

IMPROVED GENERALIZATION BY NOISE ENHANCEMENT

Anonymous authors

Paper under double-blind review

ABSTRACT

Recent studies have demonstrated that noise in stochastic gradient descent (SGD) is closely related to generalization: A larger SGD noise, if not too large, results in better generalization. Since the covariance of the SGD noise is proportional to η^2/B , where η is the learning rate and B is the minibatch size of SGD, the SGD noise has so far been controlled by changing η and/or B . However, too large η results in instability in the training dynamics and a small B prevents scalable parallel computation. It is thus desirable to develop a method of controlling the SGD noise without changing η and B . In this paper, we propose a method that achieves this goal using “noise enhancement”, which is easily implemented in practice. We expound the underlying theoretical idea and demonstrate that the noise enhancement actually improves generalization for real datasets. It turns out that large-batch training with the noise enhancement even shows better generalization compared with small-batch training.

1 INTRODUCTION

It is a big theoretical challenge in deep learning studies to understand why networks trained via stochastic gradient descent (SGD) and its variants generalize so well in the overparameterized regime, in which the number of network parameters greatly exceeds that of the training data samples (Zhang et al., 2017). This fundamental problem has been tackled from different points of view (Dziugaite & Roy, 2017; Nagarajan & Kolter, 2017; Neyshabur et al., 2017; 2019; Arora et al., 2018; Pérez et al., 2019; Jacot et al., 2018; Arora et al., 2019; D’Ascoli et al., 2020). Among them, some recent studies have pointed out the importance of an implicit regularization effect of SGD (Zhu et al., 2019; Wu et al., 2019; Smith et al., 2020). Indeed, it is empirically known that the SGD noise strength is strongly correlated with generalization of the trained network (Li et al., 2017; Jastrzębski et al., 2017; Goyal et al., 2017; Smith & Le, 2018; Hoffer et al., 2017; 2019). It has also been argued that the SGD noise prefers wide flat minima, which are considered to indicate good generalization (Keskar et al., 2017; Hoffer et al., 2017; Wu et al., 2018). From this viewpoint, not only its strength, but also the structure of the SGD noise is considered to be important since it is theoretically shown that the network can efficiently escape from bad local minima with the help of the SGD noise but not of an isotropic Gaussian noise with the same strength (Zhu et al., 2019; Wu et al., 2019).

The covariance of the SGD noise is proportional to η^2/B , where η and B denote the learning rate and the minibatch size, respectively, and hence, the SGD noise strength can be controlled by changing η and/or B . To realize good generalization, we want to increase the SGD noise strength by increasing η and/or decreasing B . However, when η becomes too large, the training dynamics often becomes unstable and the training fails. On the other hand, decreasing B prevents an efficient parallelization using multiple GPUs or TPUs.¹ It is therefore desirable to control the SGD noise without changing these hyperparameters.

The main contribution of the present paper is to show that the SGD noise can be controlled without changing η and B by a simple yet efficient method that we call *noise enhancement*. In this method, the gradient of the loss function is evaluated by using two independent minibatches. We will explain our theoretical idea in Sec. 2. We will also demonstrate that the noise enhancement improves

¹However, it is not at all trivial whether the large-batch training is really efficient even with an ideal parallelization. See Golmant et al. (2018); Hoffer et al. (2019) for scalability of large-batch training.

generalization in Sec. 3. In particular, it is empirically shown that the large-batch training using the noise enhancement even outperforms the small-batch training. This result gives us some insights into the relation between the SGD noise and generalization, which is discussed in Sec. 4. Because of its simplicity in implementation, this method would also be useful in practice.

