# The Local Learning Coefficient: A Singularity-Aware Complexity Measure

**Edmund Lau** University of Melbourne Zach Furman Timaeus George Wang Timaeus

**Daniel Murfet** University of Melbourne

#### Abstract

The Local Learning Coefficient (LLC) is introduced as a novel complexity measure for deep neural networks (DNNs). Recognizing the limitations of traditional complexity measures, the LLC leverages Singular Learning Theory (SLT), which has long highlighted the significance of singularities in the loss landscape geometry. This paper provides an extensive exploration of the LLC's theoretical underpinnings, offering both a clear definition and intuitive insights into its application. Moreover, we propose a new scalable estimator for the LLC, which is then effectively applied across diverse architectures including deep linear networks up to 100M parameters, ResNet image models, and transformer language models. Empirical evidence suggests that the LLC provides valuable insights into how training heuristics might influence the effective complexity of DNNs. Ultimately, the LLC emerges as a crucial tool for reconciling the apparent contradiction between deep learning's complexity and the principle of parsimony.

# **1 INTRODUCTION**

Occam's razor, a foundational principle in scientific inquiry, suggests that the simplest among competing hypotheses should be selected. This principle has emphasized simplicity and parsimony in scientific thinking for centuries. Yet, the advent of deep neural networks Susan Wei Monash University

(DNNs), with their multi-million parameter configurations and complex architectures, poses a stark challenge to this time-honored principle. The question arises: how do we reconcile the effectiveness of these intricate models with the pursuit of simplicity?

It is natural to pose this question in terms of model complexity. Unfortunately, many existing definitions of model complexity are problematic for DNNs. For instance, the parameter count, which in classical statistical learning theory is a commonly-used measure of the amount of information captured in a fitted model, is well-known to be inappropriate in deep learning. This is clear from technical results on generalization (Zhang et al., 2017), pruning (Blalock et al., 2020) and distillation (Hinton et al., 2015): the amount of information captured in a trained network is in some sense decoupled from the number of parameters.

As forcefully argued in (Wei et al., 2022), the gap in our understanding of DNN model complexity is induced by singularities, hence the need for a singularityaware complexity measure. This motivates our appeal to Singular Learning Theory (SLT), which recognizes the necessity for sophisticated tools in studying statistical models exhibiting singularities. Roughly speaking, a model is singular if there are many ways to vary the parameters without changing the function; the more ways to vary without a change, the more singular (and thus more degenerate). DNNs are prime examples of singular models, characterized by complex degeneracies that make them highly singular.

SLT continues a longstanding tradition of using *free* energy as the starting point for measuring model complexity (Bialek et al., 2001), and here we see the implications of singularities. Specifically, the free energy, also known as the negative log marginal likelihood, can be shown to asymptotically diverge with sample size n according to the law  $an + b \log n + o(\log \log n)$ ; the coefficient a of the linear term is the minimal loss

Proceedings of the 28<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2025, Mai Khao, Thailand. PMLR: Volume 258. Copyright 2025 by the author(s).



Figure 1: Impact of SGD learning rate (top), batch size (middle) and momentum (bottom) when training ResNet18 on CIFAR10. We plot the LLC estimate (left), test accuracy (middle) and train loss (right) across training time. As the strength of the implicit regularization increases — through higher learning rate, lower batch size and higher momentum — LLC decreases (the network gets "simpler") and test accuracy increases. Even though most training losses collapse to zero, the LLC can discern the implicit regularization pressure applied by various training heuristics.

achievable and the **coefficient** b of the remaining logarithmic divergence is then taken as the model complexity of the model class. In regular statistical models, the coefficient b is the number of parameters (divided by 2). In singular statistical models, b is not tied to the number of parameters, indicating a different kind of complexity at play.

