

---

# Temperature Balancing, Layer-wise Weight Analysis, and Neural Network Training

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Regularization in modern machine learning is crucial, and it can take various forms  
2 in algorithmic design: training set, model family, error function, regularization  
3 terms, and optimizations. The learning rate, which can be interpreted as a *tempera-*  
4 *ture* parameter within the statistical mechanics of learning, plays a crucial role in  
5 training deep neural networks. Indeed, many widely adopted strategies define the  
6 decay of the learning rate over time, using either a global learning rate or one that  
7 varies for each parameter, which can be interpreted as decreasing the temperature.  
8 This paper proposes a middle-ground approach for temperature balancing called  
9 TempBalance. It is based on the theory of heavy-tail self-regularization (HT-SR),  
10 and it is a simple yet effective layer-wise policy applicable to general global tem-  
11 perature assignments in deep learning regularization. Our main contributions are as  
12 follows: (i) In addition to following a learning rate schedule, we suggest balancing  
13 the learning rate across each layer, an approach that has received less attention com-  
14 pared to global or parameter-wise learning rate allocation. (ii) We demonstrate that  
15 HT-SR-motivated capacity control metrics characterize the layers to achieve maxi-  
16 mal temperature balance during model training, resulting in improved performance  
17 during testing. We implement TempBalance on CIFAR10, CIFAR100, SVHN,  
18 and TinyImageNet datasets using ResNets, VGGs and WideResNets with various  
19 depths and widths. Our results show that TempBalance significantly outperforms  
20 ordinary SGD and carefully-tuned spectral norm regularization, which is a closely  
21 related regularization technique. We also show that TempBalance outperforms a  
22 number of state-of-the-art optimizers and learning rate schedulers.

## 23 1 Introduction

24 Having a learning rate schedule that gradually decreases over time is crucial for the convergence and  
25 performance of state-of-the-art machine learning algorithms. Indeed, many optimization algorithms  
26 essentially boil down to designing the progression of parameter updates, as realized by different  
27 learning rate schedules [1–4]. Common schedules assign a global model learning rate per epoch,  
28 where the same learning rate is used for all layers in the model. This includes the family of cyclical  
29 learning rates [3], and parameter-wise learning rate schedules like Adam [2] and its variants [5, 6].  
30 However, such a global learning rate schedule does not take into account the structural characteristics  
31 of neural networks. At the same time, parameter-wise learning rate schedules have long been  
32 conjectured to have worse generalization performance than carefully tuned SGD optimizers [7], and  
33 storing both first and second-order moments for each parameter can lead to *significantly* increased  
34 memory consumption [8]. As mentioned in Smith et al. [9], storing the whole Megatron-Turing  
35 NLG requires 10 terabytes of aggregate memory, and the Adam optimizer’s first and second-order  
36 moments [2] consume 40% of it. Nonetheless, improving parameter-wise learning rate schedules is  
37 an active field of study [4, 5, 10, 11].

38 A largely under-explored idea is to assign layer-wise learning rates to reconcile the two extremes of  
 39 setting a single global learning rate or assigning fine-grained parameter-level learning rates. This  
 40 learning rate assignment method does not require much storage cost and can assign different training  
 41 speeds to different layers. However, existing layer-wise schemes are often introduced as an additional  
 42 part of hyperparameter sweeping, and most lack a strong theoretical foundation. For instance, layer-  
 43 wise learning rates can increase test accuracy in transfer learning [12] and domain adaptation [13], but  
 44 these learning rates are often empirically tuned. More recently, Chen et al. [14], motivated by the idea  
 45 that lower-level layers are domain-specific and higher-level layers are task-specific, automates the  
 46 search for the optimal set of learning rates. However, the authors find the nested, bi-level optimization  
 47 scheme to be too computationally expensive in practice [15]. AutoLR also automatically tunes its  
 48 layer-wise learning rates according to the “role” of each layer [16]. The method is validated almost  
 49 entirely by empirical results, further explained by layer-wise weight variations. While the authors  
 50 attempt to assign a different starting learning rate to each layer, the learning rate for each layer  
 51 continues to stay largely constant throughout training. LARS [17, 18] is another method to assign  
 52 layer-wise learning rate. It is based on the “trust ratio”, defined as the ratio of weight norm to gradient  
 53 update norm of each layer, and it is specifically used in large batch training to avoid the gradient  
 54 diverge.

55 In this paper, we propose TempBalance, a simple yet effective layer-wise learning rate assignment  
 56 regularization method. TempBalance adopts a *statistical physics viewpoint* of learning and opti-  
 57 mization [19–22], and it views the learning rate as a “temperature parameter”, which refers to some  
 58 quantity related to the empirical noise/stochasticity introduced in the learning process. From this  
 59 viewpoint, what matters in SGD training is the noise scale (which is the same as the noise scale  
 60 mentioned in Smith and Le [23], Smith et al. [24] that can be written as a function of learning rate,  
 61 batch size and momentum), instead of the learning rate *per se*. Further, a series of recent papers [25–  
 62 27] point out that neural network training can be viewed as a balance between *temperature-like*  
 63 *parameters* and *load-like parameters*, where a load-like parameter refers to some quantity related  
 64 to the quantity/quality of the data, relative to the size of the model. For instance, Yang et al. [26]  
 65 vary load and temperature parameters to provide a comprehensive taxonomy of NN loss landscapes,  
 66 showing sharp phase transitions between different types of loss landscapes. In this paper, building  
 67 on this line of research on temperature parameters, we further use ideas and measurements from  
 68 Heavy-Tail Self Regularization (HT-SR) Theory [25, 28–31] to characterize the quality of each layer  
 69 and then assign *layer-wise temperature* (i.e., layer-wise learning rate) based on their heavy-tail (HT)  
 70 characterizations. We discuss the significance of HT-SR and its connection to layer-wise temperature  
 71 in the following paragraph.

72 **HT-SR theory.** HT-SR theory suggests that as layer weight matrices train for a longer period,  
 73 they start to show strong correlations, resulting in the HT structure of the Empirical Spectrum  
 74 Density (ESD) for each layer. To obtain this ESD, we take a neural network with  $L$  layers and its  
 75 corresponding weight matrices  $\mathbf{W}_1, \mathbf{W}_2, \dots, \mathbf{W}_L$  with shape  $n \times m$  (where  $n \geq m$ ). For the  $i$ -th  
 76 layer, we calculate the eigenvalues of its correlation matrix  $\mathbf{X}_i = \mathbf{W}_i^T \mathbf{W}_i$  and then plot the ESD for  
 77 that layer, which, upon training, will gradually change to have an HT structure by HT-SR theory [28].  
 78 We can then fit a power-law (PL) distribution to the HT part of the ESD, and extract its exponent as,  
 79 namely, PL\_Alpha. The fitted PL will have the following formula:

$$p(\lambda) \propto \lambda^{-\alpha}, \quad \lambda_{min} < \lambda < \lambda_{max}. \quad (1)$$

80 The PL\_Alpha metric measures the PL exponent of the weight matrices’ ESD, and its underlying  
 81 motivation stems from random matrix theory and statistical physics [29, 30].

82 The PL\_Alpha metric has been shown to predict the trends in the test accuracy of state-of-the-art  
 83 computer vision (CV) and natural language processing (NLP) neural networks, without even the need  
 84 for access to training or testing data [30, 32]. According to Martin et al. [30], one can aggregate  
 85 PL\_Alpha’s for different layers either by simple averaging or weighted averaging, and they can all  
 86 predict test accuracy in different cases [30, 32]. Furthermore, the *layer-wise* nature of PL\_Alpha  
 87 makes it a fine-grained metric that can be used to assess the quality of individual layers of the network.  
 88 Thus, in this paper, we extend HT-SR to training, and we exploit the layer-wise information provided  
 89 by PL\_Alpha to determine the layer-wise learning rates for better test accuracy.

90 Note that PL\_Alpha is not the only way to measure the HT structure. Several recent papers [33–35]  
 91 use different HT metrics to measure the spectral of several “important matrices” (such as input/output  
 92 covariance matrices, Fisher Information Matrices and Hessian), and we show in Appendix A that

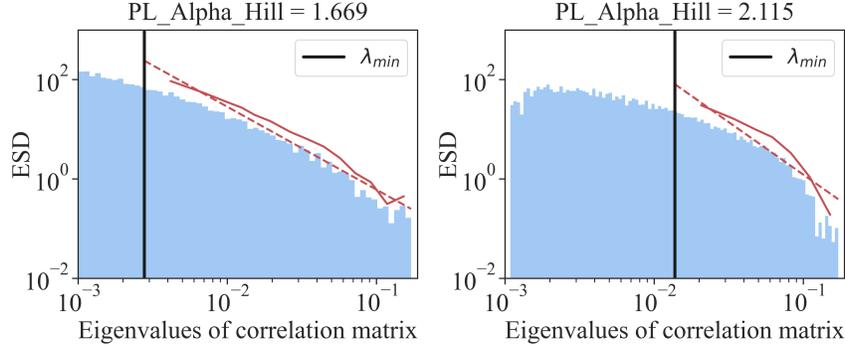


Figure 1: Examples of power-law (PL) fitting using the Hill estimator. Blue histograms depict the empirical spectral densities (ESDs). Vertical black lines indicate the lower threshold  $\lambda_{min}$  used to truncate the full ESDs and extract the tail portion. Solid red curves represent the tail part of the ESDs truncated by  $\lambda_{min}$ , while dashed red curves represent the fitted heavy-tailed (HT) distributions. The left shows a more HT ESD, requiring a relatively lower learning rate. The right one shows less HT ESD, which requires a relatively higher learning rate. Unlike prior work, we do not aim to find the “optimal” PL exponent. Instead, we use the PL exponent to *rank* ESDs to find layers that need higher/lower learning rates. These two ESDs correspond to two layers of a ResNet18 model trained on TinyImageNet.

93 these HT phenomena, measured in different ways on different matrices, are closely related to each  
 94 other. On the other hand, this also means that the “absolute value” of PL\_Alpha is unimportant,  
 95 as optimal PL exponents estimated by different algorithms can be different [30, 33]. It turns out  
 96 that what matters the most, as we show in this paper, is the layer-wise quality *ranked* by the PL  
 97 exponent: layers with a smaller PL\_Alpha to be relatively more “overtrained” and those with a larger  
 98 PL\_Alpha to be relatively more “undertrained”. This observation leads to a simple and efficient way  
 99 to *balance* layer-wise learning rates: assigning a lower learning rate to overtrained layers and a larger  
 100 learning rate to undertrained layers using PL\_Alpha (see Figure 1). On top of this, we can grid-search  
 101 the *global* learning rate on validation data, which is standard practice and is more efficient than  
 102 grid-searching the layer-wise learning rates. We use this combination of assigning layer-wise learning  
 103 rates using PL\_Alpha and grid-searching the base learning rate to avoid deciding the “optimal PL  
 104 exponent”, which can be tricky due to different ways of measuring HT. Indeed, there are different  
 105 ways to measure PL\_Alpha [28], and we use the Hill estimator [36], which shows stable performance  
 106 in our experiments. Thus, we call our version of PL\_Alpha the PL\_Alpha\_Hill metric, and we use  
 107 that for the remaining paper. Further, we use a *scale-free* way to map the estimated PL\_Alpha\_Hill  
 108 to the learning rate, meaning that arbitrary linear scaling on the estimated PL\_Alpha\_Hill (either  
 109 due to the choice of the estimator or noisy measurements) does not affect the assigned learning rates.

110 Another common way to change the ESD of weights is to constrain the spectral norm (i.e., the  
 111 largest eigenvalue) using spectral norm regularization (SNR) [37, 38]. SNR provides a different  
 112 form of regularization compared to HT-SR because it regulates the largest eigenvalue instead of  
 113 the ESD slope (i.e., the PL\_Alpha\_Hill metric). It has been demonstrated that spectral norm  
 114 and PL\_Alpha\_Hill serve distinct roles in generalization, and their combined form yields optimal  
 115 predictions for test accuracy trends [28, 30–32]. Our paper, on the other hand, demonstrates that  
 116 TempBalance outperforms SNR in training deep neural networks in most cases. Moreover, when  
 117 these two regularization methods are combined during training, they result in optimal test accuracy,  
 118 thereby confirming their complementary roles. As mentioned by Martin and Mahoney [31], Yang  
 119 et al. [32], spectral norm and PL\_Alpha\_Hill measure the scale and the shape of a ESD respectively,  
 120 and regulating both the scale and shape is crucial for achieving better ESD regularization. We  
 121 also provide ablation studies on several layer-wise metrics for assigning layer-wise learning rates,  
 122 including spectral norm, and we show that PL\_Alpha\_Hill performs the best among them.

