains the results from conditional independence test which is largely consistent with that of ranking correlation.

	batchsize	dropout	learning rate	depth	optimizer	weight decay	width	$\lvert\mathcal{S}\rvert = 0$	$\lvert\mathcal{S}\rvert = 1$	$\lvert\mathcal{S}\rvert = 2$
#-param	0.0202	0.0278	0.0259	0.0044	0.0208	0.0216	0.0379	0.0200	0.0036	0.0000
-entropy	0.0120	0.0656	0.0113	0.0086	0.0120	0.0155	0.0125	0.0117	0.0072	0.0065
1-over-sigma-pacbayes-mag	0.0884	0.1514	0.0813	0.0399	0.1004	0.1025	0.0986	0.0960	0.0331	0.0241
1-over-sigma-pacbayes	0.0661	0.1078	0.0487	0.0809	0.0711	0.0589	0.0858	0.0664	0.0454	0.0340
1-over-sigma-sharpness-mag	0.1640	0.2572	0.1228	0.1424	0.1779	0.1562	0.1786	0.1741	0.1145	0.0544
1-over-sigma-sharpness	0.1086	0.2223	0.0792	0.0713	0.1196	0.1041	0.1171	0.1159	0.0592	0.0256
cross-entropy	0.0233	0.0850	0.0118	0.0075	0.0159	0.0119	0.0183	0.0161	0.0062	0.0040
displacement	0.0462	0.0530	0.0196	0.1559	0.0502	0.0379	0.0506	0.0504	0.0183	0.0128
fisher-rao	0.0061	0.0072	0.0020	0.0713	0.0057	0.0014	0.0071	0.0059	0.0013	0.0018
fro-over-spec	0.0019	0.0065	0.0298	0.0777	0.0036	0.0015	0.0005	0.0000	0.0005	0.0013
frob-distance	0.0462	0.0530	0.0196	0.1559	0.0502	0.0379	0.0506	0.0504	0.0183	0.0128
grad-noise-epoch-1	0.0051	0.0016	0.0028	0.0633	0.0113	0.0027	0.0052	0.0036	0.0013	0.0013
grad-noise-final	0.0623	0.0969	0.0473	0.0934	0.0745	0.0577	0.0763	0.0712	0.0441	0.0329
input-grad-norm	0.0914	0.1374	0.1203	0.0749	0.1084	0.0853	0.1057	0.1042	0.0623	0.0426
margin	0.0105	0.0750	0.0078	0.0133	0.0108	0.0183	0.0119	0.0108	0.0072	0.0051
oracle-0.01	0.6133	0.5671	0.6007	0.5690	0.6171	0.6108	0.6191	0.6186	0.4727	0.2879
oracle-0.02	0.4077	0.3557	0.3929	0.3612	0.4124	0.4057	0.4154	0.4130	0.2987	0.1637
oracle-0.05	0.1475	0.1167	0.1369	0.1241	0.1515	0.1469	0.1535	0.1515	0.0980	0.0503
pacbayes-mag-init	0.0216	0.0238	0.0274	0.0046	0.0222	0.0210	0.0345	0.0202	0.0038	0.0004
pacbayes-mag-orig	0.1160	0.2249	0.1006	0.0426	0.1305	0.1316	0.1246	0.1252	0.0354	0.0221
pacbayes-orig	0.0620	0.1071	0.0392	0.0597	0.0645	0.0550	0.0977	0.0629	0.0365	0.0225
pacbayes	0.0053	0.0164	0.0084	0.0086	0.0036	0.0066	0.0185	0.0030	0.0036	0.0040
parameter-norm	0.0039	0.0197	0.0066	0.0115	0.0064	0.0049	0.0167	0.0039	0.0038	0.0047
path-norm-over-margin	0.0943	0.1493	0.1173	0.0217	0.1025	0.1054	0.1090	0.1011	0.0181	0.0139
path-norm	0.1027	0.1230	0.1308	0.0315	0.1056	0.1028	0.1160	0.1030	0.0261	0.0240
prod-of-spec-over-margin	0.2466	0.3139	0.2179	0.1145	0.2473	0.2540	0.2497	0.2481	0.0951	0.0483
prod-of-spec	0.2334	0.3198	0.2070	0.1037	0.2376	0.2470	0.2394	0.2385	0.0862	0.0415
random	0.0005	0.0002	0.0005	0.0002	0.0003	0.0006	0.0009	0.0003	0.0001	0.0004
sharpness-mag-init	0.0366	0.0460	0.0391	0.0191	0.0374	0.0373	0.0761	0.0368	0.0159	0.0134
sharpness-mag-orig	0.0125	0.0143	0.0195	0.0043	0.0120	0.0134	0.0142	0.0111	0.0036	0.0033
sharpness-orig	0.1117	0.2353	0.0809	0.0658	0.1223	0.1071	0.1254	0.1189	0.0547	0.0224
sharpness	0.0545	0.1596	0.0497	0.0156	0.0586	0.0599	0.0700	0.0583	0.0130	0.0123
spec-init	0.2536	0.3161	0.2295	0.1179	0.2532	0.2584	0.2540	0.2539	0.0980	0.0559
spec-orig-main	0.2266	0.2903	0.2072	0.0890	0.2255	0.2355	0.2262	0.2262	0.0739	0.0382
spec-orig	0.2197	0.2815	0.2045	0.0808	0.2180	0.2285	0.2181	0.2188	0.0671	0.0359
step-0.1-to-0.01	0.0125	0.0031	0.0055	0.0093	0.0074	0.0043	0.0070	0.0055	0.0026	0.0032
step-to-0.1	0.0349	0.0361	0.0397	0.1046	0.0485	0.0380	0.0568	0.0502	0.0303	0.0134
sum-of-fro-over-margin	0.1200	0.2269	0.1005	0.0440	0.1207	0.1060	0.1645	0.1227	0.0366	0.0110
sum-of-fro-over-sum-of-spec	0.0258	0.0392	0.0055	0.1111	0.0312	0.0194	0.0355	0.0297	0.0051	0.0027
sum-of-fro	0.1292	0.2286	0.1115	0.0441	0.1281	0.1134	0.1714	0.1300	0.0366	0.0119
sum-of-spec-over-margin	0.0089	0.0292	0.0406	0.0951	0.0089	0.0069	0.0054	0.0051	0.0054	0.0072
sum-of-spec	0.0127	0.0324	0.0466	0.0876	0.0117	0.0096	0.0080	0.0076	0.0079	0.0099
vc-dim	0.0422	0.0564	0.0518	0.0039	0.0422	0.0443	0.0627	0.0412	0.0033	0.0000
conditional entropy	0.9836	0.8397	0.9331	0.8308	0.9960	0.9746	0.9977	N/A	N/A	N/A

Table 10: All complexity measures (rows), hyperparameters (columns) and the **mutual information** we investigated.

