



Introducing Noise in Decentralized Training of Neural Networks

Linara Adilova^{1,2(✉)}, Nathalie Paul¹, and Peter Schlicht³

¹ Fraunhofer IAIS, Sankt Augustin, Germany

{[linara.adilova](mailto:linara.adilova@iais.fraunhofer.de),[nathalie.paul](mailto:nathalie.paul@iais.fraunhofer.de)}@iais.fraunhofer.de

² Fraunhofer Center for Machine Learning, Sankt Augustin, Germany

³ Volkswagen Group Research, Wolfsburg, Germany

peter.schlicht@volkswagen.de

Abstract. It has been shown that injecting noise into the neural network weights during the training process leads to a better generalization of the resulting model. Noise injection in the distributed setup is a straightforward technique and it represents a promising approach to improve the locally trained models. We investigate the effects of noise injection into the neural networks during a decentralized training process. We show both theoretically and empirically that noise injection has no positive effect in expectation on linear models, though. However for non-linear neural networks we empirically show that noise injection substantially improves model quality helping to reach a generalization ability of a local model close to the serial baseline.

1 Introduction

The idea of noisy optimization originates from the physical process of annealing [14], where noise helps to stabilize the state of the system through changing it by small random perturbations. Injecting noise in various ways is exploited both in convex problems [3] and non-convex ones, specifically, in the training of neural networks [1]. Research empirically shows that noise injection into the neural network objective optimization benefits the training process and improves generalization [7, 19]. In our work we consider the particular case of noise injection into the model parameters, i.e., the neural network weights.

Optimization of the objective function for various tasks can also be performed in a decentralized manner, i.e., multiple learners train on distributed data sources and synchronize according to a chosen schedule [11, 13, 17]. Decentralized training has a strong motivation coming from a growing amount of devices, e.g., mobile phones, with possibly privacy sensitive data sources and the ability to perform local computations. However, models obtained via synchronizing results of local training often do not reach the performance of serial baseline trained on a centralized dataset.

L. Adilova and N. Paul—These authors contributed equally.

© Springer Nature Switzerland AG 2019

A. Monreale et al. (Eds.): ECML PKDD 2018 Workshops, CCIS 967, pp. 37–48, 2019.

https://doi.org/10.1007/978-3-030-14880-5_4

In order to tune the quality of the distributed training setup various synchronizing protocols were investigated, which address communication timing and aggregating operators [11–13]. We consider using an existing approach for improving serial training, namely noise injection, for the case of decentralized setup in order to reach higher accuracies of the local models.

We give an overview of the related research in Sect. 2. For prior theoretical investigation we prove in Sect. 3 that in the special case of linear neural networks zero-mean noise does not have an adversarial effect on the results of training, because it cancels out in expectation. This is supported by experiments in Sect. 4.1. Since non-linear neural networks have more practical usage, in Sect. 4.2 we empirically study noise injection to non-linear neural networks trained in the decentralized manner and show that it improves performance. The simplest case of noise injection to the neural network weights is initialization noise. Experiments of McMahan et al. [17] suggest that averaging of independently trained neural networks on different local datasets results in a model that performs worse than any model of the local learners. However, when averaging periodically, our experiments show that noisy initialization is indeed beneficial. Further experiments in Sect. 4.2 are concerned with noise injection during each synchronization step and show an improved quality of the trained models compared to the non-noisy setup for two considered classification tasks. Section 5 summarizes the research results and suggests possible future work.

2 Related Work

Various ways of noise injection into the process of neural networks training were thoroughly investigated in multiple research works. For example, Bishop [4] has shown that additive noise on the inputs is equivalent to a regularization term in a loss function if the noise amplitude is sufficiently small. An [1] explored noise injection to the inputs, outputs, and weights of neural networks. He observed that additive noise injection to the weights on each update step, either in on-line or in batch gradient descent optimization, leads to higher generalization of the learned solution. It was shown that noise added to the inputs affects only the smoothness of the resulting function, while noise injected to the weights also punishes large values and activations. Injected output noise changes the loss function only by a constant value and thus does not affect the quality of the learned function. Later Wang and Principe [22] demonstrated that noise added to the desired signal affects the variance of weights and leads to faster convergence. A different approach to noise injection to the training process is disturbing the gradient updates. It is investigated, for example, in works of An [1] or Neelakantan et al. [19]. Empirical results show that adding decaying noise to gradients helps achieving a global minimum, but does not affect the generalization quality of the resulting network. Yet reaching a global minimum might be even harmful for generalization, since such solutions tend to overfit on the training data.