123 **Contributions.** The following summarizes our main contributions:

- 124 • We propose a simple yet effective layer-wise learning rate schedule TempBalance based on
- 125 HT-SR theory. We empirically found two insights. First, the mapping from PL\_Alpha\_Hill
- 126 to learning rates should be scale-free, meaning that arbitrary linear scaling on the estimated
- 127 PL exponent should not change the learning rate assignment. Second, searching for the

128 minimum eigenvalue  $\lambda_{\min}$ , a standard practice in PL fitting [28, 39, 40], leads to unstable  
129 training. We instead fix  $\lambda_{\min}$  as the medium of the ESD.

- 130 • We compare TempBalance to ordinary stochastic gradient descent (SGD) and SNR on various  
131 training tasks. This includes (1) different network architectures, such as ResNet, VGG,  
132 WideResNet, (2) different datasets, such as CIFAR10, CIFAR100, SVHN, TinyImageNet,  
133 and (3) ablation studies, such as varying widths, depths, initial learning rates and HT-SR  
134 layer-wise metrics. Compared to ordinary SGD, TempBalance achieves higher test accuracy  
135 by setting layer-wise learning rates. Compared to SNR, TempBalance performs better by  
136 providing a more fine-grained regularization on *shape/slope* instead of norm. We also show  
137 that combining TempBalance and SNR leads to further improved accuracy, verifying their  
138 complementary roles in informing deep learning training.
- 139 • We compare TempBalance to a range of state-of-the-art optimizers and learning rate sched-  
140 ulers, including SGDR [10], SGDP [41], Lookahead [42] and LARS [17, 18] on ResNet18  
141 and ResNet34 trained on CIFAR100. We show that TempBalance achieves the highest test  
142 accuracy. We do careful hyperparameter tuning for all baselines. All results are obtained  
143 from five random seeds.
- 144 • We use ablation studies to show that PL\_Alpha\_Hill provides the best test accuracy among  
145 a few layer-wise metrics considered in HT-SR [30, 32]. We also show that TempBalance  
146 maintains stable performance over SGD baselines when the model size changes. Furthermore,  
147 we show visualization results in Appendix B, verifying that TempBalance controls ESDs  
148 during training.

## 149 2 Related Work

150 Here we give an overview of the statistical mechanics of learning and recent progress in theoretical  
151 and empirical studies on generalization metrics and their applications.

### 152 2.1 Statistical mechanics of learning and HT-SR

153 Our paper is motivated by statistical mechanics of learning [43–45], and especially by works that  
154 connect load-like [43, 46, 47] and temperature-like parameters [19, 48] to neural networks. According  
155 to prior works in this area [25, 26], a temperature-like parameter represents the amount of noise and  
156 disturbance in the iteration of SGD, such as learning rate, weight decay parameter, and batch size. A  
157 load-like parameter represents the quantity and/or quality of data relative to the size of the learning  
158 model. To measure the quality of trained neural networks, Martin and Mahoney [28] introduce HT-SR  
159 theory, showing that the weight matrices of deep neural networks exhibit heavy-tailed empirical  
160 spectral densities. In subsequent papers, HT-SR has been applied to predicting trends in test accuracy  
161 of large-scale neural networks in both CV and NLP [30–32], but it has yet to be systematically  
162 incorporated to novel training algorithms. Recently, more and more papers realize the important  
163 connections between deep neural networks and statistical mechanics of learning. To name a few, Yang  
164 et al. [26] use load and temperature parameters to study a wide range of loss landscapes, providing a  
165 taxonomy from the perspective of *global structure* of a loss landscape. On the theory side, Baity-Jesi  
166 et al. [49] investigates the glassy behavior of neural networks, and Barbier et al. [50] derives the  
167 optimal generalization error of generalized linear systems. More recently, Sorscher et al. [51] studies  
168 easy versus hard samples used in training and design a “data-pruning” method. Zhou et al. [27]  
169 establishes a “three-regime model” in network pruning, unifying multiple practical hyperparameter  
170 tuning methods in a principled way.

### 171 2.2 Generalization measures

172 Note that the search for effective and robust generalization metrics has been the focus of several recent  
173 theoretical and empirical works [26, 30, 32, 52–54]. Several recent papers apply metric-informed  
174 training and architecture search, such as those based on Hessian [4, 55–57], spectral norm [37, 38],  
175 stable rank [58] and the spectrum of neural tangent kernel [59]. However, the most empirically  
176 successful generalization metrics, such as those based on the PAC-Bayes bounds [60–63], do not  
177 straightforwardly transfer to layer-wise quality metrics because such generalization metrics often  
178 study the whole neural network as an architecture-free function and lack the fine granularity to unveil  
179 the quality of each layer. Also, it has been mentioned in the literature [52] that (1) directly regularizing



197 We provide the details of TempBalance in Algorithm 1. Based on PL\_Alpha\_Hill in different  
 198 layers, we use the learning rate schedule function  $f_t$  to map the  $i_{th}$  layer to a particular learning rate  
 199  $f_t(i)$  in epoch  $t$ . We adopt  $f_t$  as a linear map between the layer-wise PL\_Alpha\_Hill and the final  
 200 layer-wise learning rate, which has the following formula:

$$f_t(i) = \eta_t \cdot \left[ \frac{\alpha_t^i - \alpha_t^{min}}{\alpha_t^{max} - \alpha_t^{min}}(s_2 - s_1) + s_1 \right], \quad (2)$$

201 where  $\eta_t$  means the base global learning rate in epoch  $t$ ,  $(s_1, s_2)$  are the minimum and maximum  
 202 learning rate scaling ratio relative to  $\eta_t$ ,  $\alpha_t^i$  represents the layer  $i$ 's PL\_Alpha\_Hill at the beginning  
 203 of epoch  $t$ , and  $(\alpha_t^{min}, \alpha_t^{max})$  denote the minimum and maximum PL\_Alpha\_Hill across all the  
 204 layers in epoch  $t$ . Using (2), we ensure that the new learning rate  $f_t(i)$  is a scaled version of the  
 205 original base learning rate  $\eta_t$  and is always inside the interval  $[s_1\eta_t, s_2\eta_t]$ . We only consider  $(s_1, s_2)$   
 206 such that  $\frac{s_1+s_2}{2} = 1$ , e.g., (0.5, 1.5) or (0.8, 1.2).

207 To fit the power law distribution  $p(\lambda)$  defined in (1), we use the famous Hill estimator [36] [64]. For  
 208 the  $i$ -th layer, suppose the weight matrix is  $\mathbf{W}_i$  and the correlation matrix  $\mathbf{W}_i^T \mathbf{W}_i$  has ascending  
 209 eigenvalues  $\{\lambda_i\}_{i=1}^n$ . Then, the Hill estimator calculates PL\_Alpha\_Hill using the following:

$$\text{PL\_Alpha\_Hill} = 1 + \frac{k}{\left(\sum_{i=1}^k \ln \frac{\lambda_{n-i+1}}{\lambda_{n-k}}\right)}, \quad (3)$$

210 where  $k$  is the adjustable parameter, and we adopt  $k = \frac{n}{2}$  in our experiments. Note that changing  $k$   
 211 essentially changes the lower eigenvalue threshold  $\lambda_{min}$  for (truncated) PL estimation, as shown by  
 212 the vertical black line in Figure 1. Choosing  $k = \frac{n}{2}$  means using the largest half of the eigenvalues to  
 213 estimate the slope. We empirically find that fixing  $k$  for all layers leads to more stable performance  
 214 than searching  $k$  for different layers (e.g., optimizing  $k$  using the Kolmogorov–Smirnov test [40].)

215 One advantage of mapping PL\_Alpha\_Hill to learning rates using (2) is that the scale of  
 216 PL\_Alpha\_Hill is unimportant, i.e., linearly scaling PL\_Alpha\_Hill arbitrarily does not change  
 217 the learning rate assignment because the linear scaling cancels each other in (2). This can maximally  
 218 reduce the artifact of estimating the ESD PL exponent/slope due to estimation noise, which has been  
 219 found to be a tricky issue in practice [28, 31].

## 220 4 Empirical results

221 In this section, we first give full details of the experimental setup (Section 4.1) and compare our  
 222 method TempBalance to a few baselines (Section 4.2). Then, in Section 4.3, we perform ablation  
 223 studies on varied initial learning rates, model widths and HT-SR layer-wise metrics.

### 224 4.1 Experimental setup

225 **Datasets.** We consider CIFAR100, CIFAR10, SVHN and Tiny ImageNet (TIN) [65–69]. CIFAR100  
 226 consists of 50000 pictures for training and 10000 pictures for testing with 100 categories. CIFAR10  
 227 consists of 50000 pictures for training and 10000 pictures for testing with 10 categories. SVHN  
 228 consists of 73257 pictures for training and 26032 pictures for testing with 10 categories. Tiny  
 229 ImageNet consists of 100000 pictures for training and 10000 images for testing with 200 classes.

230 **Models.** We mainly consider three types of deep neural networks, VGG, ResNet and WideResNet  
 231 (WRN) [70–72]. For each network, we consider two different size options. For VGG, we consider  
 232 VGG16 and VGG19. For ResNet, we consider ResNet18 and ResNet34. For WideResNet, we  
 233 consider WRN16-8 and WRN28-6. Also, for ResNet and VGG, we consider three different widths  
 234 for ablation studies.

235 **Hyperparameters.** One baseline is ordinary SGD training with a cosine annealing learning rate  
 236 schedule (CAL), which follows the formula:  $\eta_t = \frac{\eta_0}{2} \left(1 + \cos\left(\frac{t\pi}{T}\right)\right)$ , where  $t$  is the current epoch,  
 237  $T$  represents the total training epochs and  $\eta_0$  is the initial learning rate. We grid-search the optimal  
 238 initial (base) learning rate  $\eta_0$  for the CAL baseline, using the grid  $\{0.05, 0.1, 0.15\}$  for ResNet and  
 239  $\{0.025, 0.05, 0.1\}$  for VGG. The momentum and weight decay are 0.9 and  $5 \times 10^{-4}$ , respectively,  
 240 which are both standard choices.

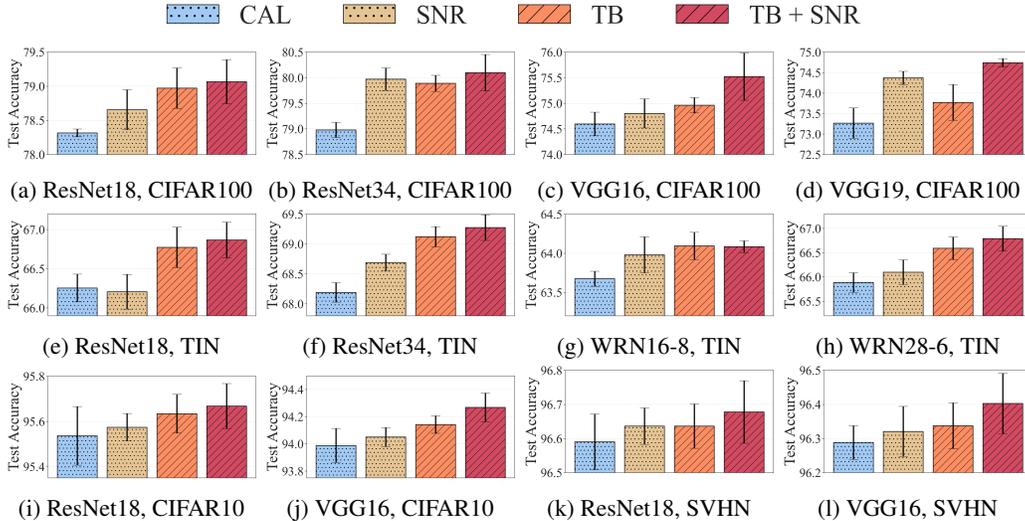


Figure 3: **(Main result)**. Comparing our method TempBalance (TB) to CAL and SNR. Our method TempBalance outperforms CAL and SNR in almost all the settings except for VGG19 and ResNet 34 on CIFAR 100. For all experiments, combining TempBalance and SNR (TB+SNR) yields the best performance. All baselines are carefully tuned. All results are obtained by running five random seeds. See Appendix C for the details in all hyperparameters.