C Extended Notation

Given any margin value $\gamma \geq 0$, we define the margin loss L_γ as follows:

$$L_\gamma(f_{\mathbf{w}}) \triangleq \mathbb{E}_{(\mathbf{X},y)\sim\mathscr{D}}\left[I\big(f_{\mathbf{w}}(\mathbf{X})[y] \leq \gamma + \max_{j\neq y} f_{\mathbf{w}}(\mathbf{X})[j]\big)\right] \tag{10}$$

and $\hat{L}_\gamma$ is defined in an analogous manner on the training set. Further, for any vector $\mathbf{v}$, we denote by $\|\mathbf{v}\|_2$ the ℓ_2 norm of $\mathbf{v}$. For any tensor $\mathbf{W}$, let $\|\mathbf{W}\|_F \triangleq \|\text{vec}(\mathbf{W})\|$. We also denote $\|\mathbf{W}\|_2$ as the spectral norm of the tensor $\mathbf{W}$ when used with a convolution operator. For convolutional operators, we compute the true singular value with the method proposed by Sedghi et al. (2018) through FFT.

We denote a tensor as $\mathbf{A}$, vector as $\boldsymbol{a}$, and scalar as A or a. For any $1 \leq j \leq k$, consider a k-th order tensor $\mathbf{A}$ and a j-th order tensor $\mathbf{B}$ where dimensions of $\mathbf{B}$ match the last j dimensions of $\mathbf{A}$. We then define the product operator $\otimes_j$:

$$(\mathbf{A} \otimes_j \mathbf{B})_{i_1,\ldots,i_{k-j}} \triangleq \langle \mathbf{A}_{i_1,\ldots,i_{k-j}}, \mathbf{B}\rangle\,, \tag{11}$$

where $i_1, \ldots, i_{k-j}$ are indices. We also assume that the input images have dimension $n \times n$ and there are κ classes. Given the number of input channels c_{in}, number of output channels c_{out}, 2D square kernel with side length k, stride s, and padding p, we define the convolutional layer $\text{conv}_{\mathbf{W},s,p}$ as follows:

$$\text{conv}_{\mathbf{W},s,p}(\mathbf{X})_{i_1,i_2} \triangleq \mathbf{W}\otimes_3 \text{patch}_{s(i_1-1)+1,s(i_2-1)+1,k}\left(\text{pad}_p(\mathbf{X})\right) \qquad \forall 1 \leq i_1, i_2 \leq \lfloor \frac{n+2p-k}{s} \rfloor \tag{12}$$

where $\mathbf{W} \in \mathbb{R}^{c_{\text{out}}\times c_{\text{in}}\times k\times k}$ is the convolutional parameter tensor, $\text{patch}_{i,j,k}(\mathbf{Z})$ is a $k \times k$ patch of $\mathbf{Z}$ starting from the point (i, j), and pad_p is the padding operator which adds p zeros to top, bottom, left and right of $\mathbf{X}$:

$$\text{pad}_p(\mathbf{X})_{i_1,i_2,j} = \begin{cases} \mathbf{X}_{i_1,i_2} & p < i_1, i_2 \leq n+p \\ 0 & \text{otherwise} \end{cases}. \tag{13}$$

We also define the max-pooling operator $\text{pool}_{k,s,p}$ as follows:

$$\text{pool}_{k,s,p}(\mathbf{X})_{i_1,i_2,j} = \max(\text{patch}_{s(i_1-1)+1,s(i_2-1)+1}\left(\text{pad}_p(\mathbf{X}_{:,:,j})\right)) \qquad \forall 1 \leq i_1, i_2 \leq \lfloor \frac{n+2p-k}{s} \rfloor \tag{14}$$

We denote by $f_{\mathbf{W},s}$ a convolutional network such that $\mathbf{W}_i \in \mathbb{R}^{c_i\times c_{i-1}\times k_i\times k_i}$ is the convolution tensor and s_i is the convolutional stride at layer i. At Layer i, we assume the sequence of convolution, ReLU and max-pooling where the max pooling has kernel k'_i and stride s'_i. Lack of max-pooling in some layers can be achieved by setting $k'_i = s'_i = 1$. We consider classification tasks and denote the number of classes by κ.

D Complexity Measures

In this section, we look at different complexity measures. When a measure μ is based on a generalization bound, we chose it so that the following is true with probability 0.99 (we choose the failure probability δ to be 0.01):

$$L \leq \hat{L} + \sqrt{\frac{\mu}{m}} \tag{15}$$

We also definecolorredconsider? measures which do not provably bound the generalization error and evaluate those.

Note that in almost all cases, the canonical ordering given based on some "common" assumptions are positively correlated with the generalization in terms of both τ and Ψ; however, for optimizer, the correlation τ is close to 0. This implies that the choice of optimizer is only essentially uncorrelated with the generalization gap in the range of models we consider. This ordering helps validate many techniques used by the practioners.

D.1 VC-Dimension Based Measures

We start by restating the theorem in (Bartlett et al., 2019) which provides an upper bound on the VC-dimension of any piece-wise linear network.