Decentralized training of neural networks on distributed data sources with different ways of synchronization is investigated in works of McMahan et al.

[17], Jiang et al. [10], Smith et al. [20] and some others. The issue closest to our research is the initialization aspect considered by McMahan et al. [17]. The process of the neural network objective optimization is sensitive to the initial state of the model as it was presented by Kolen and Pollack [15]. According to McMahan et al. [17] the averaging is harmful if initialization of the locally trained models is different, which corresponds to injection of random noise to initially equal states. Nevertheless the effect of initialization noise together with periodic averaging was not investigated.

3 Noisy Averaging

To start our theoretical investigation we consider the simplest case of models, that is linear models, and analyze the effect of noise injection into their parameters for a common training setup. Here we present the summarization of used basic concepts and a proof that noise injection in the considered special case for linear models has no detrimental effect on the final model.

3.1 Basic Concepts

The task of training a neural network is an optimization problem which is commonly solved by using *stochastic gradient descent* (SGD) or SGD-based algorithms. SGD is an on-line optimization algorithm in which a local minimum of the objective function is determined by moving in the negative direction of its gradient [6]. SGD is derived from the gradient descent algorithm and estimates the gradient by considering only one training example. Let X denote the *input space*, Y the *output space*, \mathcal{F} the *model parameters space* and $f \in \mathcal{F}$ the model parameters. The SGD update rule for minimizing an *objective function* $\ell : \mathcal{F} \times X \times Y \rightarrow \mathbb{R}_+$ reads

$$\varphi_{\eta}^{SGD}(f) = f - \eta \nabla \ell(f, x, y),$$

where the learning rate η controls the step size.

In our work we consider *mini-batch SGD* which approximates the gradient by taking batch size B many training examples into account:

$$\varphi_{B,\eta}^{mSGD}(f) = f - \eta \sum_{j=1}^B \nabla \ell^j(f), \quad (1)$$

where $\ell^j(f) = \ell(f, x_j, y_j)$.

An example of a loss function is *squared loss* $\ell^j(f) = 1/2 \left(\langle f, x_j \rangle - y_j \right)^2$.

Choosing this loss function and computing its gradient, we can write the learning algorithm update rule in the following form:

$$\varphi_{B,\eta}^{mSGD}(f) = f - \eta \sum_{j=1}^B \left(\left(\langle f, x_j \rangle - y_j \right) x_j \right).$$

3.2 Periodic Averaging Protocol with Noise Injection

As a protocol for the theoretical investigation of the effect of noise injection in the decentralized setup we consider the periodic averaging protocol with zero-mean noise injection (Algorithm 1). Let Ψ denote a probability distribution with mean zero and $\epsilon_t \geq 0$ the time-dependent noise level factor.

Algorithm 1. Periodic averaging protocol with zero-mean noise injection

Input: batch size b

Initialization:

local models $f_1^1, \dots, f_1^m \leftarrow$ one random f

At round t :

At learner i :

observe $(x_{t,1}^i, y_{t,1}^i), \dots, (x_{t,B}^i, y_{t,B}^i)$

$\psi_i \leftarrow \Psi$

update f_t^i using the learning algorithm $\varphi : f_{t+1}^i \leftarrow \varphi(f_t^i + \epsilon_t \psi_i)$

If $t \bmod b = 0$:

synchronize local models: $f_{t+1}^i \leftarrow \frac{1}{m} \sum_{i=1}^m f_{t+1}^i$

In the following we analyze the influence of noise injection on the behavior of the learning algorithm for the special case of linear models. For this case we can prove that adding zero-mean noise to the parameters optimized by mini-batch SGD with squared loss does not change the model parameters in expectation.