241 Another baseline is called spectral norm regularization (SNR). Prior work uses the following SNR  
 242 objective function [37]:

$$\min_{\Theta} \frac{1}{n} \sum_{i=1}^n l(f_{\Theta}(\mathbf{x}_i), \mathbf{y}_i) + \frac{\lambda_{sr}}{2} \sum_{l=1}^L \sigma(W_l)^2, \quad (4)$$

243 where  $\lambda_{sr}$  is the spectral norm regularization coefficient,  $\sigma(W_l)$  is the largest eigenvalue, i.e, spectrum  
 244 norm of weight matrix  $\mathbf{W}_l$ , and  $L$  is the number of layers. We use the power iteration method to  
 245 calculate  $\sigma(W_l)$  in our experiments. For SNR, we grid-search the optimal regularization coefficient  
 246  $\lambda_{sr}$ , and we again adopt the CAL schedule for SNR, similar to the CAL baseline.

247 To make our results fully reproducible, we report in Appendix C all hyperparameters, random seeds,  
 248 and all numerical values of experimental results shown in the figures.

## 249 4.2 Comparing TempBalance and multiple baseline methods.

250 First, we compare TempBalance to two baseline training methods. See results in Figure 3. In the  
 251 figure, CAL means SGD training with a CAL learning rate schedule, SNR means SGD trained with spectral  
 252 norm regularization. TB means our method TempBalance, and TB + SNR means TempBalance  
 253 combined with SNR. All error bars are obtained from five random seeds. From Figure 3, we see that  
 254 TempBalance outperforms the CAL baseline in all settings. In almost all cases, it performs better than  
 255 SNR baseline. When TempBalance does not outperform SNR, combining SNR with TempBalance  
 256 leads to better test accuracy.

257 Second, we compare our method to a number of optimizers and learning rate schedulers that are  
 258 not necessarily related to ESD of weights. These include SGDR [10], SGDP [41], Lookahead [42]  
 259 and LARS [17, 18], and we compare these baselines with TempBalance for ResNet18 and ResNet34  
 260 trained on CIFAR100. SGDR is stochastic gradient descent with warm restarts. SGDP modifies the  
 261 ordinary SGD to compensate for the effect of increasing weight norm. Lookahead [42] modifies SGD  
 262 by letting each gradient update approximate the future trajectory of multiple updates. LARS assigns  
 263 layer-wise learning rates based on the so-called “trust-ratio” and is the closest to our method. Results  
 264 in Figure 4 show that TempBalance outperforms these baselines, and TempBalance combined with  
 265 SGDP is the best-performing method. The crosses on each column represent training runs with  
 266 different hyperparameters. Note that there are several other methods based on modifying the Adam

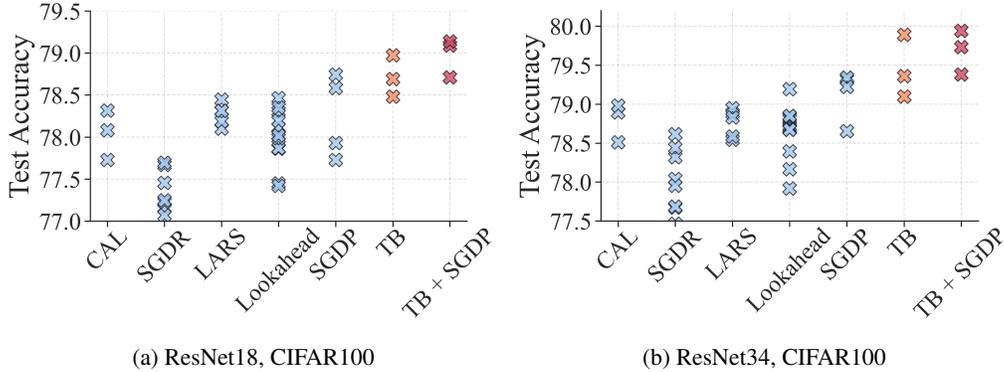


Figure 4: **(More baseline optimizers)**. Comparing our method TempBalance (TB) to cosine annealing (CAL) baseline and other state-of-the-art optimizers and learning rate schedulers for ResNet18 and ResNet34 trained on CIFAR100. Crosses for the same method represent different hyperparameter settings. Each cross represents the mean test accuracy of five random seeds. The best performing model thus far is TB combined with SGDP.

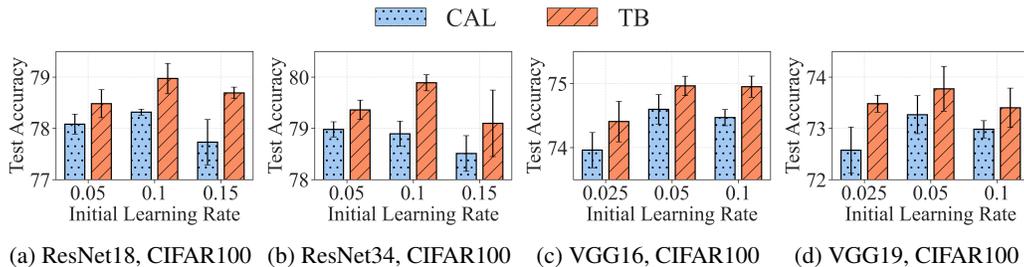


Figure 5: **(Tuning initial learning rate)**. Comparing the test accuracy of TempBalance (red) and CAL baseline (blue) for varying initial learning rate. Our method TempBalance outperforms CAL for both ResNet and VGG trained on CIFAR100. All results are obtained by running five random seeds.

267 optimizer [2], such as AdamW [11], AdamP [41] and LAMB [73]. However, we do not find them to  
 268 provide better results than the SGD baseline with cosine annealing (CAL in Figure 4).

### 269 4.3 Corroborating results and ablation studies.

270 In addition to the main results shown in Figure 3 and Figure 4, we show three ablation studies.

271 **Experiment one: tuning initial learning rate  $\eta_0$ .** We train models from scratch using TempBalance  
 272 vs. CAL with various initial learning rates. We intend on comparing TempBalance and CAL baseline  
 273 when both methods are allowed to search for the optimal hyperparameters. We again use ResNet18,  
 274 ResNet34, VGG16 and VGG19 as our architectures and show results on CIFAR100. From the results  
 275 in Figure 5, TempBalance achieves a higher test accuracy than CAL for both ResNet and VGG.

276 **Experiment two: varying channel width.** We view the fraction of model width in Experiment  
 277 one as “100%” and experiment with models with varied widths in [50%, 100%, 150%]. We again  
 278 use VGG16, VGG19, ResNet18, and ResNet 34 trained on CIFAR100, and we grid search for the  
 279 optimal learning rate for each width to get best accuracy. From Figure 6, we find that TempBalance  
 280 outperforms the baseline for all widths.

281 **Experiment three: varying HT-SR metric.** We use different HT-SR metrics to assign layer-wise  
 282 learning rates. That is, we replace the layer-wise PL\_Alpha\_Hill in (2) with other HT-SR metrics  
 283 including SpectralNorm and AlphaWeighted [30]. Results in Figure 7 show that PL\_Alpha\_Hill  
 284 achieves the optimal test accuracy.

285 **Visualization results.** We further analyze our methods by illustrating the effect of TempBalance on  
 286 regularizing ESDs. See Appendix B.

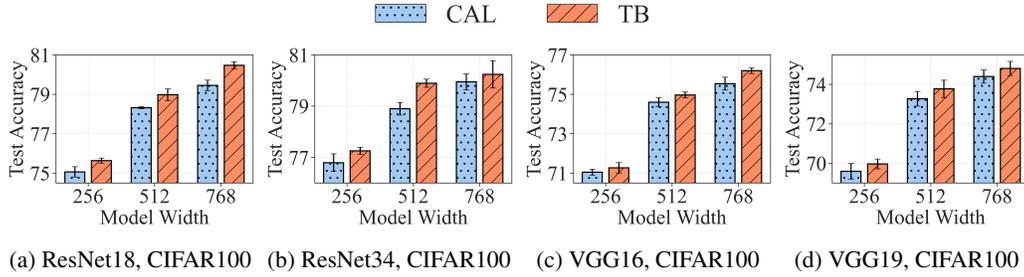


Figure 6: **(Different widths)**. Comparing TempBalance and the CAL baseline for different network widths. All results are obtained by running five random seeds.

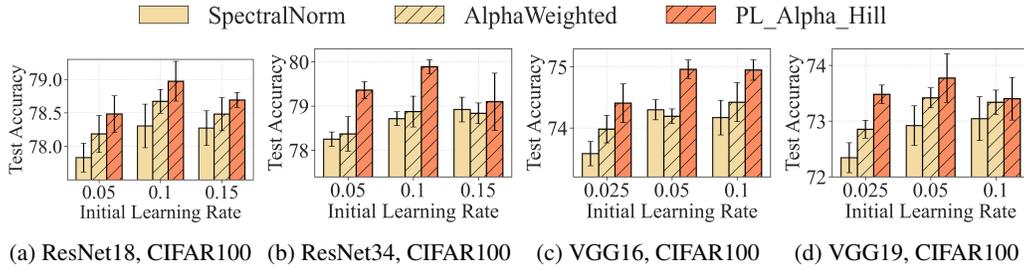


Figure 7: **(Different HT-SR metrics)**. Comparing PL\_Alpha\_Hill with multiple HT-SR metrics. PL\_Alpha\_Hill achieves the best test accuracy among these metrics. All results are obtained by running five random seeds.

## 287 5 Conclusion

288 Our extensive empirical evaluations demonstrate that TempBalance offers a straightforward yet  
 289 effective layer-wise learning rate schedule. Furthermore, our approach for balancing layer-wise  
 290 temperature confirms the following: (i) HT-SR-motivated metric PL\_Alpha\_Hill helps layers  
 291 achieve maximal temperature balance during training, exhibits strong correlations with model quality,  
 292 and yields improved performance during testing. (ii) Temperature balancing is a novel and essential  
 293 aspect of neural network training, and HT-SR theory provides a strong theoretical support for  
 294 balancing temperatures. (iii) Layer-wise learning rate schedules are cheap and effective to apply  
 295 when compared with per-parameter learning rate schedules, and it is useful to study these layer-wise  
 296 learning rate schedules further. Our method provides insights into the study of layer-wise tuning  
 297 approaches and load-temperature balancing in deep neural network training, as it serves both as a  
 298 layer-wise learning rate schedule and an effective regularization technique based on HT-SR metrics.

299 **Limitations and societal impacts.** Our paper leaves many future directions to explore, of which we  
 300 discuss a few below:

- 301 • Can HT-SR metrics be extended to parameter-wise learning rate schedules, global learning  
 302 rate schedules, or other hyperparameters? It would be great to observe how HT-SR can  
 303 assist in acquiring a comprehensive set of hyperparameter tuning tools.
- 304 • Is it possible to accelerate the computation of ESDs and PL\_Alpha\_Hill to achieve a  
 305 more adaptive learning rate scheduler? Currently, we calculate layer-wise PL\_Alpha\_Hill  
 306 once per epoch, resulting in a minimal increase in computational complexity. Consider the  
 307 example of training ResNet18 for 200 epochs on CIFAR100 or TinyImageNet. Calculating  
 308 layer-wise PL\_Alpha\_Hill takes 1.4 seconds for each epoch, leading to 4.6 minutes in total.  
 309 Training CIFAR100 on 1 V100 takes 80 minutes, and thus using TB increases 6% of training  
 310 time. Training TinyImageNet on 2 V100 takes 240 minutes, and thus using TB increase 2%  
 311 of training time. However, if we can significantly decrease the expense of computing ESDs,  
 312 it might enable an optimizer that adjusts the learning rate every few gradient updates.