Theorem 1 (Bartlett et al. (2019)) *Let $\mathcal{F}$ be the class of feed-forward networks with a fixed computation graph of depth d and ReLU activations. Let a_i and q_i be the number of activations and parameters in layer i. Then VC-dimension of $\mathcal{F}$ can be bounded as follows:*

$$VC(\mathcal{F}) \leq d + \left(\sum_{i=1}^{d}(d-i+1)q_i\right)\log_2\left(8e\sum_{i=1}^{d} ia_i \log_2\left(4e\sum_{j=1}^{d} ja_j\right)\right)$$

Theorem 2 *Given a convolutional network f, for any $\delta > 0$, with probability $1-\delta$ over the the training set:*

$$L \leq \hat{L} + 4000\sqrt{\frac{d\log_2\left(6dn\right)^3 \sum_{i=1}^{d} k_i^2 c_i c_{i-1}}{m}} + \sqrt{\frac{\log(1/\delta)}{m}} \tag{16}$$

Proof We simplify the bound in Theorem 1 using a d' to refer to the depth instead of d:

$$\begin{aligned}
\mathrm{VC}(\mathcal{F}) &\leq d' + \left(\sum_{i=1}^{d'}(d-i+1)q_i\right)\log_2\left(8e\sum_{i=1}^{d'} ia_i \log_2\left(4e\sum_{j=1}^{d'} ja_j\right)\right) \\
&\leq d' + \left(\sum_{i=1}^{d'}(d'-i+1)q_i\right)\log_2\left(8e\sum_{i=1}^{d'} ia_i\right)^2 \\
&\leq d' + 2\log_2\left(8e\sum_{i=1}^{d'} ia_i\right)\sum_{i=1}^{d'}(d'-i+1)q_i \\
&\leq 3d'\log_2\left(8e\sum_{i=1}^{d'} ia_i\right)\sum_{i=1}^{d'} q_i
\end{aligned}$$

In order to extend the above bound to a convolutional network, we need to present a pooling layer with ReLU activations. First note that maximum of two inputs can be calculated using two layers with ReLU and linear activations as $\max(x_1, x_2) = x_1 + ReLU(x_2 - x_1)$. Now, since max-pooling at layer i has kernel sizes k'_i, we need $\lceil 4\log_2(k'_i)\rceil$ layers to present that but given that the kernel size of the max-pooling layer is at most size of the image, we have

$$\lceil 4\log_2(k'_i)\rceil \leq \lceil 4\log_2(n^2)\rceil \leq \lceil 8\log_2(n)\rceil \leq 9\log_2(n)$$

Therefore, we have $d' \leq 9d\log_2(n)$. The number of activations in any of these layers is at most $n^2 c_i$ since there are at most n^2 pairs of neighbor pixels in an $n \times n$ image with c_i channels. We ignore strides when calculating the upper bound since it only reduces number of activations at a few layers and does not change the bound significantly. Using these

bounds ond', a_i and q_i the equivalent network, we can bound the VC dimension as follows:

$$\begin{aligned}
\mathrm{VC}(\mathcal{F}) &\leq 27d\log_2(n)\log_2\left(8e(9d\log_2(n))^2n^2\right)(9\log_2(n))\sum_{i=1}^{d}k_i^2c_{i-1}(c_i+1) \\
&\leq 729d\log_2(n)^2\log_2(6dn)\sum_{i=1}^{d}k_i^2c_{i-1}(c_i+1) \\
&\leq 729d\log_2(6dn)^3\sum_{i=1}^{d}k_i^2c_{i-1}(c_i+1)
\end{aligned}$$

For binary classifiers, generalization error can be in terms of Rademacher complexity (Mohri et al., 2012) which in turn can be bounded by $72\sqrt{\mathrm{VC}/m}$ (Kontorovich, 2016). Therefore, we can get the following[6] generalization bound:

$$L \leq \hat{L} + 144\sqrt{\frac{VC(\mathcal{F})}{m}} + \sqrt{\frac{\log(1/\delta)}{m}} \tag{17}$$

For multi-class classification, the generalization error can be similarly bounded by Graph dimension which is an extension of VC-dimension. A simple approach get a bound on Graph dimension is to consider all pairs of classes as binary classification problem which bounds the graph dimension by $\kappa^2\ VC(\mathcal{F})$. There, putting everything together, we get the following generalization bound:

$$L \leq \hat{L} + 4000\kappa\sqrt{\frac{d\log_2(6dn)^3\sum_{i=1}^{d}k_i^2c_{i-1}(c_i+1)}{m}} + \sqrt{\frac{\log(1/\delta)}{m}} \tag{18}$$

Inspired by Theorem 2, we define the following VC-based measure for generalization:

$$\mu_{VC}(f_{\mathbf{w}}) = \left(4000\kappa\sqrt{d\log_2(6dn)^3\sum_{i=1}^{d}k_i^2c_{i-1}(c_i+1)} + \sqrt{\log(1/\delta)}\right)^2 \tag{19}$$

Since some of the dependencies in the above measure are probably proof artifacts, we also define another measure that is nothing but the number of parameters of the model:

$$\mu_{\text{param}} = \sum_{i=1}^{d}k_i^2c_{i-1}(c_i+1) \tag{20}$$

D.1.1 Measures on the output of the network

While measures that can be calculated only based on the output of the network cannot reveal complexity of the network, they can still be very informative for predicting generalization. Therefore, we define a few measures that can be calculated solely based on the output of the network.

We start by looking at the cross-entropy over the output. Even though we used a cross-entropy based stopping criterion, the cross-entropy of the final models is not exactly the same as the stopping criterion and it could be informative. Hence we define the following measure:

$$\mu_{\text{cross-entropy}} = \frac{1}{m}\sum_{i=1}^{m}\ell(f_{\mathbf{w}}(\mathbf{X}_i), y_i) \tag{21}$$

where ℓ is the cross entropy loss.

[6]The generalization gap is bounded by two times Rademacher Complexity, hence the constant 144.

