Energy-entropy competition and the effectiveness of stochastic gradient descent in machine learning

Yao Zhang,[1] Andrew M. Saxe,[2] Madhu S. Advani,[2] and Alpha A. Lee[1, *]

[1] *Cavendish Laboratory, University of Cambridge, Cambridge CB3 0HE, United Kingdom*
[2] *Center for Brain Science, Harvard University, Cambridge, MA 02138, United States of America*

Finding parameters that minimise a loss function is at the core of many machine learning methods. The Stochastic Gradient Descent algorithm is widely used and delivers state of the art results for many problems. Nonetheless, Stochastic Gradient Descent typically cannot find the global minimum, thus its empirical effectiveness is hitherto mysterious. We derive a correspondence between parameter inference and free energy minimisation in statistical physics. The degree of undersampling plays the role of temperature. Analogous to the energy-entropy competition in statistical physics, wide but shallow minima can be optimal if the system is undersampled, as is typical in many applications. Moreover, we show that the stochasticity in the algorithm has a non-trivial correlation structure which systematically biases it towards wide minima. We illustrate our argument with two prototypical models: image classification using deep learning, and a linear neural network where we can analytically reveal the relationship between entropy and out-of-sample error.

INTRODUCTION

Methods in machine learning typically involves finding model parameters that minimize a measure of disagreement between model predictions and training data, known as the loss function. Suppose we performed P measurements of independent variables $\mathbf{x}_i \in \mathbb{R}^{N_i}$ and dependent variable y_i. Given a parametric model $f(\mathbf{x}_i, \boldsymbol{\theta})$ where $\boldsymbol{\theta} \in \mathbb{R}^N$ are the parameters, we can define a loss function $L(\{\mathbf{x}_i\}_{i=1}^P, \boldsymbol{\theta}) = R(\boldsymbol{\theta}) + \frac{1}{P}\sum_i l_i(y_i, f(\mathbf{x}_i, \boldsymbol{\theta}))$, where R is a regularising function. The Stochastic Gradient Descent (SGD) algorithm minimises the loss function by going down the direction of the steepest gradient estimated using random subsets of the data [1–3]:

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta_t \nabla_{\boldsymbol{\theta}} \left(\frac{1}{|B_t|} \sum_{i \in B_t} l_i(\mathbf{x}_i, \boldsymbol{\theta}) + R(\boldsymbol{\theta}) \right), \quad (1)$$

where we have abbreviated our notation for the loss function by hiding the dependent variables y_i, and defined η_t as the learning rate, and B_t as a randomly selected batch of training data. Models routinely used in machine learning can contain $N = O(10^7)$ parameters and are non-convex, thus finding the global minimum is computationally intractable. Nonetheless, it is empirically known that SGD yields state of the art results in machine learning, in the sense that the parameters SGD finds can give accurate predictions outside the training set (so-called generalisation/testing accuracy) [4–8]. It is highly implausible that SGD finds the global minimum, and it is mysterious why minima that SGD converges to tend to be generalisable.

There is a growing literature that sheds light on this mystery by focusing on the topology of the loss function. For Gaussian random functions, critical points with a loss function much larger than the global minimum are exponentially likely to be saddle points in the high dimensional limit [9, 10]. Therefore the value of the loss function at most local minima is similar to the global minimum, and some argue that this is why most local minima are "good enough" [11, 12]. In fact, it can be analytically shown that for some simple network architectures [13–15], there are no local minima and all critical points are either saddle points or the global minimum.

However, this theoretical picture, grounded on the topology of the loss function, has several shortcomings. First, the ultimate metric in machine learning is the generalisation accuracy rather than minimising the loss function, which only captures how well the model fits the training data. With vastly more parameters than samples, it is not difficult to completely fit the training data [16]. Therefore, theoretical results about the value of the loss function at local minima give limited insights about the effectiveness of SGD, which is its ability to find generalisable parameters. Second, it is known empirically that deeper minima have lower generalisation accuracy than shallower minima [12]. In fact, stopping SGD before complete convergence ("early stopping") can yield parameters with higher generalisation accuracy compared to a converged SGD iteration [17, 18]. Those empirical observations show that the local minimum that SGD finds is not only "good enough", but also better than deeper minima or even the global minimum that SGD misses.

Rather than focusing on the topology of the loss function, another strand of the literature focuses on the width of the minima. Pioneering works show that flatter minima lead to lower generalisation error [19–21]. The intuition is that a sharp minimum correspond to a more complex, likely overfitted, model because the parameters need to be accurately specified to define the minimum. Recent empirical studies have focused on the role of stochasticity on SGD: decreasing the batch size causes SGD to converge to wider minima and lead to better generalisation accuracy [22, 23], and in fact just adding Gaussian noise on top of SGD improves generalisation accuracy [24]. An explicit term can also be added to bias

SGD in favour of wider minima [25, 26]. However, definitions of minima width in the literature are phenomenological and not reparametrization-invariant, leading to the critique that wide minima could be made arbitrarily sharp by simply rescaling the model parameters [27].

In this paper, we will first analytically characterise the stochasticity in SGD. Our analysis will show how batch size affects stochasticity, and show that the noise has a non-trivial correlation structure which systematically biases SGD towards wide minima. We will then use a Bayesian approach to derive a correspondence between parameter inference and free energy minimisation in statistical physics with the degree of undersampling playing the role of temperature. Analogous to the energy-entropy competition in statistical physics, wide but shallow minima can be optimal if the system is undersampled. We will then apply our theoretical insight to study two prototypical non-convex models in machine learning: a numerical experiment of image classification using a deep neural network, and a linear neural network where we can derive analytical insights that explicitly relates out-of-sample error and entropy.