313 Our research centers around developing a generic algorithm for optimizing neural networks. Although  
 314 it can be applied to learning models with adverse applications, we do not see any immediate negative  
 315 societal impacts stemming from the algorithm itself.

316 **References**

- 317 [1] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning  
318 and stochastic optimization. *Journal of Machine Learning Research*, 12(61):2121–2159, 2011.
- 319 [2] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International  
320 Conference on Learning Representations*, 2014.
- 321 [3] Leslie N Smith. Cyclical learning rates for training neural networks. In *2017 IEEE winter  
322 conference on applications of computer vision (WACV)*, pages 464–472, 2017.
- 323 [4] Zhewei Yao, Amir Gholami, Sheng Shen, Mustafa Mustafa, Kurt Keutzer, and Michael Mahoney.  
324 Adahessian: An adaptive second order optimizer for machine learning. In *Proceedings of the  
325 AAAI Conference on Artificial Intelligence*, 2021.
- 326 [5] Juntang Zhuang, Tommy Tang, Yifan Ding, Sekhar C Tatikonda, Nicha Dvornek, Xenophon  
327 Papademetris, and James Duncan. Adabelief optimizer: Adapting stepsizes by the belief in  
328 observed gradients. *Advances in neural information processing systems*, 2020.
- 329 [6] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and  
330 Jiawei Han. On the variance of the adaptive learning rate and beyond. In *International  
331 Conference on Learning Representations*, 2020.
- 332 [7] Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The  
333 marginal value of adaptive gradient methods in machine learning. In *Advances in Neural  
334 Information Processing Systems*, 2017.
- 335 [8] Bharat Singh, Soham De, Yangmuzi Zhang, Thomas Goldstein, and Gavin Taylor. Layer-  
336 specific adaptive learning rates for deep networks. In *IEEE 14th International Conference on  
337 Machine Learning and Applications*, 2015.
- 338 [9] Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari,  
339 Jared Casper, Zhun Liu, Shrimai Prabhunoye, George Zerveas, Vijay Korthikanti, et al. Using  
340 deepspeed and megatron to train megatron-turing nlg 530b, a large-scale generative language  
341 model. *arXiv preprint arXiv:2201.11990*, 2022.
- 342 [10] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In  
343 *International Conference on Learning Representations*, 2017.
- 344 [11] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International  
345 Conference on Learning Representations*, 2019.
- 346 [12] Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classi-  
347 fication. In *Proceedings of the 56th Annual Meeting of the Association for Computational  
348 Linguistics*.
- 349 [13] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features  
350 with deep adaptation networks. In *International Conference on Machine Learning*, 2015.
- 351 [14] Yixiong Chen, Jingxian Li, Hua Jiang, Li Liu, and Chris Ding. Metalr: Layer-wise learning rate  
352 based on meta-learning for adaptively fine-tuning medical pre-trained models. *arXiv preprint  
353 arXiv:2206.01408*, 2022.
- 354 [15] Luca Franceschi, Paolo Frasconi, Saverio Salzo, Riccardo Grazi, and Massimiliano Pontil.  
355 Bilevel programming for hyperparameter optimization and meta-learning. In *International  
356 Conference on Machine Learning*, 2018.
- 357 [16] Youngmin Ro and Jin Young Choi. Autolr: Layer-wise pruning and auto-tuning of learning  
358 rates in fine-tuning of deep networks. In *Proceedings of the AAAI Conference on Artificial  
359 Intelligence*, 2021.
- 360 [17] Yang You, Igor Gitman, and Boris Ginsburg. Scaling SGD batch size to 32k for ImageNet  
361 training. *arXiv preprint arXiv:1708.03888*, 6(12):6, 2017.

- 362 [18] Yang You, Zhao Zhang, Cho-Jui Hsieh, James Demmel, and Kurt Keutzer. Imagenet training in  
363 minutes. In *Proceedings of the 47th International Conference on Parallel Processing*, 2018.
- 364 [19] H. S. Seung, H. Sompolinsky, and N. Tishby. Statistical mechanics of learning from examples.  
365 *Physical Review A*, 45(8):6056–6091, 1992.
- 366 [20] T. L. H. Watkin, A. Rau, and M. Biehl. The statistical mechanics of learning a rule. *Rev. Mod. Phys.*,  
367 65(2):499–556, 1993.
- 368 [21] D. Haussler, M. Kearns, H. S. Seung, and N. Tishby. Rigorous learning curve bounds from  
369 statistical mechanics. *Machine Learning*, 25(2):195–236, 1996.
- 370 [22] Andreas Engel and Christian P. L. Van den Broeck. *Statistical mechanics of learning*. Cambridge  
371 University Press, New York, NY, USA, 2001.
- 372 [23] Samuel L Smith and Quoc V Le. A bayesian perspective on generalization and stochastic  
373 gradient descent. In *International Conference on Learning Representations*, 2018.
- 374 [24] Samuel L. Smith, Pieter-Jan Kindermans, and Quoc V. Le. Don’t decay the learning rate,  
375 increase the batch size. In *International Conference on Learning Representations*, 2018.
- 376 [25] Charles H Martin and Michael W Mahoney. Rethinking generalization requires revisiting  
377 old ideas: statistical mechanics approaches and complex learning behavior. Technical Report  
378 Preprint: arXiv:1710.09553, 2017.
- 379 [26] Yaoqing Yang, Liam Hodgkinson, Ryan Theisen, Joe Zou, Joseph E Gonzalez, Kannan Ram-  
380 chandran, and Michael W Mahoney. Taxonomizing local versus global structure in neural  
381 network loss landscapes. In *Advances in Neural Information Processing Systems*, 2021.
- 382 [27] Yefan Zhou, Yaoqing Yang, Arin Chang, and Mahoney W Michael. A three-regime model of  
383 network pruning. In *International Conference on Machine Learning*, 2023.
- 384 [28] Charles H Martin and Michael W Mahoney. Implicit self-regularization in deep neural networks:  
385 Evidence from random matrix theory and implications for learning. *Journal of Machine*  
386 *Learning Research*, 22(165):1–73, 2021.
- 387 [29] Charles H Martin and Michael W Mahoney. Traditional and heavy tailed self regularization in  
388 neural network models. In *International Conference on Machine Learning*, 2019.
- 389 [30] Charles H Martin, Tongsu Serena Peng, and Michael W Mahoney. Predicting trends in the  
390 quality of state-of-the-art neural networks without access to training or testing data. *Nature*  
391 *Communications*, 12(1):1–13, 2021.
- 392 [31] Charles H Martin and Michael W Mahoney. Post-mortem on a deep learning contest: a  
393 Simpson’s paradox and the complementary roles of scale metrics versus shape metrics. Technical  
394 Report Preprint: arXiv:2106.00734, 2021.
- 395 [32] Yaoqing Yang, Ryan Theisen, Liam Hodgkinson, Joseph E Gonzalez, Kannan Ramchandran,  
396 Charles H Martin, and Michael W Mahoney. Evaluating natural language processing models  
397 with generalization metrics that do not need access to any training or testing data. In *proceedings*  
398 *of the 29th SIGKDD Conference on Knowledge Discovery and Data Mining*, 2023.
- 399 [33] Kumar K Agrawal, Arnab Kumar Mondal, Arna Ghosh, and Blake Richards.  $\alpha$ -req :  
400 Assessing representation quality in self-supervised learning by measuring eigenspectrum decay.  
401 In *Advances in Neural Information Processing Systems*, 2022.
- 402 [34] Josue Nassar, Piotr Sokol, SueYeon Chung, Kenneth D Harris, and Il Memming Park. On  $1/n$   
403 neural representation and robustness. 2020.
- 404 [35] Zeke Xie, Qian-Yuan Tang, Yunfeng Cai, Mingming Sun, and Ping Li. On the power-law  
405 hessian spectrums in deep learning. *arXiv preprint arXiv:2201.13011*, 2022.
- 406 [36] Bruce M Hill. A simple general approach to inference about the tail of a distribution. *The*  
407 *annals of statistics*, pages 1163–1174, 1975.

- 408 [37] Yuichi Yoshida and Takeru Miyato. Spectral norm regularization for improving the generaliz-  
409 ability of deep learning. *arXiv preprint arXiv:1705.10941*, 2017.
- 410 [38] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization  
411 for generative adversarial networks. In *International Conference on Learning Representations*,  
412 2018.
- 413 [39] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in  
414 empirical data. *SIAM review*, 51(4):661–703, 2009.
- 415 [40] Jeff Alstott, Ed Bullmore, and Dietmar Plenz. Powerlaw: a python package for analysis of  
416 heavy-tailed distributions. *PloS one*, 9(1):e85777, 2014.
- 417 [41] Byeongho Heo, Sanghyuk Chun, Seong Joon Oh, Dongyoon Han, Sangdoon Yun, Gyuwan  
418 Kim, Youngjung Uh, and Jung-Woo Ha. Adamp: Slowing down the slowdown for momentum  
419 optimizers on scale-invariant weights. In *International Conference on Learning Representations*,  
420 2021.
- 421 [42] Michael Zhang, James Lucas, Jimmy Ba, and Geoffrey E Hinton. Lookahead optimizer: k steps  
422 forward, 1 step back. In *Advances in Neural Information Processing Systems*, 2019.
- 423 [43] John J Hopfield. Neural networks and physical systems with emergent collective computational  
424 abilities. *Proceedings of the national academy of sciences*, 1982.
- 425 [44] Haim Sompolinsky et al. Statistical mechanics of neural networks. *Physics Today*, 41(21):  
426 70–80, 1988.
- 427 [45] LM Rasdi Rere, Mohamad Ivan Fanany, and Aniati Murni Arymurthy. Simulated annealing  
428 algorithm for deep learning. *Procedia Computer Science*, 72:137–144, 2015.
- 429 [46] Adriano Barra and Francesco Guerra. About the ergodic regime in the analogical hopfield  
430 neural networks: moments of the partition function. *Journal of mathematical physics*, 49(12):  
431 125217, 2008.
- 432 [47] Adriano Barra, Alberto Bernacchia, Enrica Santucci, and Pierluigi Contucci. On the equivalence  
433 of hopfield networks and boltzmann machines. *Neural Networks*, 34:1–9, 2012.
- 434 [48] Stephen G. Brush. History of the lenz-ising model. *Reviews of Modern Physics*, 39:883–893,  
435 1967.
- 436 [49] Marco Baity-Jesi, Levent Sagun, Mario Geiger, Stefano Spigler, Gérard Ben Arous, Chiara  
437 Cammarota, Yann LeCun, Matthieu Wyart, and Giulio Biroli. Comparing dynamics: Deep  
438 neural networks versus glassy systems. In *International Conference on Machine Learning*,  
439 2018.
- 440 [50] Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, and Lenka Zdeborová. Optimal  
441 errors and phase transitions in high-dimensional generalized linear models. *Proceedings of the  
442 National Academy of Sciences*, 116(12):5451–5460, 2019.
- 443 [51] Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, and Ari Morcos. Beyond  
444 neural scaling laws: beating power law scaling via data pruning. In *Advances in Neural  
445 Information Processing Systems*, 2022.
- 446 [52] Yiding Jiang, Behnam Neyshabur, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic  
447 generalization measures and where to find them. In *International Conference on Learning  
448 Representations*, 2019.
- 449 [53] Gintare Karolina Dziugaite, Alexandre Drouin, Brady Neal, Nitarshan Rajkumar, Ethan Ca-  
450 ballero, Linbo Wang, Ioannis Mitliagkas, and Daniel M Roy. In search of robust measures of  
451 generalization. In *Advances in Neural Information Processing Systems*, 2020.
- 452 [54] Peter Bartlett, Dylan Foster, and Matus Telgarsky. Spectrally-normalized margin bounds for  
453 neural networks. In *Advances in Neural Information Processing Systems*, 2017.

