

---

# SGD Implicitly Regularizes Generalization Error

---

**Daniel A. Roberts**  
Facebook AI Research  
New York, NY 10003  
danr@fb.com

## Abstract

We derive a simple and model-independent formula for the change in the generalization gap due to a gradient descent update. We then compare the change in the test error for stochastic gradient descent to the change in test error from an equivalent number of gradient descent updates and show explicitly that stochastic gradient descent acts to regularize generalization error by decorrelating nearby updates. These calculations depends on the details of the model only through the mean and covariance of the gradient distribution, which may be readily measured for particular models of interest. We discuss further improvements to these calculations and comment on possible implications for stochastic optimization.

## 1 Introduction

Deep learning models are almost all trained using gradient-based learning algorithms [1]. Remarkably, these algorithms continue to work well even in cases where the models have far more free parameters than the number of training examples that are used to learn their value. While it's not shocking that such models with many degrees of freedom can learn on a fixed training set [2], it's surprising that they also generalize exceedingly well on novel examples.

Furthermore, stochastic gradient descent (SGD) [3], in which only a small subset or mini-batch of examples is used to estimate the gradient, is not only known to lead to configurations that generalize well [4, 5], but in many of these large-scale modeling scenarios will far outperform full-batch gradient descent (GD) despite estimating the gradient with fewer examples [6, 7].

A central question is whether this effectiveness is due to some very general property of deep learning or instead whether it's due to the dynamics of the learning algorithm used to train them. It seems that these algorithms must employ some kind of *implicit regularization* [8], since such generalization performance does not really degrade as model size increases.

The prevailing explanation given for this empirical observation has to do with descriptions of the loss landscape [6] (also see e.g. [7, 9–13]), in particular that SGD prefers “flat” to “sharp” minima in which the loss doesn't change too much in the neighborhood of the minima. An argument originating from [14] ascribes good generalization properties to such flat minima, reasoning that the effect of changing inputs could be reinterpreted as shifting the location of the minima. (Although due to the activations having rescaling symmetries [15], this should be understood to apply only after such reparameterizations have been taken into account [16].)

A common mechanism offered for this behavior is that the stochasticity is equivalent to injecting noise into the optimization process, allowing for barriers in the landscape to be climbed and sharp minima to be avoided [6]. However, isotropic noise to full-batch GD does not help for generalization (though it might for SGD [17]?) suggesting that it's anisotropy that is important [18], and [13] argues that anisotropic noise is necessary for escaping from sharp minima. Either way, most of these

analyses rely on taking the continuum limit [10, 11, 18–21] in a way that is not necessarily justified,<sup>1</sup> and moreover assume that the covariance of the gradient does not depend strongly on the value of the model parameters. As a result, this leads many authors to conclude that some kind diffusion process in parameter space must play a fundamental role in finding minima that generalize well [10, 20].

Here, we attempt to shed light on this implicit regularization puzzle by shifting focus away from the loss landscape and instead considering the discrete update dynamics themselves. Our analysis is very general and does not make any assumptions about the model or dataset.<sup>2</sup> Instead, we derive formulas for the change in generalization error after an update and for the difference between GD and SGD gradients in terms of readily measurable quantities that characterize the gradient distribution.

Our key result is that through accumulated updates GD develops a “bias” toward overfitting, which SGD mitigates by decorrelating nearby updates.<sup>3</sup> After averaging over the datasets, we show that the SGD gradient differs from the GD gradient by the derivative of the change in the generalization error to leading order in the learning rate

$$\langle \delta g \rangle \propto \partial \langle \delta \epsilon \rangle, \quad (1)$$

where  $\langle \delta \epsilon \rangle$  is the average change in the generalization error, and  $\langle \delta g \rangle$  is the averaged difference between SGD and GD gradients. Furthermore, we show that  $\langle \delta \epsilon \rangle$  is proportional to the trace of the gradient covariance matrix  $\Sigma$

$$\langle \delta \epsilon \rangle \propto \text{tr} \Sigma, \quad (2)$$

which means we can interpret (1) as making “explicit” the implicit regularization of SGD over GD. This also means that these results would be missed in any analysis that considers a homogeneous covariance independent of the model parameters as  $\langle \delta g \rangle$  would vanish identically. Although we omit experiment results for brevity, it’s easy to check that  $\partial \text{tr} \Sigma$  tends to become very large over the course of training, and so we think that this “bias” is more central to the generalization puzzle than diffusion. The upshot is that SGD noise not only leads to diffusion, but also to drift.