Lemma 1. *For a linear model let f_t denote the model parameters attained by using mini-batch SGD with squared loss and scaled additive zero-mean weight noise injection, i.e.*

$$f_{t+1} = \varphi_{B,\eta}^{mSGD}(f_t + \epsilon_t \psi).$$

Let g_t denote the model parameters attained by using common mini-batch SGD, i.e.

$$g_{t+1} = \varphi_{B,\eta}^{mSGD}(g_t).$$

If the learning algorithms are identically initialized, it holds

$$\mathbb{E}[f_t] = g_t \quad \text{for all } t = 1, \dots, T. \quad (2)$$

Proof. We use induction over t . The case $t = 1$ follows immediately since by initialization the model parameters trained with and without noise are the same. Applying the definitions above, the expected model parameters for mini-batch SGD with noise injection read

$$\begin{aligned}
\mathbb{E}[f_t] &= \mathbb{E}\left[\varphi_{B,\eta}^{mSGD}(f_{t-1} + \epsilon_t \psi)\right] \\
&= \mathbb{E}\left[f_{t-1} + \epsilon_t \psi - \eta \sum_{j=1}^B \left(\left(\langle f_{t-1} + \epsilon_t \psi, x_{(t-1,j)} \rangle - y_{(t-1,j)}\right) x_{(t-1,j)}\right)\right] \\
&= \underbrace{\epsilon_t \mathbb{E}[\psi]}_{=0} + \mathbb{E}\left[f_{t-1} - \eta \sum_{j=1}^B \left(\left(\langle f_{t-1}, x_{(t-1,j)} \rangle - y_{(t-1,j)}\right) x_{(t-1,j)}\right)\right] \\
&\quad - \underbrace{\mathbb{E}\left[\eta \sum_{j=1}^B \left(\epsilon_t \langle \psi, x_{(t-1,j)} \rangle x_{(t-1,j)}\right)\right]}_{=0},
\end{aligned}$$

where we used that ψ is centered. By employing the induction assumption we get $\mathbb{E}[f_{t-1}] = g_{t-1}$ and conclude

$$\begin{aligned}
\mathbb{E}[f_t] &= g_{t-1} - \eta \sum_{j=1}^B \left(\left(\langle g_{t-1}, x_{(t-1,j)} \rangle - y_{(t-1,j)}\right) x_{(t-1,j)}\right) \\
&= \varphi_{B,\eta}^{mSGD}(g_{t-1}) \\
&= g_t.
\end{aligned}$$

□

The equivalence of Algorithm 1 to the non-noisy periodic averaging protocol in expectation directly follows from the linearity of the expected value.

Corollary 1. *For linear models trained with mini-batch SGD and squared loss given identical initialization, the expected model obtained by the periodic averaging protocol with zero-mean noise injection is equivalent to the model obtained by the non-noisy periodic averaging protocol.*

For the specific considered case of linear models we observe that noise injection has no influence on the model parameters in expectation both for serial and distributed learning. This theoretical result is supported by empirical evidence in the following Section. In contrast, for non-linear models the research of An [1], Edwards and Murray [7] has shown that noise injection in serial training helps improving generalization. We conjecture that in the distributed training setup injecting noise into non-linear models might also improve generalization properties of the obtained solution compared to the non-noisy training. We leave a theoretical investigation of the effect of noise injection into non-linear models in distributed training setup for future work. To substantiate our conjecture, in the following Section we perform an empirical analysis of zero-mean noise injection in a distributed setup for non-linear neural networks.

4 Empirical Evaluation

To investigate the effect of noise injection into the neural networks trained in a decentralized manner we performed a set of experiments described further.

The decentralized setup of these experiments is periodic synchronization via averaging (cf. Algorithm 1). Apart from the distributed synchronized models two baselines are trained: a local model without any synchronization with other local learners (no-sync) and a model with full data centralization (serial). These baselines are necessary to assess the performance of the synchronizing local learners, since synchronization aims to reach the performance of the serial model and outperform the no-sync baseline. Noise injection into the serial baseline is also a subject of interest, allowing to compare the possible gains in generalization ability in centralized and distributed case. For evaluation of the synchronizing learners the last averaged model obtained during the training process is used. For evaluation of the no-sync baseline we pick one random model among locally trained ones.