STOCHASTIC GRADIENT DESCENT AND ENTROPY-ENERGY COMPETITION

To facilitate an analytic treatment of SGD, we can take the continuum limit and write Equation (19) as a Langevin equation:

$$\frac{d\boldsymbol{\theta}}{dt} = -\nabla_{\boldsymbol{\theta}} L + \boldsymbol{\eta}(t), \tag{2}$$

where

$$\boldsymbol{\eta}(t) = \nabla_{\boldsymbol{\theta}} \left[\frac{1}{N} \sum_i l_i(\mathbf{x}_i, \boldsymbol{\theta}) - \frac{1}{|B_t|} \sum_{i \in B_t} l_i(\mathbf{x}_i, \boldsymbol{\theta}) \right] \tag{3}$$

is the stochastic noise. It is evident that $\langle \boldsymbol{\eta}(t) \rangle = \mathbf{0}$, where the average is performed over batches of data. However, unlike thermal noise, the SGD noise has a non-trivial correlation structure. Extending the results of [28], for a constant batch size $|B_t| = b$ and $P \to \infty$, after some cumbersome algebra yields:

$$\langle \eta_\mu(t) \eta_\nu(t) \rangle = \frac{1}{b} \left(1 - \frac{b}{P} \right) \left[\left\langle \frac{\partial l}{\partial \theta_\mu} \frac{\partial l}{\partial \theta_\nu} \right\rangle - \left\langle \frac{\partial l}{\partial \theta_\mu} \right\rangle \left\langle \frac{\partial l}{\partial \theta_\nu} \right\rangle \right], \tag{4}$$

where $\langle f \rangle = N^{-1} \sum_i f_i$ denotes the average over samples. Two insights can be gained from Equation (4). First, the variance of the noise increases as $O(1/b)$. Therefore, decreasing the batch size increases the stochasticity, explaining the numerical observation of [22]. Second, the noise is highly anisotropic and biases SGD towards wide minima. To fix ideas, we consider the noise near a minima, where $\langle \partial l / \partial \theta_\mu \rangle \approx 0$ for all μ. In the setting where the loss is the negative log probability of the

parameters and data, the first term of Equation (4) is simply the Fisher information, which equals the Hessian $H_{\mu\nu} = \langle \partial^2 l / (\partial \theta_\mu \partial \theta_\nu) \rangle$. Suppose we reparametrise the system such that the Hessian is diagonal. Equation (4) shows that "stiff" directions [29], directions in the Hessian with larger curvature thus larger eigenvalue, is forced with a higher noise, whereas sloppier direction are forced with a lower noise by SGD. This is the opposite of many optimisation methods such as the Newton method where smaller steps are taken along stiffer directions. The net consequence of taking larger steps across stiff direction is that the algorithm can easily escape narrow basin and be trapped in wide basins.

Why might wide basins be generalisable? We consider a model regression problem, where

$$y_i = f(\mathbf{x}_i, \boldsymbol{\theta}^*) + \epsilon_i, \tag{5}$$

In general this ϵ noise can be thought of as arising from degrees of freedom which our model $f(\mathbf{x}, \boldsymbol{\theta})$ does not capture because it does not have sufficient degrees of freedom to represent the true "teacher". In other words, the noise reflects the approximation error of our model. $\boldsymbol{\theta}^*$ thus represent the best possible (unknown) parameters for the model to match the true teacher function. For simplicity and generality, we now approximate the distribution of the approximation error or noise to be Gaussian: $\epsilon_i \sim \mathcal{N}(0, \sigma_i)$, which is a reasonable assumption if it is composed by summing many random, independent components. We also assume the data points are independently sampled, so that the likelihood is:

$$p(\{\mathbf{x}_i\}_{i=1}^P | \boldsymbol{\theta}) \propto \exp \left(- \sum_{i=1}^P \frac{(f(\mathbf{x}_i, \boldsymbol{\theta}) - y_i)^2}{2\sigma_i^2} \right). \tag{6}$$

Applying a Gaussian prior $p(\boldsymbol{\theta}) = \exp \left(-\lambda ||\boldsymbol{\theta}||_2^2 / 2 \right)$ over parameters, the posterior is

$$p(\boldsymbol{\theta} | \{\mathbf{x}_i\}_{i=1}^P) = \frac{1}{Z} \exp \left(- \sum_{i=1}^P \frac{(f(\mathbf{x}_i, \boldsymbol{\theta}) - y_i)^2}{2\sigma_i^2} - \frac{\lambda ||\boldsymbol{\theta}||_2^2}{2} \right). \tag{7}$$

where Z is a normalising constant. Statistical physics insights can be uncovered by first writing Equation (7) in terms of *intensive* quantities, i.e. quantitates that do not scale with the size of the data set P or the number of parameters N. To this end, we define the mean error, which can be interpreted as an "internal energy"

$$u(\boldsymbol{\theta}) = \frac{1}{P} \sum_{i=1}^P \frac{(f(\mathbf{x}_i, \boldsymbol{\theta}) - y_i)^2}{2\sigma_i^2}, \tag{8}$$

and

$$h(\boldsymbol{\theta}) = \frac{\lambda}{N} \frac{||\boldsymbol{\theta}||_2^2}{2}. \tag{9}$$

Given a new data point $\hat{\mathbf{x}}$, the goal is to predict the output $\hat{y}$. The expected $\hat{y}$ is given by:

$$\langle \hat{y} \rangle = \int p(\boldsymbol{\theta}|\{\mathbf{x}_i\}_{i=1}^P) f(\hat{\mathbf{x}}, \boldsymbol{\theta}) \, d\boldsymbol{\theta}. \qquad (10)$$