Another useful and intuitive notion that appears in generalization bounds is margin. In all measures that involve margin γ, we set the margin γ to be the 10-th percentile of the margin values on the training set and therefore ensuring $\hat{L}_\gamma \leq 0.1$. Even though margin alone is not a sensible generalization measure and can be artificially increased by scaling up the magnitude of the weights, it could still reveal information about training dynamics and therefore be informative. We report the following measure based on the margin:

$$\mu_{1/\text{margin}}(f_{\mathbf{w}}) = \frac{1}{\gamma^2} \tag{22}$$

Finally, entropy of the output is another interesting measure and it has been shown that regularizing it can improve generalization in deep learning (Pereyra et al., 2017). With a fixed cross-entropy, increasing the entropy corresponds to distribute the uncertainty of the predictions equally among the wrong labels which is connected to label smoothing and increasing the margin. We define the following measure which is the negative entropy of the output of the network:

$$\mu_{\text{neg-entropy}}(f_{\mathbf{w}}) = \frac{1}{m}\sum_{i=1}^{m}\sum_{j=1}^{\kappa} p_i[j]\log(p_i[j]) \tag{23}$$

where $p_i[j]$ is the predicted probability of the class j for the input data $\mathbf{X}_i$.

D.2 Norm/Margin Based Measures

Several generalization bounds have been proved for neural networks using margin and norm notions. In this section, we go over several such measures. For fully connected networks, Bartlett & Mendelson (2002) have shown a bound based on product of $\ell_{1,\infty}$ norm of the layer weights times a 2^d factor where $\ell_{1,\infty}$ is the maximum over hidden units of the ℓ_2 norm of the incoming weights to the hidden unit. Neyshabur et al. (2015b) proved a bound based on product of Frobenius norms of the layer weights times a 2^d factor and Golowich et al. (2017) was able to improve the factor to $\sqrt{d}$. Bartlett et al. (2017) proved a bound based on product of spectral norm of the layer weights times sum over layers of ratio of Frobenius norm to spectral norm of the layer weights and Neyshabur et al. (2018a) showed a similar bound can be achieved in a simpler way using PAC-bayesian framework.

Spectral Norm Unfortunately, none of the above founds are directly applicable to convolutional networks. Pitas et al. (2017) built on Neyshabur et al. (2018a) and extended the bound on the spectral norm to convolutional networks. The bound is very similar to the one for fully connected networks by Bartlett et al. (2017). We next restate their generalization bound for convolutional networks including the constants.

Theorem 3 (Pitas et al. (2017)) *Let B an upper bound on the ℓ_2 norm of any point in the input domain. For any $B, \gamma, \delta > 0$, the following bound holds with probability $1-\delta$ over the training set:*

$$L \leq \hat{L}_\gamma + \sqrt{\frac{\left(84B\sum_{i=1}^{d} k_i\sqrt{c_i} + \sqrt{\ln(4n^2d)}\right)^2 \prod_{i=1}^{d}\|\mathbf{W}_i\|_2^2 \sum_{j=1}^{d}\frac{\|\mathbf{W}_j-\mathbf{W}_j^0\|_F^2}{\|\mathbf{W}_j\|_2^2} + \ln(\frac{m}{\delta})}{\gamma^2 m}} \tag{24}$$

Inspired by the above theorem, we define the following spectral measure:

$$\mu_{\text{spec,init}}(f_{\mathbf{w}}) = \frac{\left(84B\sum_{i=1}^{d} k_i\sqrt{c_i} + \sqrt{\ln(4n^2d)}\right)^2 \prod_{i=1}^{d}\|\mathbf{W}_i\|_2^2 \sum_{j=1}^{d}\frac{\|\mathbf{W}_j-\mathbf{W}_j^0\|_F^2}{\|\mathbf{W}_j\|_2^2} + \ln(\frac{m}{\delta})}{\gamma^2} \tag{25}$$

The generalization bound in Theorem 3 depends on reference tensors $\mathbf{W}_i^0$. We chose the initial tensor as the reference in the above measure but another reasonable choice is the

origin which gives the following measures:

$$\mu_{\text{spec-orig}}(f_{\mathbf{w}}) = \frac{\left(84B\sum_{i=1}^{d} k_i\sqrt{c_i} + \sqrt{\ln(4n^2d)}\right)^2 \prod_{i=1}^{d} \|\mathbf{W}_i\|_2^2 \sum_{j=1}^{d} \frac{\|\mathbf{W}_j\|_F^2}{\|\mathbf{W}_j\|_2^2} + \ln(\frac{m}{\delta})}{\gamma^2} \tag{26}$$

Since some of the terms in the generalization bounds might be proof artifacts, we also measure the main terms in the generalization bound:

$$\mu_{\text{spec-init-main}}(f_{\mathbf{w}}) = \frac{\prod_{i=1}^{d} \|\mathbf{W}_i\|_2^2 \sum_{j=1}^{d} \frac{\left\|\mathbf{W}_j - \mathbf{W}_j^0\right\|_F^2}{\|\mathbf{W}_j\|_2^2}}{\gamma^2} \tag{27}$$

$$\mu_{\text{spec-orig-main}}(f_{\mathbf{w}}) = \frac{\prod_{i=1}^{d} \|\mathbf{W}_i\|_2^2 \sum_{j=1}^{d} \frac{\|\mathbf{W}_j\|_F^2}{\|\mathbf{W}_j\|_2^2}}{\gamma^2} \tag{28}$$

We further look at the main two terms in the bound separately to be able to differentiate their contributions.

$$\mu_{\text{spec-init-main}}(f_{\mathbf{w}}) = \frac{\prod_{i=1}^{d} \|\mathbf{W}_i\|_2^2 \sum_{j=1}^{d} \frac{\left\|\mathbf{W}_j - \mathbf{W}_j^0\right\|_F^2}{\|\mathbf{W}_j\|_2^2}}{\gamma^2} \tag{29}$$

$$\mu_{\text{spec-orig-main}}(f_{\mathbf{w}}) = \frac{\prod_{i=1}^{d} \|\mathbf{W}_i\|_2^2 \sum_{j=1}^{d} \frac{\|\mathbf{W}_j\|_F^2}{\|\mathbf{W}_j\|_2^2}}{\gamma^2} \tag{30}$$

$$\mu_{\text{prod-of-spec/margin}}(f_{\mathbf{w}}) = \frac{\prod_{i=1}^{d} \|\mathbf{W}_i\|_2^2}{\gamma^2} \tag{31}$$

$$\mu_{\text{prod-of-spec}}(f_{\mathbf{w}}) = \prod_{i=1}^{d} \|\mathbf{W}_i\|_2^2 \tag{32}$$

$$\mu_{\text{fro/spec}}(f_{\mathbf{w}}) = \sum_{i=1}^{d} \frac{\|\mathbf{W}_i\|_F^2}{\|\mathbf{W}_i\|_2^2} \tag{33}$$