Since we explore the effect of noise injection, we are interested in the behavior of models trained throughout several experiments that differ only by the used noise. Thus all the setups were run 10 times without using fixed random seed. This produces an indication of how distribution of possible outcomes of the training process looks like. The results are presented in the form of box plots. Here the box shows the observed values from the first to the third quartile, with a line at the median. The whiskers show the range of the results and points are representing outliers.

4.1 Linear Neural Networks

First set of experiments is performed to empirically evaluate the theoretical result obtained in Sect. 3. We employed a linear neural network with three layers having 2, 64 and 1 neuron correspondingly for approximating the target column of SUSY dataset [2]. For linear model experiments the dataset was normalized thus having -1 and 1 as targets and accuracy was calculated with 0 threshold. The optimal training parameters determined were mini-batch of $B = 10$ examples, learning rate of $\eta = 1e-5$ for the serial and no-sync baselines and $\eta = 2.5e-5$ for the local learners. We employed squared loss as explained in Sect. 3. During the training each local learner was presented 20000 examples, while serial baseline had $m \cdot 20000$ examples with $m = 10$. The noise used for this experiment is additive uniform noise in the range $[-0.5, 0.5]$ with decay factor equal to the synchronization round number. Absence of decay factor was leading to fast overflow thus making experiments not runnable.




Fig. 1. Effect on test accuracy of noise injection to local learners and to the serial baseline.

Figure 1 shows the test accuracy evaluated for 10 runs for the baselines and synchronizing learners on the independent test set of 1000000 examples. We can observe that larger noise injected into the weights of the models leads to larger variance of obtained accuracy at the same moment leaving the median value throughout different setups the same. This empirically supports the effect of noise cancellation in expectation for linear models in this training setup.

4.2 Non-linear Neural Networks

In the following we evaluate noise injection to the decentralized training of non-linear neural networks on the basis of two classification tasks. For our experiments we choose to add uniform noise in range $[-0.5, 0.5]$ and Gaussian noise in the same range.

Binary Classification. For preliminary evaluation of the approach for non-linear case we have considered the classification task on the SUSY dataset. In contrast to the linear case, the employed model architecture is a three-layered dense network with sigmoid activations. The first layer has 32 neurons, the second 64 and the output layer has 2 neurons with softmax activation. We have determined the optimal parameters of the non-noisy learning algorithm on a small fraction of the dataset. That is training mini-batch of $B = 10$ examples, learning rate of $\eta = 0.1$ for the serial and no-sync baselines and $\eta = 0.25$ for the local learners.

Initialization Noise. The simplest way of noise injection to the distributed neural networks is one step noise injection right after initialization. In McMahan et al. [17] it was shown that such noise together with one-time synchronization after local training results in a worse model than each local one in terms of training loss. We want to investigate whether periodic synchronization is capable of being more robust to initial noise injection.

With regard to Algorithm 1, initialization noise corresponds to choosing $\epsilon_1 > 0$ and $\epsilon_t = 0$ for all $t > 1$. It means that we add randomly sampled noise to each initial weight before starting the training process.

Figure 2 shows that both Gaussian and uniform initialization noise improves the serial and periodically synchronizing models in terms of cumulative training loss when using it up to some small extent ($\epsilon_t < 1$). On the contrary, higher levels of noise (e.g. $\epsilon_t = 5$) make training harder in each setup. Interestingly, a noise level of $\epsilon_t = 2$ for both Gaussian and uniform noise leads to a higher cumulative training loss in the serial model while the distributed setup benefits from it.

Even though the Gaussian distribution is a very popular choice for initializing neural networks (Glorot and Bengio [9]), adding large levels noise drawn from the normal distribution deteriorates the training process worse than uniformly distributed noise. One possible explanation is that the initial weights are already distributed normally according to the best practices and additional Gaussian noise intervenes with it in a destructive way.