Machine learning problems where SGD have been successfully applied can be characterised by the asymptotic limit where number of data points $P \to \infty$, but the number of parameters is still vast and the data-to-parameter ratio $P/N = \alpha = O(1)$; this is often referred to as the high dimensional limit [30]. For example, in image recognition challenges, the number of samples $P = O(10^7)$ but the number of parameters in a deep neural network $N = O(10^8)$ [7]. In this limit, it is justified to perform the Laplace approximation,

$$\langle \hat{y} \rangle \approx \frac{1}{Z} \sum_q \frac{\exp\left[-P(u(\boldsymbol{\theta}_q) - \frac{1}{\alpha}h(\boldsymbol{\theta}_q))\right]}{\sqrt{\det H(\boldsymbol{\theta}_q)}} f(\hat{\mathbf{x}}, \boldsymbol{\theta}_q). \quad (11)$$

where $\boldsymbol{\theta}_q$ are the local minima of the loss function,

$$\nabla_{\boldsymbol{\theta}} L(\{\mathbf{x}_i\}_{i=1}^P, \boldsymbol{\theta}_q) = \nabla_{\boldsymbol{\theta}}\left(u(\boldsymbol{\theta}_q) - \frac{1}{\alpha}h(\boldsymbol{\theta}_q)\right) = \mathbf{0}, \quad (12)$$

and $H(\boldsymbol{\theta}_q)$ is the Hessian matrix. The prefactor with the determinant of the Hessian can be written as

$$\frac{1}{\sqrt{\det H(\boldsymbol{\theta}_q)}} = \exp\left(-\frac{1}{2}\sum_{i=1}^N \log \lambda_i(\boldsymbol{\theta}_q)\right), \qquad (13)$$

where λ_i are eigenvalues of the Hessian matrix. If the log eigenvalues are all of the same order of magnitude (note that the logarithm compresses scales), the sum is $O(N)$ and therefore the corresponding intensive quantity is

$$s(\boldsymbol{\theta}_q) = -\frac{1}{2N}\sum_{i=1}^N \log \lambda_i(\boldsymbol{\theta}_q). \qquad (14)$$

The function $s(\boldsymbol{\theta}_q)$ can be interpreted as the entropy associated with that minimum. Geometrically, s is related to the "width" of a minimum — the wider a minimum is, the larger is the entropy function s at that minimum. We note that we have assumed the eigenvalues to be nonzero — this assumption is not necessarily true in deep neural networks that have more parameters than training data [31], and we stress that a full analysis of basin volume [32] is required to characterise the entropy of flat minima. Moreover, we have made a further assumption that the numerical minimisation algorithms can actually find minima – minimisation algorithms are often stuck in saddle points [11]. Our numerical experiments below suggests that for singular Hessian or saddle points, our argument still holds qualitatively if we truncate the sum in Equation (14) to include only large and positive eigenvalues.

All in all, $\langle \hat{y} \rangle$ is approximately

$$\langle \hat{y} \rangle \approx \frac{1}{Z} \sum_q e^{-PF(\boldsymbol{\theta}_q)} f(\hat{\mathbf{x}}, \boldsymbol{\theta}_q), \qquad (15)$$

where

$$F(\boldsymbol{\theta}_q) = u(\boldsymbol{\theta}_q) + \frac{1}{\alpha}(h(\boldsymbol{\theta}_q) - s(\boldsymbol{\theta}_q)) \qquad (16)$$

is the effective free energy. When $P \to \infty$ and α is fixed, the minimum that dominates is *not* the global minimum in $u(\boldsymbol{\theta})$, but rather the minimum with the lowest effective free energy. The free energy, Equation (16), has a transparent physical interpretation. The degree of undersampling $\frac{1}{\alpha}$ is analogous to the temperature, and the mean error and width of the minimum is analogous to the internal energy and entropy. The prior on the parameters, in the case of a Gaussian distribution, is analogous to an entropic spring that pulls the parameters towards the origin. Crucially, the definition of the entropy, which is derived from an asymptotic analysis of the posterior distribution, is based on the determinant of the Hessian thus is invariant to reparameterization. This overcomes one critique in the literature on the correlation between basin width and model performance [27].

We are now in a position to qualitatively understand why the Stochastic Gradient Descent algorithm works – in the high dimensional limit, wide minima, in the sense of large s, are preferable over narrow ones even at the expense of higher error in-sample error u. The extent to which wide minima are preferable is dependent on the degree of undersampling and approximation noise. Wide local minima are of course easy to find via a gradient descent algorithm, and the noise in SGD specifically biases the search toward wider minima, thus partially explaining the unreasonable effectiveness of SGD.

NUMERICAL EXPERIMENTS

Two testable predictions that emerge from our theory are: (1) The anisotropic noise in stochastic gradient descent forces it to converge to wider minima, and (2) subject to constant training error ("energy"), wider minima are more likely to have a low out-of-sample error compared to narrower minima. We will use a prototypical problem in machine learning, the CIFAR-10 challenge [33], to test the extent to which those predictions are borne out in realistic datasets.

The CIFAR-10 dataset consists of 60000 32×32 colour images of 10 different classes of objects (e.g. airplane, automobile, bird etc.), with 6000 images per class. The original machine learning problem is to parameterise a model that can classify an unseen image into one of those 10 classes. For numerical simplicity, we consider a stripped-down version of this problem where we only

consider the binary classification problem of determining whether an image is an airplane or automobile. Moreover, we down sample the images to 10×10. Those simplifications are needed because numerically computing the Hessian for industrial scale models is numerically challenging.