Finally, since product of spectral norms almost certainly increases with depth, we look at the following measure which is equal to the sum over squared spectral norms after rebalancing the layers to have the same spectral norms:

$$\mu_{\text{sum-of-spec/margin}}(f_{\mathbf{w}}) = d\left(\frac{\prod_{i=1}^{d} \|\mathbf{W}_i\|_2^2}{\gamma^2}\right)^{1/d} \tag{34}$$

$$\mu_{\text{sum-of-spec}}(f_{\mathbf{w}}) = d\left(\prod_{i=1}^{d} \|\mathbf{W}_i\|_2^2\right)^{1/d} \tag{35}$$

Frobenius Norm The generalization bound given in Neyshabur et al. (2015b) is not directly applicable to convolutional networks. However, Since for each layer i, we have $\|\mathbf{W}_i\|_2 \leq k_i^2 \|\mathbf{W}_i\|_F$ and therefore by Theorem 3, we can get an upper bound on the test error based on product of Frobenius norms. Therefore, we define the following measure based on the product of Frobenius norms:

$$\mu_{\text{prod-of-fro/margin}}(f_{\mathbf{w}}) = \frac{\prod_{i=1}^{d} \|\mathbf{W}_i\|_F^2}{\gamma^2} \tag{36}$$

$$\mu_{\text{prod-of-fro}}(f_{\mathbf{w}}) = \prod_{i=1}^{d} \|\mathbf{W}_i\|_F^2 \tag{37}$$

We also look at the following measure with correspond to sum of squared Frobenius norms of the layers after rebalancing them to have the same norm:

$$\mu_{\text{sum-of-fro/margin}}(f_{\mathbf{w}}) = d\left(\frac{\prod_{i=1}^{d}\|\mathbf{W}_i\|_F^2}{\gamma^2}\right)^{1/d} \tag{38}$$

$$\mu_{\text{sum-of-fro}}(f_{\mathbf{w}}) = d\left(\prod_{i=1}^{d}\|\mathbf{W}_i\|_F^2\right)^{1/d} \tag{39}$$

Finally, given recent evidence on the importance of distance to initialization (Dziugaite & Roy, 2017; Nagarajan & Kolter, 2019b; Neyshabur et al., 2018b), we calculate the following measures:

$$\mu_{\text{frobenius-distance}}(f_{\mathbf{w}}) = \sum_{i=1}^{d}\left\|\mathbf{W}_i - \mathbf{W}_i^0\right\|_F^2 \tag{40}$$

$$\mu_{\text{dist-spec-init}}(f_{\mathbf{w}}) = \sum_{i=1}^{d}\left\|\mathbf{W}_i - \mathbf{W}_i^0\right\|_2^2 \tag{41}$$

In case when the reference matrix $\mathbf{W}_i^0 = 0$ for all weights, eq. 40 the Frobenius norm of the parameters which also correspond to distance from the origin:

$$\mu_{\text{param-norm}}(f_{\mathbf{w}}) = \sum_{i=1}^{d}\|\mathbf{W}_i\|_F^2 \tag{42}$$

Path-norm Path-norm was introduced in Neyshabur et al. (2015b) as an scale invariant complexity measure for generalization and is shown to be a useful geometry for optimization Neyshabur et al. (2015a). To calculate path-norm, we square the parameters of the network, do a forward pass on an all-ones input and then take square root of sum of the network outputs. We define the following measures based on the path-norm:

$$\mu_{\text{path-norm/margin}}(f_{\mathbf{w}}) = \frac{\sum_i f_{\mathbf{w}^2}(\mathbf{1})[\mathbf{i}]}{\gamma^2} \tag{43}$$

$$\mu_{\text{path-norm}}(f_{\mathbf{w}}) = \sum_i f_{\mathbf{w}^2}(\mathbf{1}) \tag{44}$$

where $\mathbf{w}^2 = \mathbf{w} \circ \mathbf{w}$ is the element-wise square operation on the parameters.

Fisher-Rao Norm Fisher-Rao metric was introduced in Liang et al. (2017) as a complexity measure for neural networks. Liang et al. (2017) showed that Fisher-Rao norm is a lower bound on the path-norm and it correlates in some cases. We define a measure based on the Fisher-Rao matric of the network:

$$\mu_{\text{Fisher-Rao}}(f_{\mathbf{w}}) = \frac{(d+1)^2}{m}\sum_{i=1}^{m}\langle \mathbf{w}, \nabla_{\mathbf{w}}\ell(f_{\mathbf{w}}(\mathbf{X}_i)), y_i\rangle^2 \tag{45}$$

where ℓ is the cross-entropy loss.

D.3 Flatness-based Measures

PAC-Bayesian framework (McAllester, 1999) allows us to study flatness of a solution and connect it to generalization. Given a prior P is is chosen before observing the training set and a posterior Q which is a distribution on the solutions of the learning algorithm (and hence depends on the training set), we can bound the expected generalization error of solutions generated from Q with high probability based on the KL divergence of P and Q. The next theorem states a simplified version of PAC-Bayesian bounds.