2 NOISE ENHANCEMENT

We shall consider a classification problem. The training dataset $\mathcal{D} = \{(x^{(\mu)}, y^{(\mu)})\}_{\mu=1,2,\dots,N}$ consists of pairs of the input data vector $x^{(\mu)}$ and its label $y^{(\mu)}$. The set of all the network parameters is simply denoted by w . Then the output of the network for a given input x is denoted by $f(x; w)$. The loss function is defined as

$$L(w) = \frac{1}{N} \sum_{\mu=1}^N \ell(f(x^{(\mu)}; w), y^{(\mu)}) \equiv \frac{1}{N} \sum_{\mu=1}^N \ell_{\mu}(w), \quad (1)$$

where the function $\ell(\cdot, \cdot)$ specifies the loss (in this paper we employ the cross-entropy loss).

In the SGD, the training data is divided into minibatches of size B , and the parameter update is done by using one of them. Let $\mathcal{B}_t \subset \{1, 2, \dots, N\}$ with $|\mathcal{B}_t| = B$ be a random minibatch chosen at the t -th step, the network parameter w_t is updated as

$$w_{t+1} = w_t - \eta \nabla_w L_{\mathcal{B}_t}(w_t), \quad L_{\mathcal{B}_t}(w) = \frac{1}{B} \sum_{\mu \in \mathcal{B}_t} \ell_{\mu}(w_t) \quad (2)$$

in vanilla SGD, where $\eta > 0$ is the learning rate. It is also expressed as

$$w_{t+1} = w_t - \eta \nabla_w L(w_t) - \eta [\nabla_w L_{\mathcal{B}_t}(w_t) - \nabla_w L(w_t)] \equiv w_t - \eta \nabla_w L(w_t) - \xi_t(w_t). \quad (3)$$

Here, ξ_t corresponds to the SGD noise since its average over samplings of random minibatches is zero: $\mathbb{E}_{\mathcal{B}_t}[\xi_t] = 0$. Its covariance is also calculated straightforwardly (Zhu et al., 2019):

$$\begin{aligned} \mathbb{E}_{\mathcal{B}_t}[\xi_t \xi_t^T] &= \frac{\eta^2}{B} \frac{N-B}{N-1} \left(\frac{1}{N} \sum_{\mu=1}^N \nabla_w \ell_{\mu} \nabla_w \ell_{\mu}^T - \nabla_w L \nabla_w L^T \right) \\ &\approx \frac{\eta^2}{B} \left(\frac{1}{N} \sum_{\mu=1}^N \nabla_w \ell_{\mu} \nabla_w \ell_{\mu}^T - \nabla_w L \nabla_w L^T \right), \end{aligned} \quad (4)$$

where we assume $N \gg B$ in obtaining the last expression. This expression² shows that the SGD noise strength is controlled by η and B .

We want to enhance the SGD noise without changing η and B . Naively, it is possible just by replacing ξ_t by $\alpha \xi_t$ with a new parameter $\alpha > 1$. Equation (3) is then written as

$$\begin{aligned} w_{t+1} &= w_t - \eta \nabla_w L(w_t) - \alpha \xi_t(w_t) \\ &= w_t - \eta [\alpha \nabla_w L_{\mathcal{B}_t}(w_t) + (1 - \alpha) \nabla_w L(w_t)]. \end{aligned} \quad (5)$$

Practically, Eq. (5) would be useless because the computation of $\nabla_w L(w_t)$, i.e. the gradient of the loss function over the entire training data, is required for each iteration.³ Instead, we propose replacing $\nabla_w L(w_t)$ in Eq. (5) by $\nabla_w L_{\mathcal{B}'_t}(w_t)$, where \mathcal{B}'_t is another minibatch of the same size B that is independent of \mathcal{B}_t . We thus obtain the following update rule of the *noise-enhanced SGD*:

$$w_{t+1} = w_t - \eta [\alpha \nabla_w L_{\mathcal{B}_t}(w_t) + (1 - \alpha) \nabla_w L_{\mathcal{B}'_t}(w_t)]. \quad (6)$$

²From Eq. (4), some authors (Krizhevsky, 2014; Hoffer et al., 2017) argue that the SGD noise strength is proportional to η/\sqrt{B} , while others (Li et al., 2017; Jastrzyski et al., 2017; Smith et al., 2018) argue that it is rather proportional to $\sqrt{\eta/B}$ on the basis of the stochastic differential equation obtained for an infinitesimal $\eta \rightarrow +0$. Thus the learning-rate dependence of the noise strength is rather complicated.