In this extended abstract, we will report only on our initial theoretical results, holding experimental verification for the completed work.

## 2 Gradient Descent as Stochastic Optimization

The update equation for the parameter vector in both GD and SGD can be written as

$$\theta_{t+1}^i = \theta_t^i - \eta g_i^{(\mathcal{B})}(\theta_t), \quad g_i^{(\mathcal{B})}(\theta) \equiv \frac{1}{|\mathcal{B}|} \sum_{b=1}^{|\mathcal{B}|} g_i(\theta, x^{(b)}), \quad (3)$$

where  $\theta_t^i$  is the parameter vector after  $t$  updates with the index  $i$  ranging over all the parameters,  $g_i(\theta, x)$  is the gradient with respect to the parameters evaluated at  $\theta$  and example  $x$ , and  $\eta$  is the learning rate. We will use the shorthand  $g_i^{(\mathcal{B})}(\theta)$  to indicate the sample average of a gradient on a batch of examples  $\mathcal{B}$ . If  $\mathcal{B}$  is the entire training set, then (3) is the full-batch gradient descent update, but if  $\mathcal{B}$  is a subset of the training set, then (3) is an SGD update.

With a shift in perspective, we can treat GD and SGD on the same footing. If we consider averages over all possible realizations of the training set, then both GD and SGD are stochastic, differing just in which examples are used as estimators for the gradient. Ultimately, we are interested in learning a distribution  $p(x)$  from which our training and test examples are sampled, and we are not at all attached to the examples in our particular training set.

Furthermore, we see that the only way the input  $x$  influences the dynamics is through the batch average of the gradient  $g_i^{(\mathcal{B})}(\theta)$ . Thus, regardless of how complicated the input distribution is, we never have to characterize it directly. Instead, we can just focus on the induced gradient distribution.

In general, the gradient distribution  $p(g|\theta, x)$  will also be a complicated distribution, with the higher moments playing an important role. (The number of parameters required to describe such a

<sup>1</sup>The continuum limit implicitly assumes that the mini-batch size vanishes faster than the learning rate, despite the batch size being bounded from below by 1. For further discussion, see section 2.3.3 of [22].

<sup>2</sup>Cf. [21], which focuses on deep linear networks in the student-teacher scenario in the continuum limit.

<sup>3</sup>See [23] in which the role of “gradient diversity” for generalization is also highlighted, although it is not connected to the different dynamics between GD and SGD as we do here.

distribution will be exponential in the number of parameters required to represent the random variable  $g_i$ .) However, considering (3), we see that  $g_i^{(\mathcal{B})}(\theta)$  is a sum over many independent random variables  $g_i(\theta, x^{(b)})$ . This means that for large enough batch size the central limit theorem guarantees that the distribution of the gradient estimator  $g_i^{(\mathcal{B})}(\theta)$  will be Gaussian. Using the notation  $\langle \cdot \rangle \equiv \int dx p(x) \cdot$  to indicate an average over the data generating distribution, the gradient mean  $G_i(\theta)$  and gradient covariance matrix  $\Sigma_{ij}(\theta)$  are given by

$$G_i(\theta) \equiv \langle g_i(\theta, x) \rangle, \quad \delta^{ab} \Sigma_{ij}(\theta) \equiv \langle g_i(\theta, x^{(a)}) g_j(\theta, x^{(b)}) \rangle - G_i(\theta) G_j(\theta), \quad (4)$$

where  $\delta^{ab}$  indicates that gradients derived from different independent examples have no covariance. Then, the joint distribution of two gradient estimators  $g_i^{(\mathcal{A})}(\theta)$  and  $g_j^{(\mathcal{B})}(\theta)$  is Gaussian

$$p(g_i^{(\mathcal{A})}, g_j^{(\mathcal{B})} | \theta) \sim \mathcal{N}\left(G_i(\theta), G_j(\theta); \frac{|\mathcal{A} \cap \mathcal{B}|}{|\mathcal{A}| |\mathcal{B}|} \Sigma_{ij}(\theta)\right), \quad (5)$$

where the covariance is the covariance on a single gradient scaled by the size of the intersection of the two sets, divided by the size of both sets. Crucially, the noise  $\Sigma_{ij}$  is anisotropic and depends on  $\theta$ .