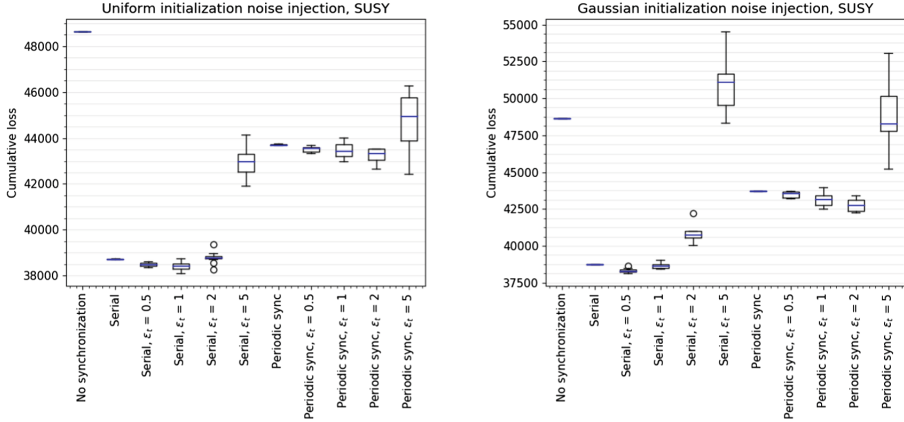


Fig. 2. Effect on cumulative training loss of uniformly and normally distributed noise added to the equal initializations of local learners and to the serial baseline.

Initialization noise experiments reveal that one-time noise injection to the initially equal model weights helps the training process since the cumulative training loss is decreasing. This motivates follow-up experiments with continuous noise injection.

Continuous Noise. Extending the initialization noise setup we now perform additional noise injection steps: Zero-mean noise gets injected to the local models’ weights after every synchronization step. Formally, in Algorithm 1 this corresponds to setting $\epsilon_1 > 0$, and $\epsilon_{t+1} > 0$ for all $t \bmod b = 0$ and $\epsilon_t = 0$ otherwise. Following the work of Murray and Edwards [18] the noise is decaying and the decaying factor is equal to the index number of the synchronization step, i.e. noise level is given by ϵ_t/t .

We want to explore whether continuous noise injection improves the generalization ability of the resulting models in the distributed setup. Therefore we calculated the evaluation accuracy on an independent test set of 1000000 examples for each of the trained models. During training, each of the local learners i is presented 1000 examples from the training dataset, while the serial model sees $m \cdot 1000$ examples. Various setups together with evaluated validation accuracies are depicted as box plots in Fig. 3.

The evaluation shows that noise injection can substantially improve the generalization ability of the models. When comparing results of the setup with 10 and 20 learners, we see that a larger amount of learners leads to a larger spread of the results of the training process that can mean either a better generalization or to the contrary convergence to a worse model.

One can also observe that with growth of the uniform noise level the resulting test accuracy becomes more unstable, i.e., for having a possibly higher median we get a larger range of values. In the experiments, Gaussian noise is more stable than uniform noise while on the other hand it has more outliers below