³If we have computational resources large enough to realize ideal parallelization for full training dataset, this naive noise enhancement would work. However, with limited computational resources, it is not desirable that we have to evaluate $\nabla_w L(w_t)$ for each iteration.

Table 1: Network configurations.

Name	Network type	Dataset	L^*	L^{**}
F1	Fully connected	Fashion-MNIST	0.01	0.001
C1	Convolutional	Cifar-10	0.01	0.001
C2	Convolutional	Cifar-100	0.02	0.001

By defining the SGD noise ξ'_t associated with \mathcal{B}'_t as

$$\xi'_t(w_t) = \eta [\nabla_w L_{\mathcal{B}'_t}(w_t) - \nabla_w L(w_t)], \quad (7)$$

Eq. (6) is rewritten as

$$w_{t+1} = w_t - \eta \nabla_w L(w_t) - \xi_t^{\text{NE}}(w_t), \quad (8)$$

where the noise ξ_t^{NE} in the noise-enhanced SGD is given by

$$\xi_t^{\text{NE}} = \alpha \xi_t + (1 - \alpha) \xi'_t. \quad (9)$$

Its mean is obviously zero, i.e. $\mathbb{E}_{\mathcal{B}_t, \mathcal{B}'_t}[\xi_t^{\text{NE}}] = 0$, and its covariance is given by

$$\begin{aligned} \mathbb{E}_{\mathcal{B}_t, \mathcal{B}'_t} [\xi_t^{\text{NE}} (\xi_t^{\text{NE}})^{\text{T}}] &= \alpha^2 \mathbb{E}_{\mathcal{B}_t} [\xi_t \xi_t^{\text{T}}] + (1 - \alpha^2) \mathbb{E}_{\mathcal{B}'_t} [\xi'_t (\xi'_t)^{\text{T}}] \\ &= [\alpha^2 + (1 - \alpha)^2] \mathbb{E}_{\mathcal{B}_t} [\xi_t \xi_t^{\text{T}}], \end{aligned} \quad (10)$$

where we have used the fact that two noises ξ_t and ξ'_t are i.i.d. random variables. In this way, the SGD-noise covariance is enhanced by a factor of $\alpha^2 + (1 - \alpha)^2 > 1$ for $\alpha > 1$. Since the size of the new minibatch \mathcal{B}'_t is same as that of the original minibatch \mathcal{B}_t , the noise enhancement does not suffer from any serious computational cost.

If we assume $N \gg B$, Eq. (10) is equivalent to Eq. (4) with an effective minibatch size

$$B_{\text{eff}} = \frac{B}{\alpha^2 + (1 - \alpha)^2}. \quad (11)$$

If the SGD noise were Gaussian, it would mean that the noise-enhanced SGD is equivalent to the normal SGD with the effective minibatch size B_{eff} . However, the SGD noise is actually far from Gaussian during training (Panigrahi et al., 2019), at least for not too large minibatch size. The noise enhancement is therefore not equivalent to reducing the minibatch size unless B_{eff} is too large.

The procedure of the noise enhancement is summarized as the follows: (i) prepare two independent minibatches \mathcal{B}_t and \mathcal{B}'_t , and (ii) replace the minibatch gradient $\nabla_w L_{\mathcal{B}_t}(w_t)$ by $\alpha \nabla_w L_{\mathcal{B}_t}(w_t) + (1 - \alpha) \nabla_w L_{\mathcal{B}'_t}(w_t)$. The numerical implementation is quite simple. It should be noted that the noise enhancement is also applicable to other variants of SGD like Adam.

3 EXPERIMENT

We shall demonstrate the efficiency of the method of the noise enhancement (NE) for several network configurations with a real dataset as listed in Table 1.