Now, let's use the distribution (5) to compute the change in generalization error incurred by a gradient descent update. A GD update induces a change in the parameters  $\delta\theta^{i(\mathcal{B})} \equiv -\eta g_i^{(\mathcal{B})}$ . Since we don't actually get to observe the loss directly, we also need a batch of samples from  $p(x)$  to make an estimator. Let  $\ell(\theta, x)$  be the loss on a particular example and define  $\ell^{(\mathcal{A})}(\theta) \equiv \frac{1}{|\mathcal{A}|} \sum_{a=1}^{|\mathcal{A}|} \ell(\theta, x^{(a)})$  as the loss estimator on a batch  $\mathcal{A}$ . Let us also define the change in the loss estimator after an update  $\delta\theta^{i(\mathcal{B})}$  as  $\delta\ell^{(\mathcal{A}, \mathcal{B})}(\theta) \equiv \ell^{(\mathcal{A})}(\theta + \delta\theta^{(\mathcal{B})}) - \ell^{(\mathcal{A})}(\theta)$ . After an update (3), the loss estimator changes by an amount

$$\delta\ell^{(\mathcal{A}, \mathcal{B})} = -\eta g^{(\mathcal{A})} \cdot g^{(\mathcal{B})} + O(\eta^2), \quad (6)$$

where we have expanded in the learning rate  $\eta$  and suppressed the dependence on  $\theta$ .

Taking an expectation over  $p(x)$ , we find that on average our loss estimator decreases to first order in  $\eta$ , with the amount determined by the second moment of the batch-gradient distribution (5)

$$\langle \delta\ell^{(\mathcal{A}, \mathcal{B})} \rangle = -\eta \left[ \|G\|^2 + \frac{|\mathcal{A} \cap \mathcal{B}|}{|\mathcal{A}| |\mathcal{B}|} \text{tr} \Sigma \right]. \quad (7)$$

The terms in the square brackets are manifestly positive; the loss manifestly decreases at order  $\eta$ . The first term, the norm squared of the mean of the gradient, is the expected decrease of an objective function from a gradient descent update. The second term, proportional to the trace of the gradient covariance matrix, is an enhancement that comes from the fact that the loss estimator and gradient estimator might have a positive correlation.

Now, if we let  $\mathcal{A}$  be a test set and  $\mathcal{B}$  be a training set of  $N$  non-overlapping examples, we find

$$\langle \delta\ell \rangle_{test} \equiv \langle \delta\ell^{(\mathcal{A}, \mathcal{B})} \rangle = -\eta \|G\|^2, \quad \langle \delta\ell \rangle_{train} \equiv \langle \delta\ell^{(\mathcal{B}, \mathcal{B})} \rangle = -\eta \|G\|^2 - \eta \frac{\text{tr} \Sigma}{N}. \quad (8)$$

Defining the change in the generalization gap to be the difference between the change in the test error compared to the training error,  $\delta\epsilon \equiv \delta\ell_{test} - \delta\ell_{train}$ , this quantity and its average are given by

$$\delta\epsilon = \eta g^{(\mathcal{B})} \cdot (g^{(\mathcal{B})} - g^{(\mathcal{A})}) + O(\eta^2), \quad \langle \delta\epsilon \rangle = \eta \frac{\text{tr} \Sigma}{N} + O(\eta^2). \quad (9)$$

The formula for the average change in the generalization gap  $\langle \delta\epsilon \rangle$  is rather sensible and even intuitive. To leading order in  $\eta$ , the difference between the test and training error is proportional to the learning rate, inversely proportional to the size of the training set, and related to the fluctuations of the gradient across different data realizations through the trace of the covariance matrix. First, if the training set was infinitely large ( $N \rightarrow \infty$ ), the training set would capture the full variance of the data generating distribution and there would be no generalization gap. Second, if there were no variance in the gradient, then the training set and test set would lead to the same updates and not overfit.

### 3 Implicit Regularization

The generalization gap from gradient descent arose because the training set is a biased estimator of the average loss, and we showed that the bias can be computed explicitly (9). In this section, we

will compare stochastic gradient descent directly with gradient descent. In order to see a meaningful difference, we need to consider more than one update.