We use a fully-connected neural network with three hidden layers and 10 units per layer. In mathematical terms, we flatten the image into a vector $\mathbf{x} \in \mathbb{R}^{100}$, and pose the model

$$y_i = \sigma_1(\mathbf{W}_5\sigma(\mathbf{W}_4\sigma(\mathbf{W}_3\sigma(\mathbf{W}_2\sigma_1(\mathbf{W}_1\mathbf{x}_i))))) \qquad (17)$$

where $\mathbf{W}_1 \in \mathbb{R}^{10\times100}$, $\mathbf{W}_2 \in \mathbb{R}^{10\times10}$, $\mathbf{W}_3 \in \mathbb{R}^{10\times10}$, $\mathbf{W}_4 \in \mathbb{R}^{10\times10}$ and $\mathbf{W}_5 \in \mathbb{R}^{10\times2}$ are model parameters that we need to infer from the data, $\sigma(x) = \max(0, x)$ is known as the `relu` function in the literature and $\sigma_1(x) = (1+e^{-x})^{-1}$ is the sigmoid function . y_i can be interpreted as a predicted probability – the model is certain that the i^{th} image is an airplane/automobile if $y_i = 0/1$. The appropriate loss function for a classification problem is the so-call cross-entropy function

$$l_i = -t_i \log y_i - (1 - t_i) \log(1 - y_i) \qquad (18)$$

where t_i is the true label (0 for airplane, 1 for automobile) and y_i is the predicted probability.

We have 12000 images of airplanes and automobiles in the dataset and 1320 parameters in our model. We will consider the case where we randomly choose 500 images as the training set. We restrict the number of images in the training data so that we can perfectly fit the model to training data (i.e. the "energy" $u = 0$). The model suffers from overfitting, but the upshot is that we can focus solely on the entropy as the energy of every minimum is zero.

To verify whether the anisotropic noise in stochastic gradient descent forces the algorithm to converge to higher entropy minima, we compare the entropy of the critical points found via SGD with an alternative model of Langevin dynamics where the noise has the same magnitude as SGD but we remove the correlation structure. Concretely, we consider the Langevin dynamics

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta G(\theta) + \boldsymbol{\zeta}(t) \qquad (19)$$

where $G(\theta) = \nabla_{\boldsymbol{\theta}} (L + R(\boldsymbol{\theta}))$ and the uncorrelated noise $\boldsymbol{\zeta}(t)$ is drawn from

$$\boldsymbol{\zeta}(t) \sim \mathcal{N}(0, \mathrm{diag}(\sigma_1(t), \sigma_2(t) \cdots, \sigma_p(t))), \qquad (20)$$

with

$$\sigma_i(t)^2 = \frac{1}{10}\sum_{m=1}^{10}(G_m^{(i)}(\boldsymbol{\theta}) - \bar{G}^{(i)}(\boldsymbol{\theta}))^2, i = 1, ..., p \qquad (21)$$

and we estimate σ_i^2, the variance in the gradient in i^{th} variable, by randomly splitting the sample into 10 minibatches of 50 samples each, estimate the minibatch gradient, and compute the variance of the gradient in the i^{th} variable across the 10 minibatches.
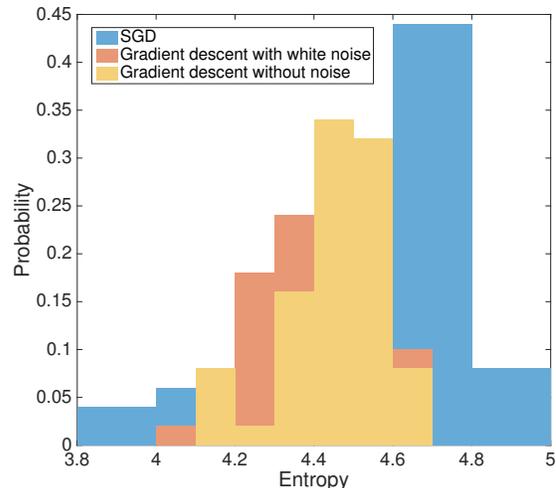
FIG. 1. SGD finds larger entropy solutions compared to Langevin dynamics or gradient descent. The plot shows a histogram of the entropy of the solutions found by the different optimization algorithms for a three-layer neural network trained on the binary classification problem of airplane and automobile images from CIFAR-10 datasets. The histogram is generated by 50 independent runs starting with random initial parameters.

To converge to a critical point, we run SGD with a minibatch size of 50 for 8000 epochs (full iteration through the dataset), followed by steepest descent minimisation until the gradient at each direction is less than 3×10^{-5}. Similarly, for the Langevin dynamics experiments, we run the modified Langevin dynamics for 80000 iterations, followed by steepest descent minimisation until the gradient at each direction is less than 3×10^{-5}. For the full gradient descent without noise experiments, we run the full gradient descent for 80000 iterations and collect the data when the gradient at each direction is less than 3×10^{-5}. We start with a different initial condition for each run, with parameters drawn from a Gaussian distribution $\mathcal{N}(0, 0.1)$. The learning rate η for both SGD and our modified Langevin dynamics is 0.1. In both cases, a small regularisation of $\lambda = 10^{-7}$ is used.

The Hessian matrix contains a small number of small negative eigenvalues (approximately 10%), as well as positive eigenvalues close to zero. It is numerically challenging to differentiate whether those small negative or small positive eigenvalues are due to numerical noise masking a true minimum, a saddle point, or a degenerate minimum where the zero eigenvalues are masked by numerical noise. To make further progress, we use a phenomenological "cleaning" procedure where eigenvalues less than the magnitude of the most negative eigenvalue is ignored. This cleaning procedure is motivated by the intuition that the most negative eigenvalue gives an indication of the "scale" of numerical error.
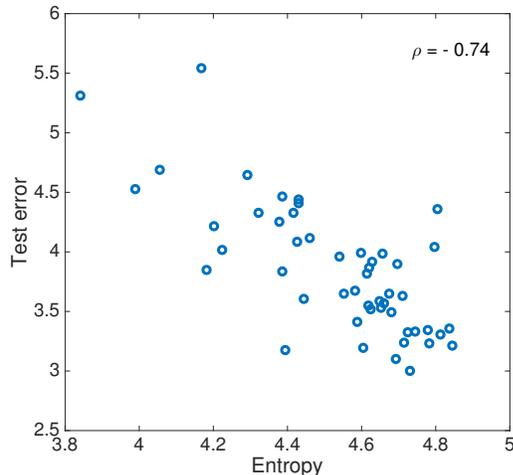
FIG. 2. Solutions with higher entropy are have lower out-of-sample error. We consider a a three-layer neural network trained and binary classification problem of airplane and automobile images discussed in the main text. The plot is generated by performing 50 independent runs of SGD starting from random initial weights. Note that the neural network is overparameterised, thus the in-sample error of the model is zero but the different solutions differs by their out-of-sample error.