The LLC arises out of a consideration of the *local* free energy. We explore the mathematical underpinnings of the LLC in Section 3, utilizing intuitive concepts like volume scaling to aid understanding. Our contributions encompass 1) the definition of the new LLC complexity measure, 2) the development of a scalable estimator for the LLC, and 3) empirical validations that underscore the accuracy and practicality of the LLC estimator. In particular, we demonstrate that the estimator is accurate and scalable to modern network size in a setting where theoretical learning coefficients are available for comparison. Furthermore, we show empirically that some common training heuristics effectively control the LLC.

On this last point, we preview some results; Figure 1 displays the LLC estimate, test accuracy, and and training loss over the course of training ResNet18 on CIFAR10. Loosely speaking, lower LLC means a less complex, and thus more degenerate, neural network. In each of the rows, lighter colors represent stronger implicit regularization, e.g., higher learning rate, lower batch size, higher SGD momentum, which corresponds to a preference for lower LLC, i.e., simpler NNs.

# 2 SETUP

Let  $W \subset \mathbb{R}^d$  be a compact space of parameters  $w \in W$ . Consider the model-truth-prior triplet

$$(p(x, y|w), q(x, y), \varphi(w)), \tag{1}$$

where q(x, y) = q(y|x)q(x) is the true data-generating mechanism, p(x, y|w) = p(y|x, w)q(x) is the posited model with parameter w representing the neural network weights, and  $\varphi$  is a prior on w. Suppose we are given a training dataset of n input-output pairs,  $\mathcal{D}_n = \{(x_i, y_i)\}_{i=1}^n$ , drawn i.i.d. from q(x, y).

To these objects, we can associate the sample negative log likelihood function defined as  $L_n(w) = -\frac{1}{n} \sum_{i=1}^{n} \log p(y_i|x_i, w)$ , and its theoretical counterpart defined as  $L(w) = -E_{q(x,y)} \log p(y|x, w)$ . It is appropriate to also call  $L_n$  and L the training and population loss, respectively, since using the negative log likelihood encompasses many classic loss functions used in machine learning and deep learning, such as mean squared error (MSE) and cross-entropy.

The behavior of the training and population losses is highly nontrivial for neural networks. To properly account for model complexity of neural network models, it is critical to engage with the challenges posed by singularities. To appreciate this, we follow (Watanabe, 2009) and make the following distinction between regular and singular statistical models. A statistical model p(x, y|w) is called **regular** if it is 1) identifiable, i.e. the parameter to distribution map  $w \mapsto p(x, y|w)$  is one-to-one, and 2) its Fisher information matrix I(w)is everywhere positive definite. We call a model **singular** if it is not regular.

For an introduction to the implications of singular learning theory for deep learning, we refer the readers to Appendix A and further reading in (Wei et al., 2022). We shall assume throughout the triplet (1) satisfies a few fundamental conditions in SLT (Watanabe, 2009). These conditions are stated and discussed in an accessible manner in A.1.

# 3 THE LOCAL LEARNING COEFFICIENT

In this paper, we introduce the Local Learning Coefficient (LLC), an extension of Watanabe (2009)'s global learning coefficient, referred to as simply learning coefficient there. Below we focus on explaining the LLC through its geometric intuition, specifically as an invariant based on the volume of the loss landscape basin.

For readers interested in the detailed theoretical foundations, we have included comprehensive explanations in the appendices. Appendix A offers a short introduction to the basics of Singular Learning Theory (SLT), and Appendix B sets out formal conditions of the welldefinedness of the LLC. This structure ensures that readers with varying levels of familiarity with SLT can engage with the content at their own pace.

# 3.1 Complexity via volume of low loss parameters

At a local minimum of the population loss landscape there is a natural notion of complexity, given by the number of bits required to specify the minimum to within a tolerance  $\epsilon$ . This idea is well-known in the literature on minimum description length (Grünwald and Roos, 2019) and was used by (Hochreiter and Schmidhuber, 1997) in an attempt to quantify the complexity of a trained neural network. However, a correct treatment has to take into consideration the degeneracy of the geometry of the population loss L.