Consider two subsequent SGD updates with minibatches  $\mathcal{B}$  and  $\mathcal{C}$  each containing  $N/2$  examples

$$\theta_1^{i(\mathcal{B})} \equiv \theta_0^i - \eta g_i^{(\mathcal{B})}(\theta_0), \quad \theta_2^{i(\mathcal{B},\mathcal{C})} \equiv \theta_1^{i(\mathcal{B})} - \eta g_i^{(\mathcal{C})}(\theta_1). \quad (10)$$

Plugging the first equation into the second and Taylor expanding, we find that the parameters after two updates  $\delta\theta_2^{i(\mathcal{B},\mathcal{C})} \equiv \theta_2^{i(\mathcal{B},\mathcal{C})} - \theta_0^i$  can be expressed as

$$\delta\theta_2^{i(\mathcal{B},\mathcal{C})} = -\eta \left[ g_i^{(\mathcal{B})} + g_i^{(\mathcal{C})} \right] + \eta^2 g_j^{(\mathcal{B})} h_{ij}^{(\mathcal{C})} + O(\eta^3), \quad (11)$$

where the gradients and Hessians are all evaluated at the initial point  $\theta_0$ ,  $h_{ij}^{(\mathcal{C})}$  is the Hessian estimator  $\frac{1}{|\mathcal{C}|} \sum_{c=1}^{|\mathcal{C}|} \partial_i \partial_j \ell(\theta, x^{(c)})$ , and we use the convention that repeated indices are summed over.

From this alone, we can see that the parameter updates for GD and SGD will be different. Let's express this in terms of the difference between the overall GD gradient and the overall SGD gradient after performing two updates so  $\delta\theta_2^{i(\mathcal{B}\cup\mathcal{C},\mathcal{B}\cup\mathcal{C})} - \delta\theta_2^{i(\mathcal{B},\mathcal{C})} \equiv -\eta \delta g_i$ . Then, making use of the fact that we can rewrite  $h_{ij} g_j = \partial_i g^2/2$ , we see that this difference in gradients is given by

$$\delta g_i = -\frac{\eta}{8} \left[ \partial_i \left\| g^{(\mathcal{B})} - g^{(\mathcal{C})} \right\|^2 + 4 \left( g_j^{(\mathcal{C})} h_{ij}^{(\mathcal{B})} - g_j^{(\mathcal{B})} h_{ij}^{(\mathcal{C})} \right) \right]. \quad (12)$$

Averaging over different data realizations and substituting in (9), we see that

$$\langle \delta g_i \rangle = -\frac{1}{2} \partial_i \langle \delta \epsilon \rangle. \quad (13)$$

This is our main result. To leading order in  $\eta$ , the difference between SGD and GD can be interpreted in terms of a difference in the gradient vector equal to the derivative of change in the average generalization error. By considering a modified loss  $\tilde{\ell}(\theta, x) \equiv \ell(\theta, x) + \frac{1}{4} \langle \delta \epsilon \rangle$ , up to  $O(\eta^3)$  this means that two full-batch GD updates of size  $N$  each can emulate two non-overlapping half-batch SGD updates of size  $N/2$  each, making the *implicit regularization* of [8] explicit.

## 4 Various Comments

Many generalizations were held out for brevity. For instance, the calculation in §3 is simple to extend to an epoch of any number of updates, and the conclusion remains the same (though the coefficient is different). We can also modify the calculations for second-order methods or to include momentum.

Our results (9) and (13) are model-independent formula that relate performance metrics of interest to simple quantities that can be measured empirically for any particular model and dataset of interest. However, for some models it's possible to evaluate these quantities directly, e.g. for a linear model with quadratic loss and Gaussian input data. In that case, we can show explicitly that in this context SGD's implicit regularization can behave like  $L_2$  regularization of the parameter vector.

Since our results are essentially expressed in terms of a dual expansion in  $\eta$  and  $1/N$  (or  $1/\text{batch-size}$ ), there are many corrections to discuss, such as next to leading order in  $\eta$ , or whether the quantities  $\delta\epsilon$  and  $\delta g_i$  self-average. These calculations are easily performed, but were also held for brevity.

One important correction concerns the fact that the gradient distributions evaluated on the test set vs. the training set will begin to diverge quickly as a model trains and overfits.<sup>4</sup> Considering (9), if these distributions are different, then the change in generalization error would become  $\eta \text{tr} \Sigma_{train}/N + \eta G_{train} \cdot (G_{test} - G_{train})$ , and thus has an additional piece that depends on the difference between test and train mean gradients. Since the second term is still present, this doesn't falsify our result (13), but shows that this divergence of gradients is also an important factor in overfitting.

Finally, we note that these results can lead to improved stochastic optimization algorithms. For instance, to leading order one can go to larger batch sizes (perhaps even full-batch) by modifying the loss to make the SGD regularization explicit. Or one could use the inverse of  $\Sigma_{ij}$  as a preconditioner to slow down learning in directions that the gradient is strongly fluctuating.<sup>5</sup> This latter method appears to work rather well, and we hope to report on it in an extended version of the present work.