Figure 1 shows that the mean entropy of minima found via SGD is larger than the mean entropy found using Langevin dynamics with the same level of noise, confirming our hypothesis that the anisotropy of the noise biases SGD towards higher entropy minima.

We next consider the correlation between entropy and the testing error, the error when the model is applied to images outside the training set. Using the same procedure discussed above, we locate critical points by running SGD for 8000 epochs followed by gradient descent minimisation. We compute the out-of-sample error by applying the model to 2000 images that are not in the training set. Figure 2 shows a discernible negative correlation between entropy and testing error. In other words, the larger the entropy is, the more likely the model can perform well outside the training set. Note that the error of the model in the test set is zero, therefore the only difference between the different solutions is entropy.

ENTROPY IN DEEP LINEAR NETWORKS

As another view onto the relationship between entropy and generalization performance, we examine the behavior of deep linear neural networks. Deep linear networks provide a simple model class that nevertheless retains important features of the learning problem faced in nonlinear deep networks. In particular, the optimization problem is nonconvex [34], and solution dynamics can be highly
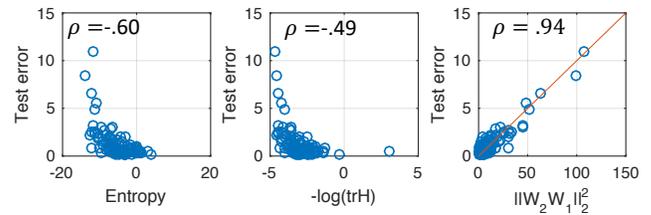
FIG. 3. Results of SGD in underdetermined deep linear networks, out-of-sample error correlates with entropy (left), the trace of the Hessian (middle) and the total weight norm (right). Correlation coefficients are given in each panel. We also demonstrate in the third panel the theoretical prediction (34) for the average scaling of the generalization error with weight norm fits the data well. Other parameters: $N_i = 10, N_h = 7, P = 5, \sigma_e = 10^{-4}$.

nonlinear [34]. We consider a three-layer deep linear network which computes the output

$$\hat{y}_i = \mathbf{W}_2\mathbf{W}_2\mathbf{x}_i \qquad (22)$$

where $\mathbf{W}_1 \in \mathbb{R}^{N_h \times N_i}$ and $\mathbf{W}_2 \in \mathbb{R}^{1 \times N_h}$. We train these networks in a student-teacher setting (see e.g. [35, 36]), where a teacher labels random input examples $\mathbf{x} \sim \mathcal{N}(0, \frac{1}{N_i}\mathbf{I})$ as $y = \bar{\mathbf{W}}\mathbf{x} + \epsilon$. The teacher's parameters $\bar{\mathbf{W}}$ are drawn independently from a unit variance Gaussian distribution, and the noise variance (representing approximation error) is $\langle \epsilon^2 \rangle = \sigma_e^2$. In this way a dataset of P examples can be drawn and collected into an input matrix $\mathbf{X} \in \mathbb{R}^{N_i \times P}$ and target vector $\mathbf{y} \in \mathbb{R}^{1 \times P}$. We assume a quadratic loss function so that the network aims to minimize $L = \sum_i (\hat{y}_i - y_i)^2$, and in our simulations, update weights with stochastic gradient descent.

The generalization behavior of such networks is dependent on the amount of training data they receive [36]. In the overdetermined case where data is plentiful ($P \geq N_i$), minimizers of the training error must compute identical input-output functions and attain identical generalization performance. However in the underdetermined case where $P < N_i$, minimizers of the training error all attain zero error on training samples, but need not compute identical input-output functions and therefore can generalize differently to new test samples. Intuitively, this is because the behavior of the input-output map is unconstrained in directions which contain no training data. We thus tested whether the correlation between entropy and generalization performance holds in the underdetermined regime. Fig. 3 confirms this correlation for deep linear networks. Here as before, eigenvalues of the Hessian that are exactly zero have been discarded. We note that for this linear case, the number of nonzero eigenvalues is simply $\min(P, N_i)$ and hence by sorting the eigenvalues and taking the top $\min(P, N_i)$, issues of numerical precision are avoided.

In the overdetermined setting, there is a manifold of global minima where the product of the weights is equal

to the least squares solution, $\mathbf{W}_2\mathbf{W}_1 = \mathbf{\Sigma}^{yx}(\mathbf{\Sigma}^x)^{-1}$ where $\mathbf{\Sigma}^{yx} = \mathbf{y}\mathbf{X}^{\mathbf{T}}$ and $\mathbf{\Sigma}^x = \mathbf{X}\mathbf{X}^T$ are the input-output covariance and input covariance respectively. All networks that attain this minimum must have the same input-output function and generalize identically. Even in this simple setting, however, the Hessian at every minimum has zero eigenvalues corresponding to directions in parameter space which point along the manifold of minima. In particular, a network with N_h hidden units will have $N_h(N_i + 1) - N_i$ zero eigenvalues.