- 454 [55] Huanrui Yang, Xiaoxuan Yang, Neil Zhenqiang Gong, and Yiran Chen. Hero: Hessian-enhanced  
455 robust optimization for unifying and improving generalization and quantization performance.  
456 In *Proceedings of the 59th ACM/IEEE Design Automation Conference*, 2022.
- 457 [56] Zhen Dong, Zhewei Yao, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. HAWQ: Hes-  
458 sian aware quantization of neural networks with mixed-precision. In *IEEE/CVF International*  
459 *Conference on Computer Vision*, 2019.
- 460 [57] Sheng Shen, Zhen Dong, Jiayu Ye, Linjian Ma, Zhewei Yao, Amir Gholami, Michael W  
461 Mahoney, and Kurt Keutzer. Q-BERT: Hessian based ultra low precision quantization of bert.  
462 In *AAAI Conference on Artificial Intelligence*, 2020.
- 463 [58] Amartya Sanyal, Philip H. Torr, and Puneet K. Dokania. Stable rank normalization for im-  
464 proved generalization in neural networks and gans. In *International Conference on Learning*  
465 *Representations*, 2020.
- 466 [59] Wuyang Chen, Xinyu Gong, and Zhangyang Wang. Neural architecture search on imagenet in  
467 four gpu hours: A theoretically inspired perspective. In *International Conference on Learning*  
468 *Representations*, 2021.
- 469 [60] David A McAllester. Some pac-bayesian theorems. In *Proceedings of the eleventh annual*  
470 *conference on Computational learning theory*, 1998.
- 471 [61] John Langford and John Shawe-Taylor. Pac-bayes & margins. In *Advances in Neural Informa-*  
472 *tion Processing Systems*, 2002.
- 473 [62] Gintare Karolina Dziugaite and Daniel M Roy. Computing nonvacuous generalization bounds  
474 for deep (stochastic) neural networks with many more parameters than training data. In *Annual*  
475 *Conference on Uncertainty in Artificial Intelligence*, 2017.
- 476 [63] Behnam Neyshabur, Srinadh Bhojanapalli, and Nathan Srebro. A PAC-Bayesian approach  
477 to spectrally-normalized margin bounds for neural networks. In *International Conference on*  
478 *Learning Representations*, 2018.
- 479 [64] Xuanzhe Xiao, Zeng Li, Chuanlong Xie, and Fengwei Zhou. Heavy-tailed regularization of  
480 weight matrices in deep neural networks. *arXiv preprint arXiv:2304.02911*, 2023.
- 481 [65] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 and cifar-100 datasets. 2009.
- 482 [66] Pierre Sermanet, Koray Kavukcuoglu, and Yann LeCun. Traffic signs and pedestrians vision  
483 with multi-scale convolutional networks. In *Snowbird Machine Learning Workshop*, 2011.
- 484 [67] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale  
485 hierarchical image database. In *IEEE conference on Computer Vision and Pattern Recognition*,  
486 2009.
- 487 [68] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng  
488 Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual  
489 recognition challenge. *International journal of computer vision*, 115:211–252, 2015.
- 490 [69] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.
- 491 [70] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale  
492 image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- 493 [71] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image  
494 recognition. In *IEEE conference on Computer Vision and Pattern Recognition*, 2016.
- 495 [72] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *Proceedings of the British*  
496 *Machine Vision Conference*, 2016.
- 497 [73] Yang You, Jing Li, Sashank Reddi, Jonathan Hseu, Sanjiv Kumar, Srinadh Bhojanapalli, Xiaodan  
498 Song, James Demmel, Kurt Keutzer, and Cho-Jui Hsieh. Large batch optimization for deep  
499 learning: Training bert in 76 minutes. In *International Conference on Learning Representations*,  
500 2020.

- 501 [74] Ryo Karakida, Shotaro Akaho, and Shun-ichi Amari. Pathological spectra of the fisher infor-  
502 mation metric and its variants in deep neural networks. *Neural Computation*, 33:2274–2307,  
503 2019.
- 504 [75] Ryo Karakida, Shotaro Akaho, and Shun-ichi Amari. Universal statistics of fisher information in  
505 deep neural networks: Mean field approach. In *the 22nd International Conference on Artificial*  
506 *Intelligence and Statistics*, 2019.
- 507 [76] Ka-Veng Yuen. *Bayesian methods for structural dynamics and civil engineering*. John Wiley &  
508 Sons, 2010.
- 509 [77] Yudi Pawitan. *In all likelihood: statistical modelling and inference using likelihood*. Oxford  
510 University Press, 2001.
- 511 [78] Bohang Zhang, Du Jiang, Di He, and Liwei Wang. Rethinking lipschitz neural networks for  
512 certified l-infinity robustness. *arXiv preprint arXiv:2210.01787*, 2022.
- 513 [79] Zhewei Yao, Amir Gholami, Kurt Keutzer, and Michael W Mahoney. PyHessian: Neural  
514 networks through the lens of the hessian. In *IEEE International Conference on Big Data (Big*  
515 *Data)*, pages 581–590, 2020.
- 516 [80] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and  
517 generalization in neural networks. In *Advances in Neural Information Processing Systems*,  
518 2018.
- 519 [81] Yongqi Du, Di Xie, Shiliang Pu, Robert Qiu, Zhenyu Liao, et al. " lossless" compression of  
520 deep neural networks: A high-dimensional neural tangent kernel approach. In *Advances in*  
521 *Neural Information Processing Systems*, 2022.
- 522 [82] Zeke Xie, Qian-Yuan Tang, Zheng He, Mingming Sun, and Ping Li. Rethinking the structure  
523 of stochastic gradients: Empirical and statistical evidence. *arXiv preprint arXiv:2212.02083*,  
524 2022.
- 525 [83] Aitor Lewkowycz, Yasaman Bahri, Ethan Dyer, Jascha Sohl-Dickstein, and Guy Gur-Ari.  
526 The large learning rate phase of deep learning: the catapult mechanism. *arXiv preprint*  
527 *arXiv:2003.02218*, 2020.
- 528 [84] Ethan Dyer and Guy Gur-Ari. Asymptotics of wide networks from feynman diagrams. In  
529 *International Conference on Learning Representations*, 2020.
- 530 [85] Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in  
531 linear regression. In *Proceedings of the National Academy of Sciences*, 2020.

## 532 A Heavy-tail phenomena in different DNN matrices are closely related

533 Recently, several papers have separately studied the HT structures in different types of “important  
 534 matrices,” including Hessian, Fisher Information Matrix (FIM) and covariance matrices [35, 74, 75].  
 535 They confirm that when neural networks are well-trained, various matrices have an HT-shaped  
 536 spectrum. Among these works, there are two major ways to characterize the HT spectrum, namely  
 537 the HT-shaped ESDs (such as PL\_Alpha), or HT-shaped decaying eigenvalues [33–35]. Our paper  
 538 mainly uses the first way of characterizing the HT structure. On the other hand, the second way  
 539 is to sort eigenvalues from largest to smallest and study the PL phenomena between the ordered  
 540 eigenvalues and their index. Our experiments show fruitful connections between the PL phenomena  
 541 manifested in different DNN matrices; if one matrix shows the PL spectrum, the other matrices often  
 542 show something similar [35]. Thus, It is meaningful to ask why and how the PL phenomena in  
 543 different prior works correlate.

544 This section first establishes the connections between input/output covariance matrices, FIM and  
 545 Hessian in subsection A.1. We find that if one of these matrices shows the PL phenomenon, the other  
 546 two matrices have a high chance to exhibit a similar PL phenomenon. Then, in subsection A.2, we  
 547 derive the connection between our metric PL\_Alpha and the PL exponent on decaying eigenvalues,  
 548 showing a simple reciprocal relationship between these two.

### 549 A.1 Connections between different matrices

550 Consider a neural network (NN)  $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^C$ , where  $\theta \in \mathbb{R}^P$  is the vectorized weights,  $d$  is the  
 551 input dimension, and  $C$  is the output dimension. When the NN is used for a classifying task,  $C$  is also  
 552 the number of classes. We denote the input data as  $\{(x_i, y_i)\}_{i=1}^n$ , where  $x_i \in \mathbb{R}^d$ , and the number of  
 553 samples is  $n$ . We denote the loss function as  $L(\theta) = \frac{1}{n} \sum_{i=1}^n l(y_i, f_\theta(x_i))$ .

554 **Covariance matrices.** We denote the output covariance matrix as  $\mathbb{E}[f_\theta(x) f_\theta^\top(x)]$ , where the expect-  
 555 ation is taken over the input distribution. We tend to consider the following empirical covariance  
 556 matrix:

$$C(\theta) := \frac{1}{n} \sum_{i=1}^n f_\theta(x_i) f_\theta^\top(x_i) \in \mathbb{R}^{C \times C}. \quad (5)$$

557 **Fisher Information Matrices.** We denote the (output) FIM as

$$\mathbb{E}[\nabla_\theta f_\theta(x) \nabla_\theta f_\theta(x)^\top] = \sum_{k=1}^C \mathbb{E}[\nabla_\theta f_\theta^{(k)}(x) \nabla_\theta f_\theta^{(k)}(x)^\top], \quad (6)$$

558 where  $f_\theta^{(k)}(x)$  is the  $k$ -th entry of the vector function  $f(x)$ . We also consider the empirical version of  
 559 the FIM:

$$F(\theta) := \sum_{k=1}^C \frac{1}{n} \sum_{i=1}^n \nabla_\theta f_\theta^{(k)}(x_i) \nabla_\theta f_\theta^{(k)}(x_i)^\top \in \mathbb{R}^{P \times P}. \quad (7)$$

560 Note that (7) can be equally written as

$$F(\theta) := \frac{1}{n} \nabla_\theta \tilde{f}_\theta(x) \nabla_\theta \tilde{f}_\theta(x)^\top, \quad (8)$$

where  $\nabla_\theta \tilde{f}_\theta(x)$  has the following form:

$$\begin{bmatrix} \frac{\partial f_\theta^{(1)}(x_1)}{\partial \theta_1} \dots \frac{\partial f_\theta^{(1)}(x_n)}{\partial \theta_1} & \dots & \frac{\partial f_\theta^{(C)}(x_1)}{\partial \theta_1} \dots \frac{\partial f_\theta^{(C)}(x_n)}{\partial \theta_1} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_\theta^{(1)}(x_1)}{\partial \theta_P} \dots \frac{\partial f_\theta^{(1)}(x_n)}{\partial \theta_P} & \dots & \frac{\partial f_\theta^{(C)}(x_1)}{\partial \theta_P} \dots \frac{\partial f_\theta^{(C)}(x_n)}{\partial \theta_P} \end{bmatrix} \in \mathbb{R}^{P \times Cn}.$$

561 **Hessian Matrices.** We denote the Hessian as  $\mathbb{E} \left[ \frac{\partial^2 l(y, f_\theta(x))}{\partial \theta^2} \right]$ , and we tend to consider the empirical  
 562 Hessian Matrices:

$$H(\theta) := \frac{\partial^2 L(\theta)}{\partial \theta^2} \in \mathbb{R}^{P \times P}, \quad (9)$$

563 where  $L(\theta)$  is the empirical loss function  $L(\theta) = \frac{1}{n} \sum_{i=1}^n l(y_i, f_\theta(x_i))$ .

564 **Hessian and FIM are equivalent under certain conditions.** FIM can be defined in alternative ways  
565 different from (6). For instance, from classic statistical knowledge, we have the standard FIM (sFIM)  
566 in the following form:

$$sFIM := \mathbb{E}[\nabla_\theta \log P(y|x; \theta) \nabla_\theta \log P(y|x; \theta)^T], \quad (10)$$

567 where  $P(y|x; \theta)$  represents the likelihood; after simple derivations, one can show that sFIM also has  
568 the following form [76, 77]:

$$sFIM = -\mathbb{E} \left[ \frac{\partial^2 \log P(y|x; \theta)}{\partial \theta^2} \right]. \quad (11)$$

569 Therefore, when the loss function is defined as the negative log-likelihood, the sFIM in (11) is  
570 equivalent to Hessian defined in (9).