Consider a local minimum  $w^*$  of population loss L and a sufficiently small neighborhood  $B(w^*)$ , such that for all  $w \in B(w^*)$  we have  $L(w) \ge L(w^*)$ . Given a tolerance  $\epsilon > 0$  we can consider the set of parameters  $B(w^*, \epsilon) = \{w \in B(w^*) \mid L(w) - L(w^*) < \epsilon\}$  whose loss is within the tolerance, the volume of which we define to be

$$V(\epsilon) := \operatorname{Vol}(B(w^*, \epsilon)) = \int_{B(w^*, \epsilon)} dw \,. \tag{2}$$

The minimal number of bits to specify this set within the neighbourhood  $B(w^*)$  is

$$-\log_2(V(\epsilon)/\operatorname{Vol}(B(w^*))). \tag{3}$$

This can be taken as a measure of the complexity of the set of low loss parameters near  $w^*$ . However, as it stands this notion depends on  $\epsilon$  and has no intrinsic meaning. Classically, this is addressed as follows: if a model is regular and thus L(w) is locally quadratic around  $w^*$ , the volume satisfies a law of the form  $V(\epsilon) \approx c\epsilon^{d/2}$ , where c is a constant that depends on the curvature of the basin around the local minimum  $w^*$ , and d is the dimension of W.

This explains why  $\frac{d}{2}$  is a valid measure of complexity in the regular case, albeit one that cannot distinguish  $w^*$  from any other local minima. The curvature c is less significant than the scaling exponent, but *can* distinguish local minima and log c is sometimes used as a complexity measure.

The population loss of a neural network is *not* locally quadratic near its local minima, from which it follows that the functional form of the volume  $V(\epsilon)$  is more complicated than in the regular case. The correct functional form was discovered by Watanabe (2009) and we



Figure 2: The LLC  $\lambda(w^*)$  measures volume scaling around a local minimum of the loss. Top: the area (or "2D-volume") under the loss  $L(w_1, w_2) = w_1^2 + w_2^2$  scales as  $V(\epsilon) \propto \epsilon$  around  $w^* = (0, 0)$ , so  $\lambda(w^*) = 1$ . Bottom: the area under the loss  $L(w_1, w_2) = w_1^2 w_2^4$  scales as  $V(\epsilon) \propto \epsilon^{1/4}$  around  $w^* = (0, 0)$ , so  $\lambda(w^*) = \frac{1}{4}$ . Lower  $\lambda(w^*)$  means more volume: as we approach the minimum ( $\epsilon \to 0$ ), there is much more volume available for the bottom loss function compared to the top for any given small tolerance  $\epsilon$ . Reproduced with permission from (Hoogland and van Wingerden, 2023).

adapt it here to a local neighborhood of a parameter (see Appendix A for details):

**Definition 1** (The Local Learning Coefficient (LLC),  $\lambda(w^*)$ ). There exists a unique rational<sup>1</sup> number  $\lambda(w^*)$ , a positive integer  $m(w^*)$  and some constant c > 0 such that asymptotically as  $\epsilon \to 0$ ,

$$V(\epsilon) = c \,\epsilon^{\lambda(w^*)} (-\log \epsilon)^{m(w^*)-1} + o(\epsilon^{\lambda(w^*)} (-\log \epsilon)^{m(w^*)-1}).$$

$$(4)$$

We call  $\lambda(w^*)$  the Local Learning Coefficient (LLC), and  $m(w^*)$  the local multiplicity.

The volume scaling relationship in Equation (4) is an important result in singularity theory, see for example Theorem 7.6 of Arnold et al. (2012). In the context of statistical learning theory, Theorem 7.1 and Main Theorem 6.2 in Watanabe (2009) establish a deep connection between the scaling exponent  $\lambda(w^*)$  and learning-theoretic quantities of singular models.