As the simplest example, a deep linear chain with scalar weights w_2 and w_1 (i.e. $N_i = N_h = 1$) has Hessian

$$H = \begin{bmatrix} w_2^2 \sigma^x & 2w_2 w_1 \sigma^x - \sigma^{yx} \\ 2w_2 w_1 \sigma^x - \sigma^{yx} & w_1^2 \sigma^x \end{bmatrix}, \quad (23)$$

yielding eigenvalues $\lambda_1 = w_2^2 \sigma^x + w_1^2 \sigma^x$ and $\lambda_2 = 0$ on the manifold of global minima. Hence the spectral norm of the Hessian coincides with the trace of the Hessian and is simply λ_1. Further, our entropy measure corresponds to taking the negative log of λ_1. The manifold of global minima is defined by the hyperbola $w_1 w_2 = \sigma^{yx}/\sigma^x$. Substituting this into λ_1, we have $\lambda_1 = (\sigma^{yx})^2/(\sigma^x w_1^2) + w_1^2 \sigma^x$ on the solution manifold, which retains a dependence on w_1. In particular, taking w_1 to infinity or zero will cause the entropy to go to infinity; but all points on the solution manifold implement the same function and generalize identically. Thus a question remains: why do we observe a strong correlation between entropy and generalization error in typical training settings in spite of the parameter dependence of these Hessian-based metrics?

Some insight can be gained by examining the trace of the Hessian, which for these deep linear networks we find to be

$$\mathrm{tr}H = \|\mathbf{W}_2\|_2^2 \|\mathbf{X}\|_F^2 + \|\mathbf{W}_1\mathbf{X}\|_F^2, \quad (24)$$

where F denotes a Frobenius norm. As we show in Figure 5, the negative log of the trace of the Hessian closely tracks the entropy, and correlates about equally well with the generalization performance (Figure 3 middle), so understanding the behavior of the trace can yield insight into the entropy measure as well. (We note, for instance, that in the linear chain considered above the entropy is exactly equal to $-\log(\mathrm{tr}H)$). With this in mind, the reason for the correlation becomes evident: standard training procedures initialize the norm of each layer's weights to be approximately equal, and when so initialized, this persists through training. That is, under batch gradient descent learning (where every sample is used in each update $|B| = N$), the difference $\|\mathbf{W}_2\|_2^2 - \|\mathbf{W}_1\|_F^2$ is an invariant of the dynamics.

We can prove this fact by deriving the dynamics for this setting in the continuous limit by computing the derivatives of the loss function with respect to the
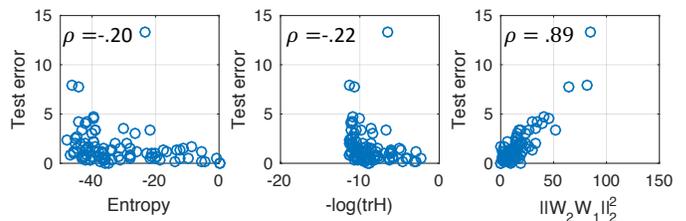
FIG. 4. Results of SGD in networks with extremely imbalanced initial weights ($\|\mathbf{W}_2\|_2^2 \geq \|\mathbf{W}_1\|_F^2$ or vice versa), the correlation between out-of-sample error and entropy (left) or the trace of the Hessian (center) decreases, but the correlation with total weight norm remains unchanged (c.f. Fig. 3). Weights were initialized with random asymmetries between 1 and 100. Other parameters as in Fig. 3.

weights. This yields the coupled differential equations:

$$\tau \frac{d}{dt}\mathbf{W}_1 = \mathbf{W}_2^T \left(\mathbf{\Sigma}^{yx} - \mathbf{W}_2\mathbf{W}_1\mathbf{\Sigma}^{xx} \right), \quad (25)$$

$$\tau \frac{d}{dt}\mathbf{W}_2 = \left(\mathbf{\Sigma}^{yx} - \mathbf{W}_2\mathbf{W}_1\mathbf{\Sigma}^{xx} \right)\mathbf{W}_1^T. \quad (26)$$

We can write our invariant as:

$$\frac{d}{dt}\left(\|\mathbf{W}_2\|_2^2 - \|\mathbf{W}_1\|_F^2 \right) = \frac{d}{dt}Tr\left(\mathbf{W}_2^T\mathbf{W}_2\right) - \frac{d}{dt}Tr\left(\mathbf{W}_1\mathbf{W}_1^T\right). \quad (27)$$

Applying the derivative yields

$$Tr\left(\mathbf{W}_2^T\dot{\mathbf{W}}_2 - \dot{\mathbf{W}}_1\mathbf{W}_1^T \right) + Tr\left(\dot{\mathbf{W}}_2^T\mathbf{W}_2 - \mathbf{W}_1\dot{\mathbf{W}}_1^T \right) = 0, \quad (28)$$

where the final equality holds by substituting (25) and (26) into the LHS of the expression above.

Batch gradient descent is equivalent to SGD for small learning rates and so this invariance can be expected to approximately hold in that setting. It follows that a balanced initial condition implicitly restricts the ultimate solution to be balanced as well, resolving the dependence of the trace and entropy metrics on the asymmetry in weight norms. Balanced solutions can still vary widely in norm, and from Equation (24), solutions with lower norm will have lower trace and entropy. For the linear networks considered here, the optimal out-of-sample error is attained by the minimum norm solution. Hence for deep linear networks, the correlation between entropy and generalization performance is stronger for balanced initial conditions (which are standard in practice). To demonstrate this, in Figure 4 we modify our initialization procedure to introduce asymmetry, yielding a reduced correlation.