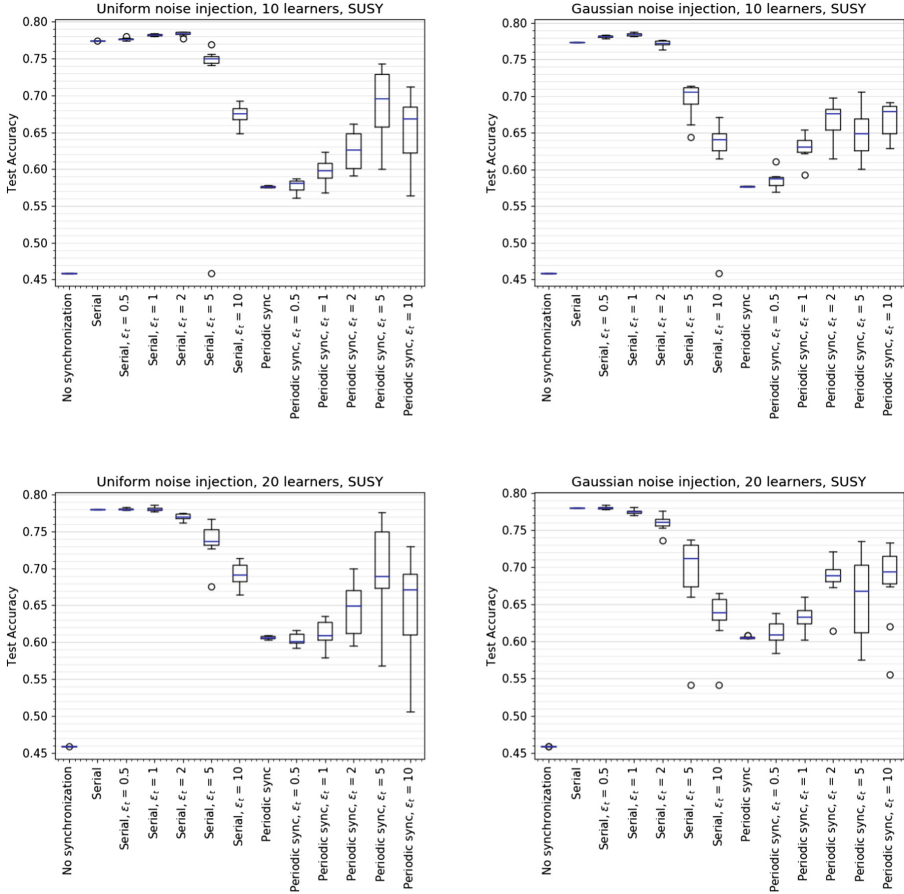


Fig. 3. Effect on test accuracy of injecting uniformly or normally distributed noise every b time steps throughout the training process of local learners and the serial baseline.

the median. The spread of the serial baseline with noise injection is very small compared to the distributed models. It might be interesting to investigate the reasons why noise has a more pronounced effect in the distributed setup than in the serial one.

Image Classification. To further investigate the effect of noise injection into non-linear models in a distributed setup we have chosen the classification task on the MNIST dataset [16]. The model architecture is more complicated than in the previous experiment. It has two dense layers with 512 neurons each and a dropout layer after each of them. The output layer performs a softmax activation to predict one of the ten classes. The activation of the dense layers is ReLU. We have determined the optimal parameters of the non-noisy learning algorithm




Fig. 4. Effect on test accuracy of injecting uniformly and normally distributed noise every b time steps throughout the training process of local learners and the serial baseline.

on a small fraction of the dataset. That is, equivalently to the first experiment, training mini-batch of $B = 10$ examples, learning rate of $\eta = 0.1$ for the serial and no-sync baselines and $\eta = 0.25$ for the local learners. During training each learner is presented 500 examples from the training set and evaluation is performed on the 10000 images of the test set.

We observe that in this experiment, compared to the previous, noise injection requires different level values in order to obtain the improved quality of the models (Fig. 4). More precisely, for noise levels $\epsilon_t \geq 1$ we get prohibitively low test accuracies for both uniform and Gaussian noise. One possible explanation is that the ReLU activation is much more sensitive to noise disturbance of the weights compared to the bounded sigmoid function which we have used in the first experiment. Moreover the dropout layers might be contributing to this effect since dropout is also supposed to prevent a model from overfitting [21]. More refined research on this effect is left for future work. Concentrating on lower levels of noise ($\epsilon_t < 0.5$), we see that we can again improve the generalization ability of the trained models. In the distributed setup this effect is more pronounced than in the serial one similarly to the first set of experiments.

5 Conclusion

The research presented in this paper investigates noise injection into neural networks, trained in a decentralized manner on distributed data sources. We have proven that for linear models in a common training setup zero-mean noise injection retains the results of the non-noisy setup. We have performed experiments to empirically underline this theoretical statement. Further experiments show that with non-linear models in a distributed setup noise injection improves the quality of the models. The evaluation shows that indeed carefully chosen levels of noise have a positive effect on the generalizing abilities of the synchronized models. It might be explained by the fact that noise enforces wider exploration of the space of solutions [5, 8] which is an interesting subject for further investigation. Also, experiments show that the impact of noise in the distributed training is even greater than in the serial case.