Theorem 4 *For any $\delta > 0$, distribution D, prior P, with probability $1-\delta$ over the training set, for any posterior Q the following bound holds:*

$$\mathbb{E}_{\mathbf{v}\sim Q}\left[L(f_{\mathbf{v}})\right] \leq \mathbb{E}_{\mathbf{w}\sim Q}\left[\hat{L}(f_{\mathbf{v}})\right] + \sqrt{\frac{\mathrm{KL}(Q||P) + \log\left(\frac{m}{\delta}\right)}{2(m-1)}} \tag{46}$$

If P and Q are Gaussian distributions with $P = \mathscr{N}(\mu_P, \Sigma_P)$ amd $Q = \mathscr{N}(\mu_Q, \Sigma_Q)$, then the KL-term can be written as follows:

$$\mathrm{KL}(\mathscr{N}(\mu_Q, \Sigma_Q)||\mathscr{N}(\mu_P, \Sigma_P)) = \frac{1}{2}\left[\mathrm{tr}\left(\Sigma_P^{-1}\Sigma_Q\right) + \left(\mu_Q - \mu_P\right)^\top \Sigma_P^{-1}\left(\mu_Q - \mu_P\right) - k + \ln(\frac{\det \Sigma_P}{\det \Sigma_Q})\right].$$

Setting $Q = \mathscr{N}(\mathbf{w}, \sigma^2 I)$ and $P = \mathscr{N}(\mathbf{w}^0, \sigma^2 I)$ similar to Neyshabur et al. (2017), the KL term will be simply $\frac{\|\mathbf{w}-\mathbf{w}^0\|_2^2}{2\sigma^2}$. However, since σ belongs to prior, if we search to find a value for σ, we need to adjust the bound to reflect that. Since we search over less than 20000 predefined values of σ in our experiments, we can use the union bound which changes the logarithmic term to $\log(20000m/\delta)$ and we get the following bound:

$$\mathbb{E}_{\mathbf{u}\sim\mathscr{N}(u,\sigma^2 I)}\left[L(f_{\mathbf{w}+\mathbf{u}})\right] \leq \mathbb{E}_{\mathbf{u}\sim\mathscr{N}(u,\sigma^2 I)}\left[\hat{L}(f_{\mathbf{w}+\mathbf{u}})\right] + \sqrt{\frac{\frac{\|\mathbf{w}-\mathbf{w}^0\|_2^2}{4\sigma^2} + \log(\frac{m}{\sigma}) + 10}{m-1}} \tag{47}$$

Based on the above bound, we define the following measures using the origin and initialization as reference tensors:

$$\mu_{\text{pac-bayes-init}}(f_{\mathbf{w}}) = \frac{\left\|\mathbf{w}-\mathbf{w}^0\right\|_2^2}{4\sigma^2} + \log(\frac{m}{\sigma}) + 10 \tag{48}$$

$$\mu_{\text{pac-bayes-orig}}(f_{\mathbf{w}}) = \frac{\|\mathbf{w}\|_2^2}{4\sigma^2} + \log(\frac{m}{\delta}) + 10 \tag{49}$$

where σ is chosen to be the largest number such that $\mathbb{E}_{\mathbf{u}\sim\mathscr{N}(u,\sigma^2 I)}\left[\hat{L}(f_{\mathbf{w}+\mathbf{u}})\right] \leq 0.1$.

The above framework captures flatness in the expected sense since we add Gaussian perturbations to the parameters. Another notion of flatness is the worst-case flatness where we search for the direction that changes the loss the most. This is motivated by (Keskar et al., 2016) where they observe that this notion would correlate to generalization in the case of different batch sizes. We can use PAC-Bayesian framework to give generalization bounds for worst-case perturbations as well. The magnitude of a Gaussian variable with with variance σ^2 is at most $\sigma\sqrt{2\log(2/\delta)}$ with probability $1-\delta/2$. Applying a union bound on all parameters, we get that with probability $1-\delta/2$ the magnitude of the Gaussian noise is at most $\alpha = \sigma\sqrt{2\log(2\omega/\delta)}$ where ω is the number of parameters of the model. Therefore, we can get the following generalization bound:

$$\mathbb{E}_{\mathbf{u}\sim\mathscr{N}(u,\sigma^2 I)}\left[L(f_{\mathbf{w}+\mathbf{u}})\right] \leq \max_{|u_i|\leq\alpha} \hat{L}(f_{\mathbf{w}+\mathbf{u}}) + \sqrt{\frac{\frac{\|\mathbf{w}-\mathbf{w}^0\|_2^2 \log(2\omega/\delta)}{2\alpha^2} + \log(\frac{2m}{\delta}) + 10}{m-1}} \tag{50}$$

Inspired by the above bound, we define the following measures:

$$\mu_{\text{sharpness-init}}(f_{\mathbf{w}}) = \frac{\left\|\mathbf{w}-\mathbf{w}^0\right\|_2^2 \log(2\omega)}{4\alpha^2} + \log(\frac{m}{\sigma}) + 10 \tag{51}$$

$$\mu_{\text{sharpness-orig}}(f_{\mathbf{w}}) = \frac{\|\mathbf{w}\|_2^2 \log(2\omega)}{4\alpha^2} + \log(\frac{m}{\delta}) + 10 \tag{52}$$

where α is chosen to be the largest number such that $\max_{|u_i|\leq\alpha} \hat{L}(f_{\mathbf{w}+\mathbf{u}}) \leq 0.1$.

To understand the importance of the flatness parameters σ and α, we also define the following measures:

$$\mu_{\text{pac-bayes-flatness}}(f_{\mathbf{w}}) = \frac{1}{\sigma^2} \tag{53}$$

$$\mu_{\text{sharpness-flatness}}(f_{\mathbf{w}}) = \frac{1}{\alpha^2} \tag{54}$$

where α and σ are computed as explained above.

Magnitude-aware Perturbation Bounds The magnitude of perturbation in (Keskar et al., 2016) was chosen so that for each parameter the ratio of magnitude of perturbation to the magnitude of the parameter is bounded by a constant α'[7]. Following a similar approach, we can choose the posterior for parameter i in PAC-Bayesian framework to be $\mathcal{N}(w_i, \sigma'^2|w_i|^2 + \epsilon^2)$. Now, substituting this in the Equation equation D.3 and solving for the prior $\mathcal{N}(\mathbf{w}^0, \sigma_P^2)$ that minimizes the KL term by setting the gradient with respect to σ_2^P to zero, KL can be written as follows:

$$
\begin{aligned}
2\mathrm{KL}(Q||P) &= \omega \log\left(\frac{\sigma'^2+1}{\omega}\left\|\mathbf{w}-\mathbf{w}^0\right\|_2^2+\epsilon^2\right) - \sum_{i=1}^{\omega}\log\left(\sigma'^2|w_i-w_i^0|^2+\epsilon^2\right) \\
&= \sum_{i=1}^{\omega}\log\left(\frac{\epsilon^2+(\sigma'^2+1)\left\|\mathbf{w}-\mathbf{w}^0\right\|_2^2/\omega}{\epsilon^2+\sigma'^2|w_i-w_i^0|^2}\right)
\end{aligned}
$$