The exact generalization error is given by $e_g = \frac{1}{N_i}\left\|\bar{\mathbf{W}} - \mathbf{W}_2\mathbf{W}_1\right\|_2^2 + \sigma_e^2$. Now consider splitting these matrices into the component in the subspace of the training data, and the component perpendicular to the training data. We thus have

$$e_g = \frac{1}{N_i}\left\| \bar{\mathbf{W}}^{\|} + \bar{\mathbf{W}}^{\perp} - \mathbf{W}_2\mathbf{W}_1^{\|} - \mathbf{W}_2\mathbf{W}_1^{\perp} \right\|_2^2 + \sigma_e^2. \quad (29)$$
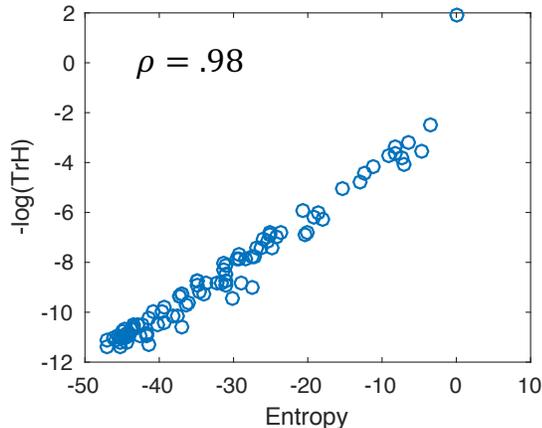
FIG. 5. The negative log of the trace of the Hessian correlates strongly with entropy.

If training is successful, then $\bar{\mathbf{W}}^{\parallel} - \mathbf{W}_2\mathbf{W}_1^{\parallel} \approx 0$. Therefore

$$e_g \approx \frac{1}{N_i} \left\| \bar{\mathbf{W}}^{\perp} - \mathbf{W}_2\mathbf{W}_1^{\perp} \right\|_2^2 + \sigma_e^2. \qquad (30)$$

Given P examples, and N_i input units, the dimension of the null space of the data covariance will be $N_i - P$. Also, there will be no correlation between $\bar{\mathbf{W}}^{\perp}$ and $\mathbf{W}_2\mathbf{W}_1^{\perp}$ because the student weights are initialized randomly and are exposed to no information about the true weights in the direction orthogonal to the data. It follows that:

$$\langle e_g \rangle = \frac{1}{N_i} \left\langle \|\bar{\mathbf{W}}^{\perp}\|_2^2 \right\rangle + \frac{1}{N_i} \left\langle \|\mathbf{W}_2\mathbf{W}_1^{\perp}\|_2^2 \right\rangle + \sigma_e^2. \qquad (31)$$

To relate the previous expression to the norm of the product of the weights, we use the fact that

$$\|\mathbf{W}_2\mathbf{W}_1\|_2^2 = \|\mathbf{W}_2\mathbf{W}_1^{\parallel} + \mathbf{W}_2\mathbf{W}_1^{\perp}\|_2^2 \approx \|\bar{\mathbf{W}}^{\parallel} + \mathbf{W}_2\mathbf{W}_1^{\perp}\|_2^2. \qquad (32)$$

Rearranging and averaging the expression above yields:

$$\left\langle \|\mathbf{W}_2\mathbf{W}_1^{\perp}\|_2^2 \right\rangle = - \left\langle \|\bar{\mathbf{W}}^{\parallel}\|_2^2 \right\rangle + \left\langle \|\mathbf{W}_2\mathbf{W}_1\|_2^2 \right\rangle. \qquad (33)$$

Substituting this into our expression for generalization error yields:

$$\langle e_g \rangle = \frac{1}{N_i} \left(\left\langle \|\bar{\mathbf{W}}^{\perp}\|_2^2 \right\rangle - \left\langle \|\bar{\mathbf{W}}^{\parallel}\|_2^2 \right\rangle \right) + \frac{1}{N_i} \left\langle \|\mathbf{W}_2\mathbf{W}_1\|_2^2 \right\rangle + \sigma_e^2. \qquad (34)$$

Thus, we have an expression for how the generalization performance of this student network is linearly related to the norm of its weights. The fit of this prediction may be seen in simulations in Figure 3.

In the under-determined setting of a deep linear network, when the Hessian has many zero directions, we have shown that high entropy solutions are correlated with lower test error. In this simple setting, the reason for this connection is that the higher entropy solutions tend to have lower weight norms, and the weight norm is strongly correlated with the test error. In more complex non-linear networks, a similar mechanism may play a role, but there may also be additional causes behind why the entropy metric is related to generalization performance.

CONCLUSION

We showed that the SGD algorithm, a workhorse minimisation algorithm in machine learning, preferentially finds minima that are wide because the noise is correlated and anisotropic. By performing a Laplace approximation, we show that the training error and log determinant of the Hessian matrix plays the role of energy and entropy in statistical physics, whilst the degree of undersampling plays the role of temperature. In the undersampled, "high temperature", regime the Bayes optimal parameters are determined by a balance between the training error and the basin width. This provides a physical justification of why stochastic gradient descent can train a model with low out-of-sample error despite almost certainly not finding the global minimum in the loss function. We tested the energy-entropy competition hypothesis with two prototypical machine learning models: a simple image classification problem and a linear neural network. The numerical experiments show that: (1) SGD locates minima with higher entropy compared to gradient descent with white noise (Langevin dynamics) or gradient descent without noise; (2) subject to constant training error, higher entropy solutions enjoy lower testing error. For the linear neural network model, we can analytically explain the relationship between entropy and test error.

However, our analysis thus far focused on an harmonic expansion around minima, leading to a characterisation of entropy based on the log determinant of the Hessian matrix. There are two shortcomings with this characterisation. First, as is shown in the linear case, the Hessian matrix is singular for overparameterised systems, and zero eigenvalues are present also in overparametrised nonlinear neural networks models [31]. Although we heuristically justified our phenomenological clipping of low eigenvalues in the regime of balanced weight initialisation, a more nuanced metric is needed to characterise flat minima with singular Hessian matrices. Second, the harmonic expansion approximation ignores higher order geometric features of the basins of attraction. The molecular physics community has developed a rich framework for describing energy landscapes [37], and techniques from molecular simulations that can compute the basin volume of high dimensional energy landscape will likely yield significantly greater insights into the connection between basin volume and test error [32, 38].