We describe the details of the network architecture below:

- **F1:** A fully-connected feed-forward network with 7 hidden layers, each of which has 500 neurons with the ReLU activation. The output layer consists of 10 neurons with the softmax activation.
- **C1:** A modified version of the VGG configuration (Simonyan & Zisserman, 2014). Following Keskar et al. (2017), let us denote a stack of n convolutional layers of a filters and a kernel size of $b \times c$ with the stride length of d by $n \times [a, b, c, d]$. The C1 network uses the configuration: $3 \times [64, 3, 3, 1]$, $3 \times [128, 3, 3, 1]$, $3 \times [256, 3, 3, 1]$, where a MaxPool(2) is applied after each stack. To all layers, the ghost-batch normalization of size 100 and the ReLU activation are applied. Finally, an output layer consists of 10 neurons with the softmax activation.




Figure 1: Minibatch-size dependence of the test accuracy (left) and the convergence time (right) for each fixed value of α in C1.

Table 2: Best test accuracy for each value of α .

Name	α	B_{opt}	test accuracy (%)	convergence time
C1	1	100	88.05 ± 0.18	24500 ± 2775
	1.5	200	88.69 ± 0.11	20825 ± 2113
	2.0	300	88.77 ± 0.30	15932 ± 1870
	2.5	500	88.66 ± 0.22	10040 ± 1153
F1	1	900	90.17 ± 0.14	10934 ± 816
	1.5	2000	90.39 ± 0.17	7914 ± 528
C2	1	600	61.40 ± 0.54	5292 ± 935
	1.5	1000	61.75 ± 0.48	5175 ± 748

- **C2**: It is similar to but larger than C1. The C2 network uses the configuration: $3 \times [64, 3, 3, 1]$, $3 \times [128, 3, 3, 1]$, $3 \times [256, 3, 3, 1]$, $2 \times [512, 3, 3, 1]$, where a MaxPool(2) is applied after each stack. To all layers above, the ghost-batch normalization of size 100 and the ReLU activation are applied. The last stack above is followed by a 1024-dimensional dense layer with the ReLU activation, and finally, an output layer consists of 10 neurons with the softmax activation.

For all experiments, we used the cross-entropy loss and the Adam optimizer with the default hyperparameters. Neither data augmentation nor weight decay is applied in our experiment. To aid the convergence, we halve the learning rate when the training loss reaches the value L^* . Training finishes when the training loss becomes smaller than the value L^{**} . Our choices of L^* and L^{**} are also described in Table 1. The convergence time is defined as the number of iteration steps until the training finishes. Training is repeated 10 times starting from different random initializations (the Glorot initialization is used), and we measure the mean test accuracy and the mean convergence time as well as their standard deviations.

3.1 EFFECT OF THE NOISE ENHANCEMENT

First we demonstrate how the noise enhancement affects the generalization and the convergence time for C1 (similar results are obtained for F1 and C2 as we show later). For each fixed value of $\alpha = 1, 1.5, 2.0, 2.5$ ($\alpha = 1$ means no NE applied) we calculated the mean test accuracy and the mean convergence time for varying minibatch sizes B . The result is presented in Fig. 1. We can see that the NE improves generalization for a not too large α . It is also observed that the generalization gap between small-batch training and large-batch training diminishes by increasing α . The NE with large α is therefore efficient for large-batch training. On the other hand, the convergence time increases with α for a fixed B .