Future research could investigate the theoretical background of noise injection to non-linear models in a decentralized training setup as well as the effect of various network architectures and training parameters. A promising framework for studying noise injection effects is regularization theory, since injected noise can be described as a regularization term that is added to the loss function [1, 4].

References

1. An, G.: The effects of adding noise during backpropagation training on a generalization performance. *Neural Comput.* **8**(3), 643–674 (1996)
2. Baldi, P., Sadowski, P., Whiteson, D.: Searching for exotic particles in high-energy physics with deep learning. *Nat. Commun.* **5**, 4308 (2014)
3. Ben-Tal, A., Nemirovski, A.: Robust convex optimization. *Math. Oper. Res.* **23**(4), 769–805 (1998)
4. Bishop, C.M.: Training with noise is equivalent to Tikhonov regularization. *Neural Comput.* **7**(1), 108–116 (1995)
5. Chen, B., Deng, W., Du, J.: Noisy softmax: improving the generalization ability of DCNN via postponing the early softmax saturation. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
6. Chung, K.L., et al.: On a stochastic approximation method. *Ann. Math. Stat.* **25**(3), 463–483 (1954)
7. Edwards, P.J., Murray, A.F.: Fault-tolerance via weight-noise in analogue VLSI implementations—a case study with Epsilon. *IEEE Proc. Circ. Syst. II: Analog Digit. Sig. Process.* **45**(9), 1255–1262 (1998)
8. Gan, Z., Li, C., Chen, C., Pu, Y., Su, Q., Carin, L.: Scalable Bayesian learning of recurrent neural networks for language modeling. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics: Long Papers*, vol. 1, pp. 321–331 (2017)
9. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 249–256 (2010)
10. Jiang, Z., Balu, A., Hegde, C., Sarkar, S.: Collaborative deep learning in fixed topology networks. In: *Advances in Neural Information Processing Systems*, pp. 5904–5914 (2017)

11. Kamp, M., Boley, M., Keren, D., Schuster, A., Sharfman, I.: Communication-efficient distributed online prediction by dynamic model synchronization. In: Calders, T., Esposito, F., Hüllermeier, E., Meo, R. (eds.) ECML PKDD 2014. LNCS (LNAI), vol. 8724, pp. 623–639. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44848-9_40
12. Kamp, M., Boley, M., Missura, O., Gärtner, T.: Effective parallelisation for machine learning. In: Advances in Neural Information Processing Systems, pp. 6477–6488 (2017)
13. Kamp, M., Bothe, S., Boley, M., Mock, M.: Communication-efficient distributed online learning with kernels. In: Frasconi, P., Landwehr, N., Manco, G., Vreeken, J. (eds.) ECML PKDD 2016. LNCS (LNAI), vol. 9852, pp. 805–819. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46227-1_50
14. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by simulated annealing. *Science* **220**(4598), 671–680 (1983)
15. Kolen, J.F., Pollack, J.B.: Back propagation is sensitive to initial conditions. In: Advances in Neural Information Processing Systems, pp. 860–867 (1991)
16. LeCun, Y.: The MNIST database of handwritten digits (1998). <http://yann.lecun.com/exdb/mnist/>
17. McMahan, H.B., Moore, E., Ramage, D., Hampson, S., et al.: Communication-efficient learning of deep networks from decentralized data. arXiv preprint [arXiv:1602.05629](https://arxiv.org/abs/1602.05629) (2016)
18. Murray, A.F., Edwards, P.J.: Enhanced MLP performance and fault tolerance resulting from synaptic weight noise during training. *IEEE Trans. Neural Netw.* **5**(5), 792–802 (1994)
19. Neelakantan, A., et al.: Adding gradient noise improves learning for very deep networks. arXiv preprint [arXiv:1511.06807](https://arxiv.org/abs/1511.06807) (2015)
20. Smith, V., Chiang, C.K., Sanjabi, M., Talwalkar, A.S.: Federated multi-task learning. In: Advances in Neural Information Processing Systems, pp. 4424–4434 (2017)
21. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**(1), 1929–1958 (2014)
22. Wang, C., Principe, J.C.: Training neural networks with additive noise in the desired signal. *IEEE Trans. Neural Netw.* **10**(6), 1511–1517 (1999)