Therefore, the generalization bound can be written as follows

$$
\mathbb{E}_{\mathbf{u}}\left[L(f_{\mathbf{w}+\mathbf{u}})\right] \leq \mathbb{E}_{\mathbf{u}}\left[\hat{L}(f_{\mathbf{w}+\mathbf{u}})\right] + \sqrt{\frac{\frac{1}{4}\sum_{i=1}^{\omega}\log\left(\frac{\epsilon^2+(\sigma'^2+1)\|\mathbf{w}-\mathbf{w}^0\|_2^2/\omega}{\epsilon^2+\sigma'^2|w_i-w_i^0|^2}\right)+\log(\frac{m}{\delta})+10}{m-1}} \tag{55}
$$

where $u_i \sim \mathcal{N}(0, \sigma'^2|w_i| + \epsilon^2)$, $\epsilon = 1e-3$ and σ' is chosen to be the largest number such that $\mathbb{E}_{\mathbf{u}}\left[\hat{L}(f_{\mathbf{w}+\mathbf{u}})\right] \leq 0.1$. We define the following measures based on the generalization bound:

$$
\mu_{\text{pac-bayes-mag-init}}(f_{\mathbf{w}}) = \frac{1}{4}\sum_{i=1}^{\omega}\log\left(\frac{\epsilon^2+(\sigma'^2+1)\left\|\mathbf{w}-\mathbf{w}^0\right\|_2^2/\omega}{\epsilon^2+\sigma'^2|w_i-w_i^0|^2}\right)+\log(\frac{m}{\delta})+10 \tag{56}
$$

$$
\mu_{\text{pac-bayes-mag-orig}}(f_{\mathbf{w}}) = \frac{1}{4}\sum_{i=1}^{\omega}\log\left(\frac{\epsilon^2+(\sigma'^2+1)\left\|\mathbf{w}\right\|_2^2/\omega}{\epsilon^2+\sigma'^2|w_i-w_i^0|^2}\right)+\log(\frac{m}{\delta})+10 \tag{57}
$$

We also follow similar arguments are before to get a similar bound on the worst-case sharpness:

$$
\mathbb{E}_{\mathbf{u}}\left[L(f_{\mathbf{w}+\mathbf{u}})\right] \leq \max_{|u_i|\leq\alpha'|w_i|+\epsilon}\hat{L}(f_{\mathbf{w}+\mathbf{u}}) + \sqrt{\frac{\frac{1}{4}\sum_{i=1}^{\omega}\log\left(\frac{\epsilon^2+(\alpha'^2+4\log(2\omega/\delta))\|\mathbf{w}-\mathbf{w}^0\|_2^2/\omega}{\epsilon^2+\alpha'^2|w_i-w_i^0|^2}\right)+\log(\frac{m}{\delta})+10}{m-1}} \tag{58}
$$

We look at the following measures based on the above bound:

$$
\mu_{\text{pac-sharpness-mag-init}}(f_{\mathbf{w}}) = \frac{1}{4}\sum_{i=1}^{\omega}\log\left(\frac{\epsilon^2+(\alpha'^2+4\log(2\omega/\delta))\left\|\mathbf{w}-\mathbf{w}^0\right\|_2^2/\omega}{\epsilon^2+\alpha'^2|w_i-w_i^0|^2}\right)+\log(\frac{m}{\delta})+10 \tag{59}
$$

$$
\mu_{\text{pac-sharpness-mag-orig}}(f_{\mathbf{w}}) = \frac{1}{4}\sum_{i=1}^{\omega}\log\left(\frac{\epsilon^2+(\alpha'^2+4\log(2\omega/\delta))\left\|\mathbf{w}\right\|_2^2/\omega}{\epsilon^2+\alpha'^2|w_i-w_i^0|^2}\right)+\log(\frac{m}{\delta})+10 \tag{60}
$$

Finally, we look at measures that are only based the sharpness values computed above:

$$
\mu_{\text{pac-bayes-mag-flat}}(f_{\mathbf{w}}) = \frac{1}{\sigma'^2} \tag{61}
$$

$$
\mu_{\text{sharpness-mag-flat}}(f_{\mathbf{w}}) = \frac{1}{\alpha'^2} \tag{62}
$$

where α and σ are computed as explained above.

[7] They actually used a slightly different version which is a combination of the two perturbation bounds we calculated here. Here, for more clarity, we decomposed it into two separate perturbation bounds.

D.4 Optimization-based Measures

There are mixed results about how the optimization speed is relevant to generalization. On one hand we know that adding Batch Normalization or using shortcuts in residual architectures help both optimization and generalization and Hardt et al. (2015) suggests that faster optimization results in better generalization. On the other hand, there are empirical results showing that adaptive optimization methods that are faster, usually generalize worse (Wilson et al., 2017b). Here, we put these hypothesis into test by looking at the number of steps to achieve cross entropy 0.1 and the number of steps needed to go from cross-entropy 0.1 to 0.01:

$$\mu_{\text{\#steps-0.1-loss}}(f_{\mathbf{w}}) = \text{\#steps from initialization to 0.1 cross-entropy} \tag{63}$$

$$\mu_{\text{\#steps-0.1-0.01-loss}}(f_{\mathbf{w}}) = \text{\#steps from 0.1 to 0.01 cross-entropy} \tag{64}$$

The above measures tell us if the speed of optimization at early or late stages can be informative about generalization. We also define measures that look at the SGD gradient noise after the first epoch and at the end of training at cross-entropy 0.01 to test the gradient noise can be predictive of generalization:

$$\mu_{\text{grad-noise-epoch1}}(f_{\mathbf{w}}) = \text{Var}_{(\mathbf{X},y)\ S}\left(\nabla_{\mathbf{w}}\ell(f_{\mathbf{w}^1}(\mathbf{X}), y)\right) \tag{65}$$

$$\mu_{\text{grad-noise-final}}(f_{\mathbf{w}}) = \text{Var}_{(\mathbf{X},y)\ S}\left(\nabla_{\mathbf{w}}\ell(f_{\mathbf{w}}(\mathbf{X}), y)\right) \tag{66}$$

where $\mathbf{w}^1$ is the weight vector after the first epoch.