<sup>4</sup>This is easy to measure and could serve as a new statistic for overfitting that is worth investigating further.

<sup>5</sup>In general  $\Sigma_{ij}$  is very anisotropic, and it's exceptionally computationally efficient to compute its leading eigenvalues and eigenvectors. Together, these facts make such an algorithm very practical.

## Acknowledgments

We are grateful to Léon Bottou, Boris Hanin, Max Kleiman-Weiner, Yann LeCun, Tengyu Ma, DJ Strouse, and Sho Yaida for comments and discussions. This extended abstract was brought to you by the letter  $\Sigma$  after averaging over many different realizations.

## References

- [1] Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *arXiv preprint arXiv:1606.04838*, 2016.
- [2] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. 2016.
- [3] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- [4] Léon Bottou and Yann Le Cun. On-line learning for very large data sets. *Applied stochastic models in business and industry*, 21(2):137–151, 2005.
- [5] Léon Bottou and Olivier Bousquet. The tradeoffs of large scale learning. In *Advances in neural information processing systems*, pages 161–168, 2008.
- [6] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.
- [7] Samuel L. Smith and Quoc V. Le. A bayesian perspective on generalization and stochastic gradient descent. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=BJij4yg0Z>.
- [8] Behnam Neyshabur, Ryota Tomioka, Ruslan Salakhutdinov, and Nathan Srebro. Geometry of optimization and implicit regularization in deep learning. *arXiv preprint arXiv:1705.03071*, 2017.
- [9] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: Biasing gradient descent into wide valleys. *arXiv preprint arXiv:1611.01838*, 2016.
- [10] 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.
- [11] Wenqing Hu, Chris Junchi Li, Lei Li, and Jian-Guo Liu. On the diffusion approximation of nonconvex stochastic gradient descent. *arXiv preprint arXiv:1705.07562*, 2017.
- [12] Robert Kleinberg, Yuanzhi Li, and Yang Yuan. An alternative view: When does sgd escape local minima? *arXiv preprint arXiv:1802.06175*, 2018.
- [13] Zhanxing Zhu, Jingfeng Wu, Lei Wu, Jinwen Ma, and Bing Yu. The regularization effects of anisotropic noise in stochastic gradient descent. *arXiv preprint arXiv:1803.00195*, 2018.
- [14] Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural Computation*, 9(1):1–42, 1997.
- [15] Behnam Neyshabur, Ruslan R Salakhutdinov, and Nati Srebro. Path-sgd: Path-normalized optimization in deep neural networks. In *Advances in Neural Information Processing Systems*, pages 2422–2430, 2015.
- [16] Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. Sharp minima can generalize for deep nets. *arXiv preprint arXiv:1703.04933*, 2017.
- [17] Arvind Neelakantan, Luke Vilnis, Quoc V Le, Ilya Sutskever, Lukasz Kaiser, Karol Kurach, and James Martens. Adding gradient noise improves learning for very deep networks. *arXiv preprint arXiv:1511.06807*, 2015.

- [18] Yao Zhang, Andrew M. Saxe, Madhu S. Advani, and Alpha A. Lee. Energy–entropy competition and the effectiveness of stochastic gradient descent in machine learning. *Molecular Physics*, 116(21-22):3214–3223, 2018. doi: 10.1080/00268976.2018.1483535.
- [19] Stephan Mandt, Matthew D Hoffman, and David M Blei. Stochastic gradient descent as approximate bayesian inference. *The Journal of Machine Learning Research*, 18(1):4873–4907, 2017.
- [20] Pratik Chaudhari and Stefano Soatto. Stochastic gradient descent performs variational inference, converges to limit cycles for deep networks. *arXiv preprint arXiv:1710.11029*, 2017.
- [21] Andrew K Lampinen and Surya Ganguli. An analytic theory of generalization dynamics and transfer learning in deep linear networks. *arXiv preprint arXiv:1809.10374*, 2018.
- [22] Sho Yaida. Fluctuation-dissipation relations for stochastic gradient descent. *arXiv preprint arXiv:1810.00004*, 2018.
- [23] Dong Yin, Ashwin Pananjady, Max Lam, Dimitris Papailiopoulos, Kannan Ramchandran, and Peter Bartlett. Gradient diversity: a key ingredient for scalable distributed learning. *arXiv preprint arXiv:1706.05699*, 2017.