571 **Why is the FIM defined in (6) equivalent to (10).** Back to deep learning, FIM is often defined  
572 as (6). It is thus meaningful to derive the equivalence between these two forms. Suppose  $P(y|x; \theta)$   
573 here means the conditional probability distribution of output  $y$  given input data  $x$ . If  $P(y|x; \theta)$  is  
574 assumed to take the following form:

$$P(y|x; \theta) = \frac{1}{\sqrt{2\pi}} \exp \left( -\frac{1}{2} \|y - f_\theta(x)\|^2 \right), \quad (12)$$

575 then the MSE estimator  $\min_\theta \frac{1}{2} \|y - f_\theta(x)\|^2$  is equivalent to the maximum likelihood estimation of  
576  $P(y|x; \theta)$ . Then, plugging (12) into (10), we have:

$$sFIM_{mse} = \mathbb{E}[\|y - f_\theta(x)\|^2 \nabla_\theta f_\theta(x) \nabla_\theta f_\theta(x)^T]. \quad (13)$$

577 We now expand  $sFIM_{mse}$  by the definition of expectation, and we have the following [74]:

$$sFIM_{mse} = \int_{\mathbb{R}} \int_{\mathbb{R}} \|y - f_\theta(x)\|^2 \nabla_\theta f_\theta(x) \nabla_\theta f_\theta(x)^T p(x, y; \theta) dy dx \quad (14)$$

$$= \int_{\mathbb{R}} \int_{\mathbb{R}} \|y - f_\theta(x)\|^2 \nabla_\theta f_\theta(x) \nabla_\theta f_\theta(x)^T P(y|x; \theta) q(x) dy dx \quad (15)$$

$$= \int_{\mathbb{R}} \left[ \int_{\mathbb{R}} \frac{1}{\sqrt{2\pi}} \|y - f_\theta(x)\|^2 \exp \left( -\frac{1}{2} \|y - f_\theta(x)\|^2 \right) dy \right] \nabla_\theta f_\theta(x) \nabla_\theta f_\theta(x)^T q(x) dx \quad (16)$$

$$= \int_{\mathbb{R}} \nabla_\theta f_\theta(x) \nabla_\theta f_\theta(x)^T q(x) dx \quad (17)$$

$$= \mathbb{E}[\nabla_\theta f_\theta(x) \nabla_\theta f_\theta(x)^T], \quad (18)$$

578 where (14) follows from the definition of expectation,  $q(x)$  is input distribution, and (17) holds  
579 because the integral of  $y$  in the brackets [] equals 1 due to the property of Gamma function  $\Gamma(\cdot)$ .

580 Therefore, from (18), we find that  $sFIM_{mse}$  is just equal to  $FIM$  defined in (6). Also, plugging (12)  
581 into  $\mathbb{E} \left[ \frac{\partial^2 \log P(y|x; \theta)}{\partial \theta^2} \right]$  and taking the loss function  $L(\theta)$  as the mean-square loss, we will again find  
582  $\mathbb{E} \left[ \frac{\partial^2 \log P(y|x; \theta)}{\partial \theta^2} \right]$  is equal to  $H(\theta)$ . Therefore, jointly considering (11), we can see that FIM is equal  
583 to Hessian  $H(\theta)$ .

584 **PL in the covariance matrix and PL in Hessian are tightly correlated.** Next, we consider the  
585 relationship between the covariance matrix and Hessian. Suppose the NN function  $f_\theta$  is a Lipchitz  
586 function [78]. Then, it can be seen that the covariance matrix (5) may be controlled and estimated by  
587 FIM defined in (6), which is equivalent to be controlled by Hessian.

588 Although deriving an exact equivalent between these two can be hard, we numerically show that the  
589 PL in one matrix informs the PL in the other. To visualize their relationship in the presence of PL, we  
590 train a simple MLP with one hidden layer and 2000 neurons for 50 epochs. We leverage the spectral  
591 regularization from Nassar et al. [34] to make the output covariance matrix exhibit a PL spectrum.  
592 Meanwhile, we calculate the top eigenvalues of covariance and hessian [79], fit the PL exponent  $s$  for

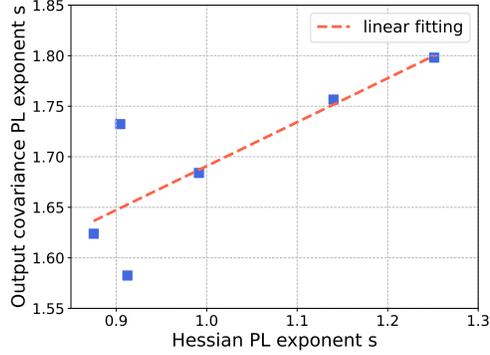


Figure 8: We train a MLP for 50 epochs and fit PL exponent  $s$  for both the output covariance and the Hessian. For models trained with epochs [1,10,20,30,40,50], we see their PL exponents  $s$  show a strong correlation.

593 each matrix, and compare the PL exponents against each other. More specifically, we take trained  
 594 NNs from epochs [1,10,20,30,40,50] and plot the Hessian PL exponent  $s$  versus the output covariance  
 595 PL exponent  $s$ . From the results shown in figure 8, we can see that their PL exponent  $s$  shows a  
 596 strong correlation, which supports our claim that the PL phenomena in one matrix can inform the  
 597 other.

598 **Connections to the NTK matrix.** Interestingly, if we ignore the constant in (8) and switch the two  
 599 matrices multiplied together, we obtain  $\nabla_{\theta} \tilde{f}_{\theta}(x)^T \nabla_{\theta} \tilde{f}_{\theta}(x)$ . This matrix is equal to Neural Tangent  
 600 Kernel(NTK) [80], which is a kernel used to approximate the deep neural network when NN’s  
 601 width is infinite. We thus conjecture that NTK should show PL when the NN is well trained [81].  
 602 Indeed, Karakida et al. [74] and Karakida et al. [75] study the eigenvalues of NTK, showing a PL  
 603 trend. Some other work on stochastic gradient [82] claim that the so-called “stochastic gradient  
 604 matrix” (which is similar to the NTK matrix) shows a PL spectrum as well, which matches our  
 605 expectations. Also, Lewkowycz et al. [83], Dyer and Gur-Ari [84] show that the eigenvalues of NTK  
 606 are similar to those in the Hessian, which again meets our expectation because the Hessian tends to  
 607 be PL when neural networks are well-trained [35].

608 In summary, the derivations above indicate that different “important matrices” are tightly correlated  
 609 to each other in terms of the PL trends: if one matrix shows a PL spectrum, there is a high chance  
 610 that the other ones show something similar.

## 611 A.2 Connections between PL in ESD and PL in decaying eigenvalues

612 Next, we derive the connection between our metric PL\_Alpha and the exponent of PL distribution on  
 613 decaying eigenvalues. Take the covariance matrix (5) as an instance. According to Nassar et al. [34],  
 614 the HT phenomenon in the output covariance matrix is similar to the layer-wise covariance matrices.  
 615 Thus, without the loss of generality, we can consider the case when there is only one layer in the  
 616 neural network. We assume the weight matrix  $L$  is in  $\mathbf{R}^{N \times M}$ . According to prior works, when  $L$  is  
 617 well-trained, the ESD follows a PL distribution:

$$p(\lambda) = \frac{1}{H} \lambda^{-\alpha}, \quad \lambda_{min} < \lambda < \lambda_{max}. \quad (19)$$

618 Here,  $H$  is a normalizing constant, and  $\alpha$  is the PL exponent.

619 Another way to characterize the PL phenomenon is to consider eigenvalues directly following a PL  
 620 series. For example, Xie et al. [35] show that the decaying eigenvalues follows the following PL  
 621 series:

$$\lambda_k = \lambda_1 k^{-s}, \quad k = 1, 2, \dots, M, \quad (20)$$

622 where  $\lambda_1$  is the same as  $\lambda_{max}$  used in the main paper.

623 Now, we will analytically and empirically show that these two ways of characterizing PL are  
 624 essentially equivalent. Furthermore, the two PL coefficients satisfy  $s = \frac{1}{\alpha-1}$ .

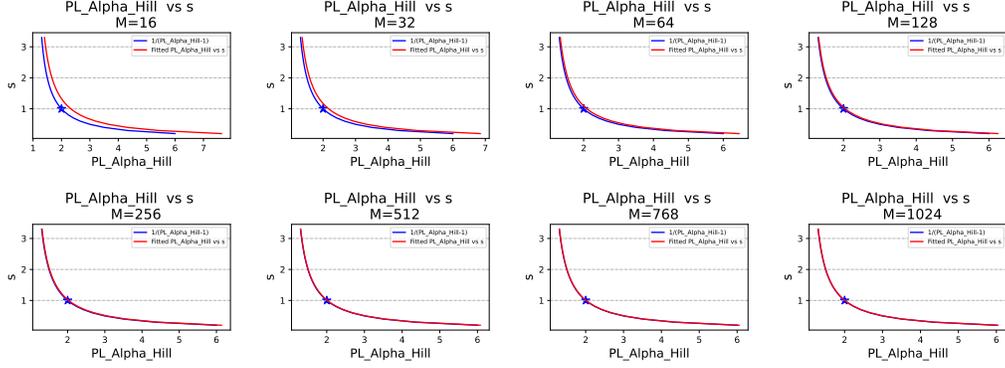


Figure 9: We show the connection between PL\_Alpha, estimated by PL\_Alpha\_Hill, and the PL exponent of the decaying eigenvalues (denoted as  $s$ ) satisfy  $s = \frac{1}{\text{PL\_Alpha\_Hill}-1}$ . Results are shown for different matrix size  $M$ . In particular, we see that PL\_Alpha = 2 [28] is equivalent to  $s = 1$  [33] in the linear case.

625 **An analytical way to show that  $s = \frac{1}{\alpha-1}$ .** The derivation is actually quite simple. Consider the case  
 626 that  $\lambda_k = \lambda_1 k^{-s}$  (i.e., (20) holds), and suppose  $\Lambda$  is a random variable distributed according to the  
 627 empirical distribution from these eigenvalues  $\lambda_k = \lambda_1 k^{-s}$ . Now, from (20), we easily see that the  
 628 distribution function takes the following form:

$$\mathbb{P}(\Lambda > \lambda_1 k^{-s}) = \frac{k}{M}. \quad (21)$$

629 By changing variables  $\lambda_1 k^{-s} = \lambda$ , we get the cumulative distribution function of  $\Lambda$ :

$$\mathbb{P}(\Lambda > \lambda) \sim \lambda^{-\frac{1}{s}}. \quad (22)$$

630 After that, we take the derivative with respect to  $\lambda$ , and we get the ESD:

$$p(\lambda) \sim \lambda^{-\left(\frac{1}{s}+1\right)}. \quad (23)$$

631 In other words, we have  $\lambda^{-\left(\frac{1}{s}+1\right)} = \lambda^{-\alpha}$ , which means  $s = \frac{1}{\alpha-1}$ .

632 **An empirical way to show that  $s = \frac{1}{\alpha-1}$ .** We consider matrices of size  $M \times M$ , where we choose  $M$   
 633 in [16, 32, 64, 128, 256, 512, 768, 1024] and assign the parameters such that the decaying eigenvalues  
 634 obey the formula  $\lambda_1 k^{-s}$ , for  $s$  in  $[0.2, 0.3, 0.4, \dots, 3.2]$ . Then, we fit the ESD and get our estimate  
 635 PL\_Alpha\_Hill. We plot the relationship between PL\_Alpha\_Hill and  $s$  in figure9. From figure  
 636 9, we find that the connection between  $\alpha$  and  $s$  shows a good fit with the formula  $s = \frac{1}{\alpha-1}$ . With  
 637 increasing matrix size  $M$ , the fitting becomes increasingly accurate.