For each fixed α , there is an optimal minibatch size B_{opt} , which increases with α . In Table 2, we list $B_{\text{opt}} \in \{100, 200, 300, 400, 500, 600, 700, 800, 900, 1000, 2000, 3000, 5000\}$ as well as the test

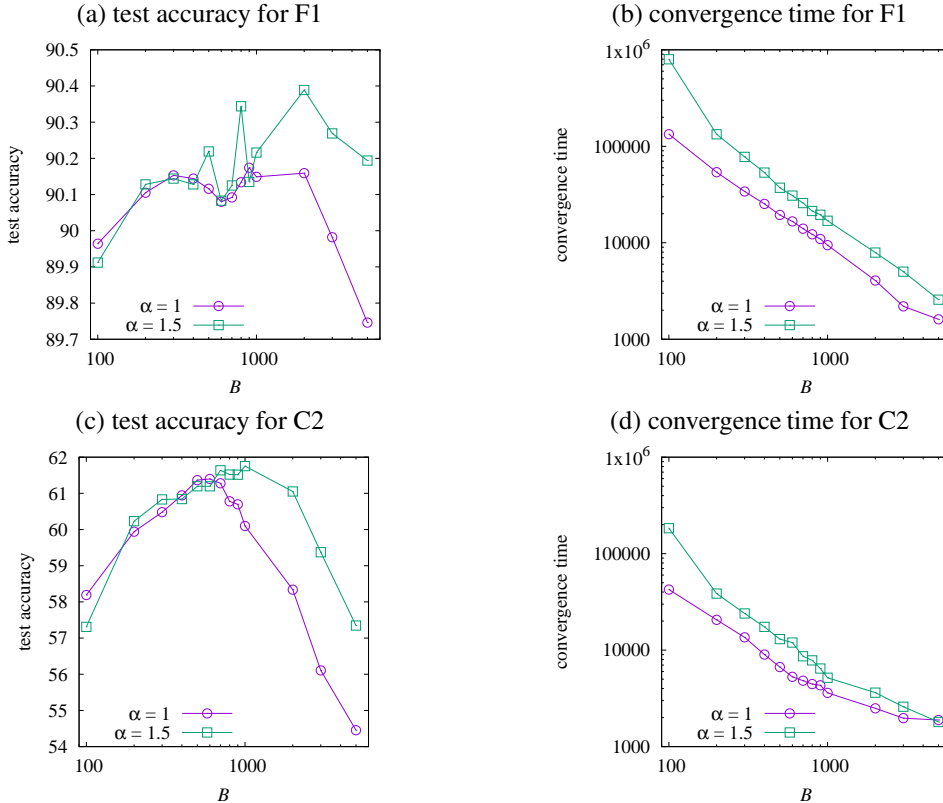


Figure 2: Minibatch-size dependence of the test accuracy and the convergence time for each fixed value of α in F1 and C2.

accuracy and the convergence time at $B = B_{\text{opt}}$. We see that the test accuracy at B_{opt} is improved by the NE. Moreover, the NE shortens the convergence time at B_{opt} without hurting generalization performance.⁴ This experimental observation shows practical efficiency of the method of the NE.

Although we have focused on C1, other configurations F1 and C2 also show similar results. For F1 and C2, we compare the result for $\alpha = 1$ with that for $\alpha = 1.5$. In Fig. 2, the minibatch-size dependences of the test accuracy and the convergence time are shown for F1 and C2. In Table 2, we also show the test accuracy and the convergence time at $B = B_{\text{opt}}$ for each α in F1 and C2. These results are qualitatively same as those in C1 (Fig. 1 and Table 2).

3.2 COMPARISON BETWEEN THE NOISE ENHANCEMENT AND REDUCING THE MINIBATCH SIZE

It is pointed out that reducing the minibatch size B with $\alpha = 1$ has a similar effect as the NE with a fixed B ; it results in better generalization but a longer convergence time.⁵ We shall compare the large-batch training with the NE to the small-batch training without the NE. First we calculate the test accuracy and the convergence time for varying B and a fixed $\alpha = 1$ (no NE). We then calculate the test accuracy for varying $\alpha > 1$ and a fixed $B = 5000$, which corresponds to large minibatch training. In other words, we compare the effect of the NE with that of reducing B .

The comparison between reducing B with $\alpha = 1$ and increasing α with $B = 5000$ is given in Fig. 3. We see that both give similar curves; increasing the convergence time with a peaked test accuracy.