E Algorithms

We first lay out some common notations used in the pseudocode:

1. f: the architecture that takes parameter θ and input x and map to $f(x;\theta)$ which is the predicted label of x
2. θ: parameters
3. M: Some kind of iteration; M_1: binary search depth; M_2: Monte Carlo Estimation steps; M_3: Iteration for estimating the loss
4. $\mathscr{D} = \{(x_i, y_i)\}_{i=0}^n$ the dataset the model is trained on; $\mathcal{B}$ as a uniformly sampled minibatch from the dataset.

Both search algorithm relies on the assumption that the loss increases monotonically with the perturbation magnitude σ around the final weight. This assumption is quite mild and in reality holds across almost all the models in this study.

Algorithm 1 EstimateAccuracy

Inputs: model f, parameter θ, dataset $\mathscr{D}$, estimate iteration M
Initialize Accuracy $= 0$
for episode $i = 1$ **to** M **do**
  $\mathcal{B} \sim \text{sample}(\mathscr{D})$
  Accuracy $+= \frac{1}{|\mathcal{B}|}\sum_i \delta(y_i = f(\mathcal{B}_i;\theta))$
end for
return Accuracy/M

Note that for finding the sharpness σ, we use the cross entropy as the differentiable surrogate object instead of the 1-0 loss which is in general not differentiable. Using gradient ascent brings another additional challenge that is for a converged model, the local gradient signal is usually weak, making gradient ascent extremely inefficient. To speed up thie process, we add a uniform noise with range being $[-\sigma_{new}/N_{\mathbf{w}}, \sigma_{new}/N_{\mathbf{w}}]$ to lift the weight off the flat minima where $N_{\mathbf{w}}$ is the number of parameters. This empirical greatly accelerates the search.

Algorithm 2 Find σ for PAC-Bayesian Bound

1: **Inputs:** f, θ_0, model accuracy ℓ, target accuracy deviation d, Upper bound $\sigma_{\max}$, Lower bound $\sigma_{\min}$, M_1, M_2, M_3
2: **Initialize**
3: **for** episode $i = 1$ **to** M_1 **do**
4: $\quad \sigma_{new} = (\sigma_{\max} + \sigma_{\min})/2$
5: $\quad \hat{\ell} = 0$
6: $\quad$ **for** step $j = 0$ **to** M_2 **do**
7: $\quad\quad \theta \leftarrow \theta_0 + \mathcal{N}(0, \sigma_{new}^2 I)$
8: $\quad\quad \hat{\ell} = \hat{\ell} + \text{EstimateAccuracy}(f, \theta_{new}, \mathscr{D}, M_3)$
9: $\quad$ **end for**
10: $\quad \hat{\ell} = \hat{\ell}/M_2$
11: $\quad \hat{d} = |\ell - \hat{\ell}|$
12: $\quad$ **if** $\hat{d} < \epsilon_d$ **or** $\sigma_{\max} - \sigma_{\min} < \epsilon_\sigma$ **then**
13: $\quad\quad$ **return** σ_{new}
14: $\quad$ **end if**
15: $\quad$ **if** $\hat{d} > d$ **then**
16: $\quad\quad \sigma_{\max} = \sigma_{new}$
17: $\quad$ **else**
18: $\quad\quad \sigma_{\min} = \sigma_{new}$
19: $\quad$ **end if**
20: **end for**

Algorithm 3 Find σ for Sharpness Bound

1: **Inputs:** f, θ_0, loss function $\mathcal{L}$, model accuracy ℓ, target accuracy deviation d, Upper bound $\sigma_{\max}$, Lower bound $\sigma_{\min}$, M_1, M_2, M_3, gradient steps M_4
2: **Initialize**
3: **for** episode $i = 1$ **to** M_1 **do**
4: $\quad \sigma_{new} = (\sigma_{\max} + \sigma_{\min})/2$
5: $\quad \hat{\ell} = \infty$
6: $\quad$ **for** step $j = 0$ **to** M_2 **do**
7: $\quad\quad \theta = \theta_0 + \mathcal{U}(\sigma_{new}/2)$
8: $\quad\quad$ **for** step $k = 0$ **to** M_4 **do**
9: $\quad\quad\quad \mathcal{B} \sim \text{sample}(\mathscr{D})$
10: $\quad\quad\quad \theta = \theta + \eta \nabla_\theta \ell(f, \mathcal{B}, \theta)$
11: $\quad\quad\quad$ **if** $||\theta|| > \sigma_{new}$ **then**
12: $\quad\quad\quad\quad \theta = \sigma_{new} \cdot \frac{\theta}{||\theta||}$
13: $\quad\quad\quad$ **end if**
14: $\quad\quad$ **end for**
15: $\quad\quad \hat{\ell} = \min(\hat{\ell}, \text{EstimateAccuracy}(f, \theta_{new}, \mathscr{D}, M_3))$
16: $\quad$ **end for**
17: $\quad \hat{d} = |\ell - \hat{\ell}|$
18: $\quad$ **if** $\hat{d} < \epsilon_d$ **or** $\sigma_{\max} - \sigma_{\min} < \epsilon_\sigma$ **then**
19: $\quad\quad$ **return** σ_{new}
20: $\quad$ **end if**
21: $\quad$ **if** $\hat{d} > d$ **then**
22: $\quad\quad \sigma_{\max} = \sigma_{new}$
23: $\quad$ **else**
24: $\quad\quad \sigma_{\min} = \sigma_{new}$
25: $\quad$ **end if**
26: **end for**

Further, for magnitude aware version of the bounds, the overall algorithm stays the same with the exception that now covariance matrices at line 7 of Algorithm 2 become as diagonal matrix containing w_i^2 on the diagonal; similarly, for line 12 of Algorithm 3, the weight

clipping of each w_i is conditioned on $\sigma_{new}|w_i|$, i.e. clipped to $[-\sigma_{new}|w_i|, \sigma_{new}|w_i|]$. Here w_i denotes the ith parameter of flattened $\mathbf{w}$.