638 **When  $s = \frac{1}{\alpha-1}$ ,  $s = 1$  corresponds to  $\alpha = 2$ .** Some prior works Nassar et al. [34], Xie et al.  
 639 [35], Bartlett et al. [85] measure the HT phenomena from the perspective of decaying eigenvalues  
 640 with PL exponent  $s$ , and they show either theoretically or empirically that  $s = 1$  is the *optimal*  
 641 exponent. Now that we have  $s = \frac{1}{\alpha-1}$  in the linear case, and from the theory of NTK[80], the infinite  
 642 wide neural network is approximated as a linear model, we tend to believe that  $\alpha = 2$  satisfies a  
 643 similar property. Indeed, one of the main contributions of Martin and Mahoney [28] is to establish  
 644 different HT families of ESDs, and 2 is believed to be the boundary between “moderately HT” and  
 645 “very HT.” Martin and Mahoney [28] further argue that the optimal exponent for PL\_Alpha is in  
 646 the range [2,4]. Combining the perspective from Nassar et al. [34], Xie et al. [35], Bartlett et al.  
 647 [85] and those from Martin and Mahoney [28], it is reasonable to believe that the optimal exponent  
 648 for PL\_Alpha is around 2. When PL\_Alpha is much higher or lower than two, the NN probably  
 649 has some issue in training. Although we argued in the main paper that the “absolute value” of  
 650 PL\_Alpha is unimportant in implementing our TempBalance algorithm, it is, however, helpful to  
 651 have an “optimal” PL\_Alpha value to test if our algorithm actually works in controlling the ESDs.  
 652 We will show visualization results in Appendix B that TempBalance leads to a better distribution of  
 653 our estimated PL\_Alpha\_Hill.

654 **B Visualization results: how does TempBalance control ESDs**

655 In this section, we demonstrate that the proposed method, TempBalance, effectively controls the  
 656 shape of the empirical spectral densities (ESDs), resulting in a more favorable distribution of  
 657 PL\_Alpha\_Hill among the layers of neural networks (NNs) compared to the baseline method  
 658 CAL. This observation elucidates the superior performance of TempBalance over CAL in our main  
 659 experiment, as presented in Section 4.2.

660 We evaluate the models reported in Figure 3. For each individual NN, we compute and aggregate  
 661 PL\_Alpha\_Hill values across all layers, excluding the first and last layers that have an extremely  
 662 small number of eigenvalues and thus cause inaccurate PL\_Alpha\_Hill estimation. We aggregate  
 663 the PL\_Alpha\_Hill values from five models trained using different random seeds for each method.  
 664 Figure 10 shows the distribution of PL\_Alpha\_Hill of TempBalance and the baseline CAL. Comparing  
 665 TempBalance with CAL, we see that TempBalance consistently yields a more concentrated  
 666 distribution. Furthermore, TempBalance causes the median and mean of the distribution to approach  
 667 2 (shown in each subplot respectively as the middle vertical line and the red star). The value 2  
 668 represents the theoretically optimal PL\_Alpha\_Hill value, as we have justified in Appendix A.

669 Next, in Figure 11, we group the models into different subgroups based on their architectures and/or  
 670 datasets, aggregating the PL\_Alpha\_Hill values and comparing the distributions of the two methods  
 671 TempBalance and CAL. Once again, we observe that TempBalance results in a more concentrated  
 672 distribution, with a larger number of samples (layers) having PL\_Alpha\_Hill values closer to 2.

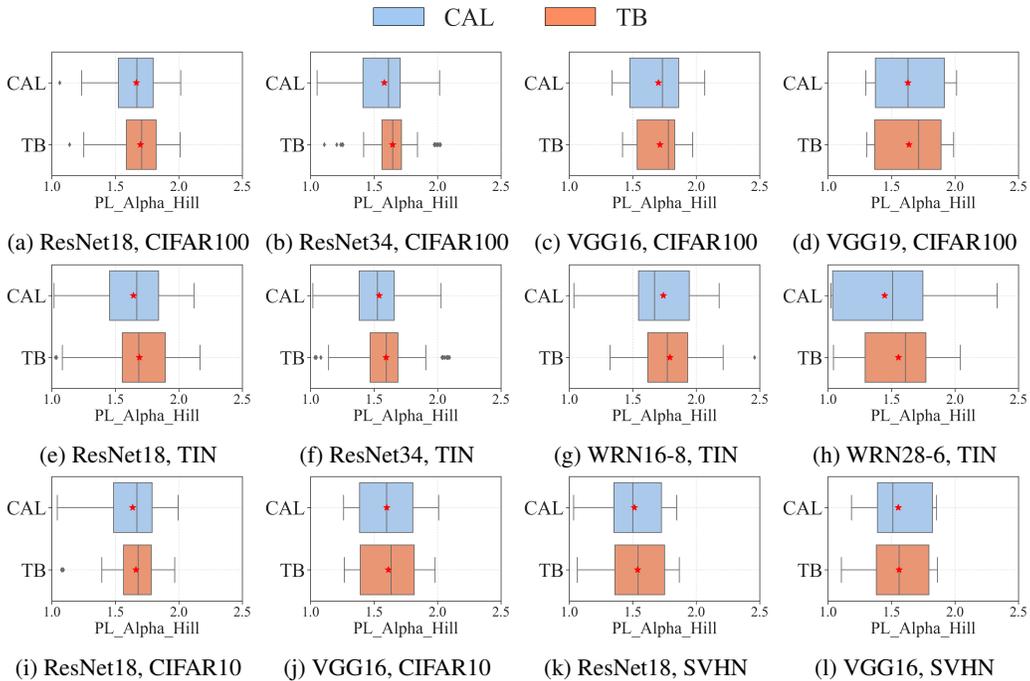


Figure 10: Comparing the distribution of PL\_Alpha\_Hill of NNs trained by our method TempBalance (TB) and CAL. The mean of each distribution is indicated by a red star marker. Each distribution aggregates the PL\_Alpha\_Hill values from models trained using five different random seeds. Across all experiments, our method TempBalance consistently yields a more concentrated distribution, resulting in the mean and median approaching the theoretically optimal PL\_Alpha\_Hill value of 2, as supported in Appendix A.

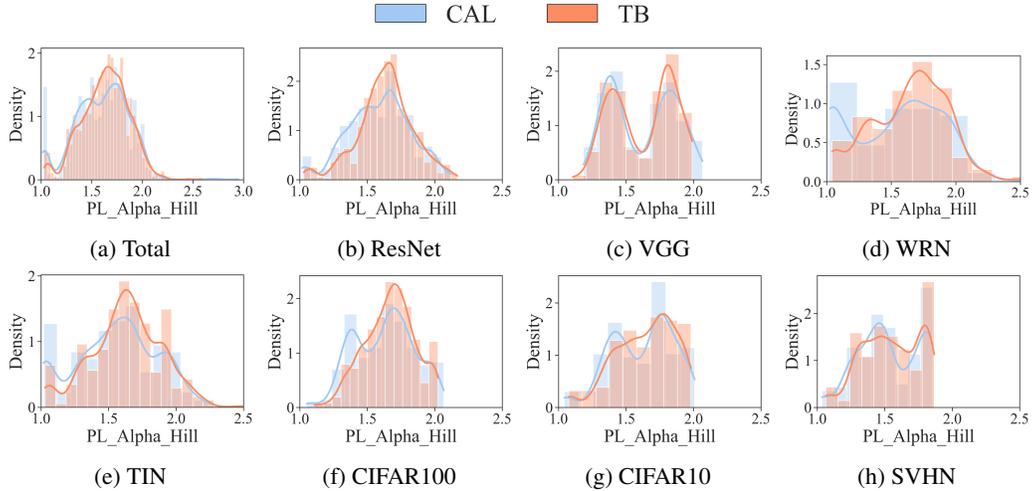


Figure 11: Comparing our method TempBalance (TB) to CAL in terms of the distribution of PL\_Alpha\_Hill of aggregating NNs into different architectures and datasets. Each distribution aggregates the PL\_Alpha\_Hill of models trained with five random seeds. Across all subgroups, our method TempBalance consistently exhibits a more concentrated distribution, accompanied by a higher number of layers approaching a PL\_Alpha\_Hill value close to 2. This value of 2 corresponds to the theoretically optimal PL\_Alpha\_Hill value, as justified in Appendix A.

## 673 C Hyperparameter settings for reproducing our results

674 In this section, we report all hyperparameters, random seeds and all numerical values of experimental  
 675 results shown in the main paper (in Section 4).

676 First, we report the common hyperparameters shared by all the experiments: the default optimizer is  
 677 SGD, trained with batch size 128, number of training epochs 200, weight decay  $5e-4$ , and momentum  
 678 0.9. The default HT-SR metric used in TempBalance is PL\_Alpha\_Hill. For each experimental  
 679 setting, we use five random seeds, which are always 43, 37, 13, 51, 71, and we report the mean and  
 680 standard deviation of the test accuracy across these seeds.

681 First, Table 1 reports the details of experiments shown in Figure 3. We carefully tune the initial  
 682 learning rate  $\eta_0$  and  $\lambda_{sr}$  for the two baseline methods CAL and SNR. Then, Table 2 reports the  
 683 detailed hyperparameter settings of the experiments shown in Figure 4. We again carefully tune the  
 684 hyperparameters of various baseline optimizers and schedulers, as specified in their papers. Finally,  
 685 Table 3, Table 4 and Table 5 respectively report the details of the experiments shown in Figure 5,  
 686 Figure 6 and Figure 7.

Table 1: Parameter settings of the experiment reported in Section 4.2 Figure 3. The hyperparameter in bold is the best hyperparameter selection reported in the main paper. The five random seeds for each setting are {43, 37, 13, 51, 71}, and the means and standard deviations of the test accuracy among the five seeds are reported.

Index	Dataset	Model	Method	Initial learning rate $\eta_0$	$\lambda_{sr}$	Test Acc (best hyperparam.)	scaling ratio ( $s_1, s_2$ )
0	CIFAR100	ResNet18	CAL	0.05, <b>0.1</b> , 0.15	-	78.31 $\pm$ 0.05	-
1		ResNet18	SNR	0.1	0.001, 0.005, <b>0.01</b> , 0.015	78.65 $\pm$ 0.29	-
2		ResNet18	TB	0.1	-	78.97 $\pm$ 0.29	(0.5, 1.5)
3		ResNet18	TB + SNR	0.1	0.001	79.06 $\pm$ 0.32	(0.6, 1.4)
4		ResNet34	CAL	<b>0.05</b> , 0.1, 0.15	-	78.98 $\pm$ 0.14	-
5		ResNet34	SNR	0.1	0.001, 0.005, 0.01, <b>0.015</b>	79.97 $\pm$ 0.21	-
6		ResNet34	TB	0.1	-	79.89 $\pm$ 0.15	(0.5, 1.5)
7		ResNet34	TB + SNR	0.1	0.005	80.09 $\pm$ 0.35	(0.6, 1.4)
8		VGG16	CAL	0.025, <b>0.05</b> , 0.1	-	74.59 $\pm$ 0.23	-
9		VGG16	SNR	0.05	0.001, <b>0.005</b> , 0.01, 0.015	74.80 $\pm$ 0.28	-
10		VGG16	TB	0.05	-	74.96 $\pm$ 0.15	(0.5, 1.5)
11		VGG16	TB + SNR	0.05	0.005	75.52 $\pm$ 0.46	(0.6, 1.4)
12		VGG19	CAL	0.025, <b>0.05</b> , 0.1	-	73.26 $\pm$ 0.37	-
13		VGG19	SNR	0.05	0.001, 0.005, <b>0.01</b> , 0.015	74.37 $\pm$ 0.16	-
14		VGG19	TB	0.05	-	73.77 $\pm$ 0.43	(0.5, 1.5)
15	VGG19	TB + SNR	0.05	0.01	74.74 $\pm$ 0.10	(0.5, 1.5)	
16	TinyImageNet	ResNet18	CAL	0.05, <b>0.1</b> , 0.15	-	66.25 $\pm$ 0.17	-
17		ResNet18	SNR	0.1	0.001, 0.005, <b>0.01</b> , 0.015	66.20 $\pm$ 0.22	-
18		ResNet18	TB	0.1	-	66.77 $\pm$ 0.25	(0.6, 1.4)
19		ResNet18	TB + SNR	0.1	0.001	66.86 $\pm$ 0.22	(0.6, 1.4)
20		ResNet34	CAL	0.05, <b>0.1</b> , 0.15	-	68.19 $\pm$ 0.16	-
21		ResNet34	SNR	0.1	0.001, 0.005, <b>0.01</b> , 0.015	68.69 $\pm$ 0.13	-
22		ResNet34	TB	0.1	-	69.12 $\pm$ 0.16	(0.6, 1.4)
23		ResNet34	TB + SNR	0.1	0.001	69.27 $\pm$ 0.21	(0.6, 1.4)
24		WRN16-8	CAL	0.05, <b>0.1</b> , 0.15	-	63.67 $\pm$ 0.09	-
25		WRN16-8	SNR	0.1	0.00005, <b>0.0001</b> , 0.001	63.98 $\pm$ 0.23	-
26		WRN16-8	TB	0.1	-	64.09 $\pm$ 0.17	(0.6, 1.4)
27		WRN16-8	TB + SNR	0.1	0.0001	64.08 $\pm$ 0.07	(0.6, 1.4)
28		WRN28-6	CAL	<b>0.05</b> , 0.1, 0.15	-	65.88 $\pm$ 0.20	-
29		WRN28-6	SNR	0.1	0.00005, <b>0.0001</b> , 0.001	66.09 $\pm$ 0.25	-
30		WRN28-6	TB	0.1	-	66.58 $\pm$ 0.23	(0.6, 1.4)
31	WRN28-6	TB + SNR	0.1	0.0001	66.79 $\pm$ 0.25	(0.6, 1.4)	
32	CIFAR10	ResNet18	CAL	0.05, <b>0.1</b> , 0.15	-	95.53 $\pm$ 0.12	-
33		ResNet18	SNR	0.1	<b>0.001</b> , 0.005, 0.01, 0.015	95.57 $\pm$ 0.06	-
34		ResNet18	TB	0.1	-	95.63 $\pm$ 0.08	(0.5, 1.5)
35		ResNet18	TB + SNR	0.1	0.001	95.66 $\pm$ 0.09	(0.6, 1.4)
36		VGG16	CAL	0.025, 0.05, <b>0.1</b>	-	93.98 $\pm$ 0.12	-
37		VGG16	SNR	0.05	0.001, <b>0.005</b> , 0.01, 0.015	94.04 $\pm$ 0.07	-
38		VGG16	TB	0.05	-	94.14 $\pm$ 0.06	(0.5, 1.5)
39	VGG16	TB + SNR	0.05	0.005	94.26 $\pm$ 0.10	(0.6, 1.4)	
40	SVHN	ResNet18	CAL	0.05, <b>0.1</b> , 0.15	-	96.59 $\pm$ 0.08	-
41		ResNet18	SNR	0.1	0.001, <b>0.005</b> , 0.015, 0.01	96.65 $\pm$ 0.12	-
42		ResNet18	TB	0.1	-	96.63 $\pm$ 0.06	(0.5, 1.5)
43		ResNet18	TB + SNR	0.1	0.01	96.67 $\pm$ 0.09	(0.6, 1.4)
44		VGG16	CAL	0.025, <b>0.05</b> , 0.1	-	96.28 $\pm$ 0.04	-
45		VGG16	SNR	0.05	0.001, 0.005, <b>0.015</b> , 0.01	96.32 $\pm$ 0.07	-
46		VGG16	TB	0.05	-	96.33 $\pm$ 0.06	(0.5, 1.5)
47	VGG16	TB + SNR	0.05	0.005	96.40 $\pm$ 0.08	(0.6, 1.4)	