⁴The NE for a fixed B increases the convergence time, but B_{opt} also increases, which decreases the convergence time.

⁵As was already mentioned, under the Gaussian noise approximation, increasing α is indeed equivalent to reducing B to B_{eff} given by Eq. (11).

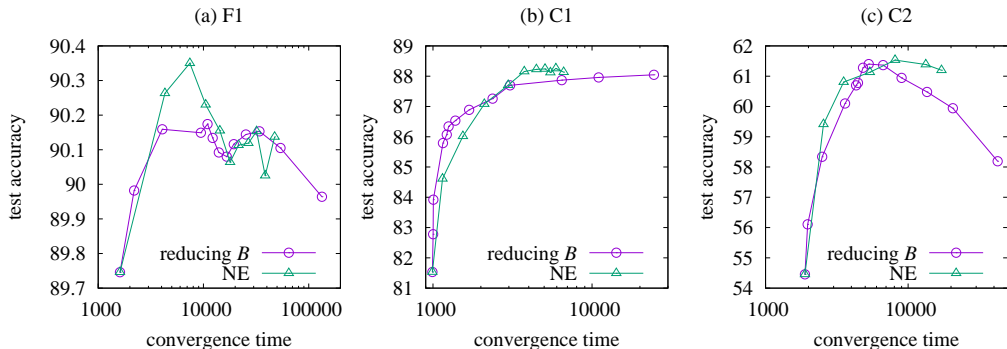


Figure 3: Comparison between the effects of reducing the minibatch size B with $\alpha = 1$ and of increasing α with $B = 5000$. The longitudinal axis and the horizontal axis represent the test accuracy and the convergence time, respectively. Circle data points (reducing B with $\alpha = 1$) correspond to $B = 5000, 3000, 2000, 1000, 900, 800, 700, 600, 500, 400, 300, 200, 100$ from left to right. Triangle data points (increasing α with $B = 5000$) correspond to $\alpha = 1, 2, \dots, 11$ for F1 and C1, and $\alpha = 1, 2, \dots, 7$ for C2, from left to right.

Table 3: Comparison of best test accuracies for varying B with $\alpha = 1$ (without the noise enhancement) and for varying α with $B = 5000$ (with the noise enhancement). The range of varying B and α is the same as in Fig. 3.

Name		B	α	Best test accuracy (%)
F1	without NE	900	1	90.17 ± 0.14
	with NE	5000	3	90.35 ± 0.05
C1	without NE	100	1	88.05 ± 0.18
	with NE	5000	10	88.26 ± 0.23
C2	without NE	600	1	61.40 ± 0.54
	with NE	5000	5	61.53 ± 0.35

However, in every case of F1, C1, and C2, the NE (increasing α) results in better accuracy compared with reducing B if α is properly chosen.

In Table 3, we compare the best test accuracies between varying B with $\alpha = 1$ (without the NE) and increasing α with $B = 5000$ (with the NE). In all cases, the large-batch training with the NE outperforms the small-batch training without the NE.

4 DISCUSSION

We have shown that the method of the NE for gradient-based optimization algorithms improves generalization. In particular, large-batch training with the NE even outperforms small-batch training without the NE, which clearly shows that the NE is not equivalent to reducing the minibatch size B .

In this section, we shall discuss two fundamental questions raised here:

- (i) *Why does a stronger SGD noise result in a better generalization?*
- (ii) *How is the inequivalence between the NE and reducing B theoretically understood?*

We first consider (i). When the SGD noise strength is inhomogeneous in the parameter space, network parameters will be likely to evolve to a minimum of the loss landscape with a weaker SGD noise.⁶ That is, if the SGD noise is strong enough near a minimum, the network parameters will easily escape from it with the help of the SGD noise. As a result, only minima around which the

⁶In physics, similar phenomena are known; Brownian particles in a medium with inhomogeneous temperature tend to gather in a colder region (Soret effect) (Duhr & Braun, 2006; Sancho, 2015).