A perhaps more vexing issue is that neural networks seldom converge to a minima – in fact, it can be shown analytically that the number of saddle points proliferates in high dimensional energy landscape functions [11, 12]. Therefore, an open question is to go beyond characterising minima and develop a statistical framework that relate geometric properties of saddle points to the out-of-sample performance of machine learning models.

It is our pleasure to dedicate this work to Daan Frenkel on the occasion of his 70th birthday. Alpha Lee is particularly indebted to Daan for his continuing scientific inspiration and mentionship, from his first interactions with Daan as a teaching assistant in Park City to being colleagues in Cambridge. AAL is supported by the Winton Programme for the Physics of Sustainability, and AMS and MSA are supported by the Swartz Program in Theoretical Neuroscience.

* aal44@cam.ac.uk

[1] H. Robbins and S. Monro, The annals of mathematical statistics , 400 (1951).

[2] L. Bottou, in *Proceedings of COMPSTAT'2010* (Springer, 2010) pp. 177–186.

[3] L. Bottou, F. E. Curtis, and J. Nocedal, arXiv preprint arXiv:1606.04838 (2016).

[4] A. Krizhevsky, I. Sutskever, and G. E. Hinton, in *Advances in neural information processing systems* (2012) pp. 1097–1105.

[5] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, *et al.*, International Journal of Computer Vision **115**, 211 (2015).

[6] Y. LeCun, Y. Bengio, and G. Hinton, Nature **521**, 436 (2015).

[7] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. Alemi, arXiv preprint arXiv:1602.07261 (2016).

[8] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, *et al.*, Nature **529**, 484 (2016).

[9] Y. V. Fyodorov and I. Williams, Journal of Statistical Physics **129**, 1081 (2007).

[10] A. J. Bray and D. S. Dean, Physical review letters **98**, 150201 (2007).

[11] Y. N. Dauphin, R. Pascanu, C. Gulcehre, K. Cho, S. Ganguli, and Y. Bengio, in *Advances in neural information processing systems* (2014) pp. 2933–2941.

[12] A. Choromanska, M. Henaff, M. Mathieu, G. B. Arous, and Y. LeCun, in *AISTATS* (2015).

[13] P. Baldi and K. Hornik, Neural networks **2**, 53 (1989).

[14] K. Kawaguchi, in *Advances in Neural Information Processing Systems* (2016) pp. 586–594.

[15] H. Lu and K. Kawaguchi, arXiv preprint arXiv:1702.08580 (2017).

[16] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, arXiv preprint arXiv:1611.03530 (2016).

[17] L. Prechelt, in *Neural Networks: Tricks of the Trade* (Springer, 2012) pp. 53–67.

[18] D. Duvenaud, D. Maclaurin, and R. Adams, in *Artificial Intelligence and Statistics* (2016) pp. 1070–1077.

[19] G. E. Hinton and D. Van Camp, in *Proceedings of the sixth annual conference on Computational learning theory* (ACM, 1993) pp. 5–13.

[20] S. Hochreiter and J. Schmidhuber, Advances in Neural Information Processing Systems , 529 (1995).

[21] S. Hochreiter and J. Schmidhuber, Neural Computation **9**, 1 (1997).

[22] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang, arXiv preprint arXiv:1609.04836 (2016).

[23] S. Jastrzebski, Z. Kenton, D. Arpit, N. Ballas, A. Fischer, Y. Bengio, and A. Storkey, arXiv preprint arXiv:1711.04623 (2017).

[24] A. Neelakantan, L. Vilnis, Q. V. Le, I. Sutskever, L. Kaiser, K. Kurach, and J. Martens, arXiv preprint arXiv:1511.06807 (2015).

[25] C. Baldassi, C. Borgs, J. T. Chayes, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina, Proceedings of the National Academy of Sciences **113**, E7655 (2016).

[26] P. Chaudhari, A. Choromanska, S. Soatto, and Y. LeCun, arXiv preprint arXiv:1611.01838 (2016).

[27] L. Dinh, R. Pascanu, S. Bengio, and Y. Bengio, arXiv preprint arXiv:1703.04933 (2017).

[28] C. J. Li, L. Li, J. Qian, and J.-G. Liu, arXiv preprint arXiv:1705.07562 (2017).

[29] J. J. Waterfall, F. P. Casey, R. N. Gutenkunst, K. S. Brown, C. R. Myers, P. W. Brouwer, V. Elser, and J. P. Sethna, Physical review letters **97**, 150601 (2006).

[30] M. Advani, S. Lahiri, and S. Ganguli, Journal of Statistical Mechanics: Theory and Experiment **2013**, P03014 (2013).

[31] L. Sagun, U. Evci, V. U. Guney, Y. Dauphin, and L. Bottou, arXiv:1706.04454 (2018).

[32] S. Martiniani, K. J. Schrenk, J. D. Stevenson, D. J. Wales, and D. Frenkel, Physical Review E **94**, 031301 (2016).

[33] A. Krizhevsky and G. Hinton, *Learning multiple layers of features from tiny images*, Tech. Rep. (2009).

[34] A. M. Saxe, J. L. McClelland, and S. Ganguli, arXiv preprint arXiv:1312.6120 (2013).

[35] H. Seung, H. Sompolinsky, and N. Tishby, Physical Review A **45**, 6056 (1992).

[36] M. S. Advani and A. M. Saxe, arXiv preprint arXiv:1710.03667 (2017).

[37] D. Wales, *Energy landscapes: Applications to clusters, biomolecules and glasses* (Cambridge University Press, 2003).

[38] A. J. Ballard, R. Das, S. Martiniani, D. Mehta, L. Sagun, J. D. Stevenson, and D. J. Wales, Physical Chemistry Chemical Physics **19**, 12585 (2017).