Table 2: Parameter settings of the experiment reported in Section 4.2 Figure 4. The hyperparameter in bold is the best hyperparameter selection reported in the main paper. The five random seeds for each setting are {43, 37, 13, 51, 71}, and the means and standard deviations of the test accuracy among the five seeds are reported.

Index	Dataset	Model	Method	Initial learning rate $\eta_0$	SGDR ( $T_0, T_{mut}$ )	Lookahead $k$	Lookahead $\alpha$	Test Acc (best hyperparams.)	scaling ratio ( $s_1, s_2$ )
0	CIFAR100	ResNet18	CAL	0.05, <b>0.1</b> , 0.15	-	-	-	78.31 $\pm$ 0.05	-
1		ResNet18	SGDR	0.05, <b>0.1</b> , 0.15	( <b>100,1</b> ), (10, 2), (1, 2)	-	-	77.69 $\pm$ 0.20	-
2		ResNet18	LARS	26, <b>28</b> , 30, 32, 34	-	-	-	78.44 $\pm$ 0.12	-
3		ResNet18	Lookahead	0.05, <b>0.1</b> , 0.15	-	-	<b>10, 5</b>	78.46 $\pm$ 0.18	-
4		ResNet18	SGDP	0.01, 0.05, <b>0.1</b> , 0.15, 0.2	-	-	-	78.74 $\pm$ 0.11	-
5		ResNet18	TB	0.05, <b>0.1</b> , 0.15	-	-	-	78.97 $\pm$ 0.29	(0.5, 1.5)
6		ResNet18	TB + SGDP	0.05, <b>0.1</b> , 0.15	-	-	-	79.13 $\pm$ 0.15	(0.5, 1.5)
7		ResNet34	CAL	<b>0.05</b> , 0.1, 0.15	-	-	-	78.98 $\pm$ 0.14	-
8		ResNet34	SGDR	<b>0.05</b> , 0.1, 0.15	( <b>100,1</b> ), (10, 2), (1, 2)	-	-	78.61 $\pm$ 0.20	-
9		ResNet34	LARS	26, 28, 30, <b>32</b> , 34	-	-	-	78.94 $\pm$ 0.19	-
10		ResNet34	Lookahead	0.05, 0.1, <b>0.15</b>	-	-	<b>10, 5</b>	79.19 $\pm$ 0.12	-
11		ResNet34	SGDP	0.01, 0.05, <b>0.1</b> , 0.15, 0.2	-	-	-	79.34 $\pm$ 0.21	-
12		ResNet34	TB	0.05, <b>0.1</b> , 0.15	-	-	-	79.89 $\pm$ 0.15	(0.5, 1.5)
13	ResNet34	TB + SGDP	0.05, <b>0.1</b> , 0.15	-	-	-	79.94 $\pm$ 0.30	(0.5, 1.5)	

Table 3: Parameter settings of the experiment reported in Section 4.3 Figure 5. The five random seeds for each setting are {43, 37, 13, 51, 71}, and the means and standard deviations of the test accuracy among the five seeds are reported.

Index	Dataset	Model	Method	Initial		Test Acc	scaling ratio ( $s_1, s_2$ )
				learning rate $\eta_0$			
0	CIFAR100	ResNet18	CAL	0.05, 0.1, 0.15		$78.08 \pm 0.19, 78.31 \pm 0.05, 77.72 \pm 0.44$	-
1		ResNet18	TB	0.05, 0.1, 0.15		$78.48 \pm 0.27, 78.97 \pm 0.29, 78.69 \pm 0.11$	(0.5, 1.5)
2		ResNet34	CAL	0.05, 0.1, 0.15		$78.98 \pm 0.14, 78.89 \pm 0.24, 78.51 \pm 0.34$	-
3		ResNet34	TB	0.05, 0.1, 0.15		$79.36 \pm 0.18, 79.89 \pm 0.15, 79.09 \pm 0.64$	(0.5, 1.5)
4		VGG16	CAL	0.025, 0.05, 0.1		$73.96 \pm 0.27, 74.59 \pm 0.23, 74.46 \pm 0.12$	-
5		VGG16	TB	0.025, 0.05, 0.1		$74.40 \pm 0.31, 74.96 \pm 0.15, 74.94 \pm 0.16$	(0.5, 1.5)
6		VGG19	CAL	0.025, 0.05, 0.1		$72.57 \pm 0.45, 73.26 \pm 0.37, 72.98 \pm 0.16$	-
7		VGG19	TB	0.025, 0.05, 0.1		$73.47 \pm 0.16, 73.77 \pm 0.43, 73.40 \pm 0.38$	(0.5, 1.5)

Table 4: Parameter settings of the experiment reported in Section 4.3 Figure 6. The five random seeds for each setting are {43, 37, 13, 51, 71}, and the means and standard deviations of the test accuracy among the five seeds are reported.

Index	Dataset	Model	Method	Initial		Test Acc	scaling ratio ( $s_1, s_2$ )
				learning rate $\eta_0$	Width		
0	CIFAR100	ResNet18	CAL	0.1	256, 512, 768	$75.05 \pm 0.26, 78.31 \pm 0.05, 79.44 \pm 0.26$	-
1		ResNet18	TB	0.1	256, 512, 768	$75.63 \pm 0.12, 78.97 \pm 0.29, 80.47 \pm 0.18$	(0.5, 1.5)
2		ResNet34	CAL	0.1	256, 512, 768	$76.79 \pm 0.34, 78.89 \pm 0.24, 79.94 \pm 0.31$	-
3		ResNet34	TB	0.1	256, 512, 768	$77.25 \pm 0.14, 79.89 \pm 0.15, 80.23 \pm 0.53$	(0.5, 1.5)
4		VGG16	CAL	0.05	256, 512, 768	$71.04 \pm 0.14, 74.59 \pm 0.23, 75.53 \pm 0.32$	-
5		VGG16	TB	0.05	256, 512, 768	$71.26 \pm 0.26, 74.96 \pm 0.15, 76.19 \pm 0.14$	(0.5, 1.5)
6		VGG19	CAL	0.05	256, 512, 768	$69.58 \pm 0.39, 73.26 \pm 0.37, 74.39 \pm 0.33$	-
7		VGG19	TB	0.05	256, 512, 768	$69.96 \pm 0.25, 73.77 \pm 0.43, 74.80 \pm 0.35$	(0.5, 1.5)

Table 5: Parameter settings of the experiment reported in Section 4.3 Figure 7. The five random seeds for each setting are {43, 37, 13, 51, 71}, and the means and standard deviations of the test accuracy among the five seeds are reported.

Index	Dataset	Model	Method	HT-SR Metric	Initial		Test Acc	scaling ratio ( $s_1, s_2$ )
					learning rate $\eta_0$			
0	CIFAR100	ResNet18	TB	SpectralNorm	0.05, 0.1, 0.15		$77.83 \pm 0.21, 78.30 \pm 0.32, 78.27 \pm 0.25$	(0.5, 1.5)
1		ResNet18	TB	AlphaWeighted	0.05, 0.1, 0.15		$78.18 \pm 0.27, 78.67 \pm 0.17, 78.48 \pm 0.24$	(0.5, 1.5)
1		ResNet18	TB	PL_Alpha_Hill	0.05, 0.1, 0.15		$78.48 \pm 0.27, 78.97 \pm 0.29, 78.69 \pm 0.11$	(0.5, 1.5)
2		ResNet34	TB	SpectralNorm	0.05, 0.1, 0.15		$78.25 \pm 0.16, 78.71 \pm 0.15, 78.92 \pm 0.28$	(0.5, 1.5)
3		ResNet34	TB	AlphaWeighted	0.05, 0.1, 0.15		$78.36 \pm 0.39, 78.87 \pm 0.34, 78.83 \pm 0.23$	(0.5, 1.5)
3		ResNet34	TB	PL_Alpha_Hill	0.05, 0.1, 0.15		$79.36 \pm 0.18, 79.89 \pm 0.15, 79.09 \pm 0.64$	(0.5, 1.5)
4		VGG16	TB	SpectralNorm	0.025, 0.05, 0.1		$73.58 \pm 0.19, 74.29 \pm 0.16, 74.17 \pm 0.28$	(0.5, 1.5)
5		VGG16	TB	AlphaWeighted	0.025, 0.05, 0.1		$73.97 \pm 0.22, 74.19 \pm 0.11, 74.42 \pm 0.31$	(0.5, 1.5)
5		VGG16	TB	PL_Alpha_Hill	0.025, 0.05, 0.1		$74.40 \pm 0.31, 74.96 \pm 0.15, 74.94 \pm 0.16$	(0.5, 1.5)
6		VGG19	TB	SpectralNorm	0.025, 0.05, 0.1		$72.34 \pm 0.26, 72.91 \pm 0.35, 73.04 \pm 0.39$	(0.5, 1.5)
7		VGG19	TB	AlphaWeighted	0.025, 0.05, 0.1		$72.85 \pm 0.16, 73.41 \pm 0.17, 73.33 \pm 0.21$	(0.5, 1.5)
7		VGG19	TB	PL_Alpha_Hill	0.025, 0.05, 0.1		$73.47 \pm 0.16, 73.77 \pm 0.43, 73.40 \pm 0.38$	(0.5, 1.5)