SGD noise is weak enough survive. Since the covariance of the SGD noise is given by Eq. (4), or Eq. (10) for the NE, the strong SGD noise is considered to have an implicit regularization effect toward minima with a small variance of $\{\nabla_w \ell_\mu\}$. Some previous studies have introduced various measures which express an implicit regularization effect of SGD (Keskar et al., 2017; Yin et al., 2018; Wu et al., 2018). Among them, the “gradient diversity” introduced by Yin et al. (2018) is closely related to the above argument.

A small variance of the sample-dependent gradients $\{\nabla_w \ell_\mu\}$ around a minimum of the loss function implies that the loss landscape $L_{\mathcal{B}}(w)$ for a minibatch \mathcal{B} does not largely depend on \mathcal{B} . Such a minimum would contain information on common features among training data samples, which would be relevant for a given classification, but not contain information on sample-specific features which lead to overfitting. This is our intuitive picture that explains why the strong SGD noise results in good generalization performance.

The above consideration is solely based on Eq. (4), i.e., the covariance structure of the SGD noise, and the effect of non-Gaussian noise has been ignored. However, when the SGD noise is strengthened by reducing B , the SGD noise deviates from Gaussian and the above argument should be somehow modified. As we have already mentioned, the inequivalence between the NE and reducing B results from the non-Gaussian nature of the SGD noise, which is therefore a key ingredient to answer the question (ii). The method of the NE can increase the noise strength without changing B , and hence it is considered to suppress the non-Gaussianity compared with the case of just reducing B . The experimental result presented in Sec. 3 then indicates that the non-Gaussian nature of the SGD noise has a negative impact on generalization. A possible interpretation is that sample-specific features show up and are overestimated, which results in overfitting, when the central limit theorem is strongly violated.⁷ However, the relation between the non-Gaussianity of the SGD noise and generalization remains unclear (Wu et al., 2019), and it would be an important future problem to make this point clear.

In this way, we now have intuitive arguments which might be relevant to answer the questions (i) and (ii), but theoretical solid explanations are still lacking. Our results will not only be useful in practice, but also give theoretical insights into those fundamental questions, which merit further study.

REFERENCES

- Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. Stronger generalization bounds for deep nets via a compression approach. In *International Conference on Machine Learning*, 2018. doi: arXiv.
- Sanjeev Arora, Simon S. Du, Wei Hu, Zhiyuan Li, Ruslan Salakhutdinov, and Ruosong Wang. On Exact Computation with an Infinitely Wide Neural Net. In *Neural Information Processing Systems*, 2019.
- Stéphane D’Ascoli, Maria Refinetti, Giulio Biroli, and Florent Krzakala. Double Trouble in Double Descent : Bias and Variance(s) in the Lazy Regime. *arXiv preprint arXiv:2003.01054*, 2020.
- Stefan Duhr and Dieter Braun. Thermophoretic depletion follows boltzmann distribution. *Physical Review Letters*, 96(16):168301, 2006.
- Gintare Karolina Dziugaite and Daniel M Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *Uncertainty in Artificial Intelligence*, 2017.
- Noah Golmant, Nikita Vemuri, Zhewei Yao, Vladimir Feinberg, Amir Gholami, Kai Rothauge, Michael W. Mahoney, and Joseph Gonzalez. On the Computational Inefficiency of Large Batch Sizes for Stochastic Gradient Descent. *arXiv preprint arXiv:1811.12941*, 2018.

⁷At a certain stage of training, some training data samples have been confidently classified correctly but others have not. This fact suggests that the distributions of $\ell_\mu(w)$ and $\nabla_w \ell_\mu(w)$ have a long tail and that the variance of $\{\nabla_w \ell_\mu\}$ is not small enough to justify the central limit theorem unless B is sufficiently large. Indeed, Panigrahi et al. (2019) have demonstrated that the SGD noise looks Gaussian only in an early stage of training for a not too large B .

- Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour. *arXiv preprint arXiv:1706.02677*, 2017.
- Elad Hoffer, Itay Hubara, and Daniel Soudry. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. In *Advances in Neural Information Processing Systems*, 2017.
- Elad Hoffer, Tal Ben-Nun, Itay Hubara, Niv Giladi, Torsten Hoefer, and Daniel Soudry. Augment your batch: better training with larger batches. *arXiv preprint arXiv:1901.09335*, 2019.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, 2018.
- Stanisław Jastrzębski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three Factors Influencing Minima in SGD. *arXiv preprint arXiv:1711.04623*, 2017.
- Nitish Shirish Keskar, Jorge Nocedal, Ping Tak Peter Tang, Dheevatsa Mudigere, and Mikhail Smelyanskiy. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations*, 2017.
- Alex Krizhevsky. One weird trick for parallelizing convolutional neural networks. *arXiv preprint arXiv:1404.5997*, 2014.
- Qianxiao Li, Cheng Tai, and E. Weinan. Stochastic modified equations and adaptive stochastic gradient algorithms. In *International Conference on Machine Learning*, 2017.
- Vaishnavh Nagarajan and J. Zico Kolter. Generalization in Deep Networks: The Role of Distance from Initialization. In *Advances in Neural Information Processing Systems*, 2017.
- Behnam Neyshabur, Srinadh Bhojanapalli, David Mcallester, and Nathan Srebro. Exploring Generalization in Deep Learning. In *Advances in Neural Information Processing Systems*, 2017.
- Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. The role of over-parametrization in generalization of neural networks. In *International Conference on Learning Representations*, 2019.
- Abhishek Panigrahi, Raghav Somani, Navin Goyal, and Praneeth Netrapalli. Non-Gaussianity of Stochastic Gradient Noise. *arXiv preprint arXiv:1910.09626*, 2019.
- Guillermo Valle Pérez, Ard A Louis, and Chico Q Camargo. Deep learning generalizes because the parameter-function map is biased towards simple functions. In *International Conference on Learning Representations*, 2019.
- J. M. Sancho. Brownian colloids in underdamped and overdamped regimes with nonhomogeneous temperature. *Physical Review E*, 92(6):062110, 2015.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Samuel L. Smith and Quoc V. Le. A Bayesian perspective on generalization and stochastic gradient descent. In *International Conference on Learning Representations*, 2018.
- Samuel L Smith, Pieter Jan Kindermans, Chris Ying, and Quoc V Le. Don't decay the learning rate, increase the batch size. In *International Conference on Learning Representations*, 2018.
- Samuel L. Smith, Erich Elsen, and Soham De. On the Generalization Benefit of Noise in Stochastic Gradient Descent. *arXiv preprint arXiv:2006.15081*, 2020.
- Jingfeng Wu, Wenqing Hu, Haoyi Xiong, Jun Huan, Vladimir Braverman, and Zhanxing Zhu. On the Noisy Gradient Descent that Generalizes as SGD. *arXiv preprint arXiv:1906.07405*, 2019.

Lei Wu, Chao Ma, and E. Weinan. How SGD selects the global minima in over-parameterized learning: A dynamical stability perspective. In *Advances in Neural Information Processing Systems*, 2018.

Dong Yin, Ashwin Pananjady, Max Lam, Dimitris Papailiopoulos, Kannan Ramchandran, and Peter L. Bartlett. Gradient diversity: A key ingredient for scalable distributed learning. In *International Conference on Artificial Intelligence and Statistics*, 2018.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding Deep Learning Requires Rethinking of Generalization. In *International Conference on Learning Representations*, 2017.

Zhanxing Zhu, Jingfeng Wu, Bing Yu, Lei Wu, and Jinwen Ma. The anisotropic noise in stochastic gradient descent: Its behavior of escaping from sharp minima and regularization effects. In *International Conference on Machine Learning*, 2019.