In the case where  $m(w^*) = 1$ , the formula simplifies, and

$$V(\epsilon) \propto \epsilon^{\lambda(w^*)} \tag{5}$$

Thus, the LLC  $\lambda(w^*)$  is the (asymptotic) volume scaling exponent near a minimum  $w^*$  in the loss landscape: increasing the error tolerance by a factor of aincreases the volume by a factor of  $a^{\lambda(w^*)}$ . Applying Equation (3), the number of bits needed to specify  $V(\epsilon)$  within  $B(w^*)$ , for sufficiently small  $\epsilon$  and in the case  $m(w^*) = 1$ , is approximated by

$$-\lambda(w^*)\log_2\epsilon + O(\log_2\log_2\epsilon)$$

Informally, the LLC tells us the number of additional bits needed to halve an already small error of  $\epsilon$ :

$$-\log_2 \left[ V(\frac{\epsilon}{2})/V(\epsilon) \right] \approx -\lambda(w^*) \log_2 \frac{\epsilon}{2} + \lambda(w^*) \log_2 \epsilon$$
$$= \lambda(w^*)$$

See Figure 2 for simple examples of the LLC as a scaling exponent. We note that, in contrast to the regular case, the scaling exponent  $\lambda(w^*)$  depends on the underlying data distribution q(x, y).

The (global) learning coefficient  $\lambda$  was defined by Watanabe (2001). If  $w^*$  is taken to be a global minimum of the population loss L(w), and the neighbourhood  $B(w^*)$  is taken to be the entire parameter space W, then we obtain the global learning coefficient  $\lambda$  as the scaling exponent of the volume  $V(\epsilon)$ . The learning coefficient and related quantities like the WBIC (Watanabe, 2013) have historically seen significant application in Bayesian model selection (e.g. Endo et al., 2020; Fontanesi et al., 2019; Hooten and Hobbs, 2015; Sharma, 2017; Kafashan et al., 2021; Semenova et al., 2020).

In Appendix C, we prove that the LLC is invariant to *local diffeomorphism*: that is, roughly, a locally smooth and invertible change of variables. This property is motivated by the desire that a good complexity measure should not be confounded by superficial differences in how a model is represented: two models which are essentially the same should have the same complexity.

<sup>&</sup>lt;sup>1</sup>The fact that  $\lambda(w^*)$  is *rational-valued*, and not real-valued, as one would naively assume, is a deep fact of algebraic geometry derived from the celebrated Hironaka's resolution of singularities.

It is important to note that the LLC can differ between networks implementing the same function, even between different parameter choices within the same architecture that result in the same input-output map. This directly relates to the information-theoretic intuition of the LLC: it quantifies the number of bits needed to locally improve parameter precision for the desired function, reflecting how discoverable or learnable the function is from data. Consider how specifying the function y = x in a zero-parameter model that contains only this hypothesis is much simpler than specifying the same function in a 3-parameter model  $y = a + bx + cx^2$  or within a deep network. We view this as a feature rather than a limitation. The LLC measures the complexity of a network's specific implementation given a parameter setting, rather than measuring the complexity based solely on the inputoutput behavior. This aligns with the intuition that if two networks achieve the same task but one does so through a more convoluted process, the more convoluted implementation should be recognized as more complex.

# 4 LLC ESTIMATION

Having established the intuition behind the theoretical LLC in terms of volume scaling, we now turn to the task of estimating it. As described in the introduction, the LLC is a coefficient in the asymptotic expansion of the local free energy. It is this fact that we leverage for estimation. There is a mathematically rigorous link between volume scaling and the appearance of the LLC in the local free energy which we will not discuss in detail; see (Watanabe, 2009, Theorem 7.1).

We first introduce what we call the idealized LLC estimator which is theoretically sound, albeit intractable from a computational point of view. In the sections that follow the introduction of the idealized LLC estimator, we walk through the steps taken to engineer a practically implementable version of the idealized LLC estimator.

#### 4.1 Idealized LLC estimator

Consider the following integral

$$Z_n(B_\gamma(w^*)) = \int_{B_\gamma(w^*)} \exp\{-nL_n(w)\}\varphi(w)\,dw, \quad (6)$$

where  $L_n(w)$  is again the sample negative log likelihood,  $B_{\gamma}(w^*)$  denotes a small ball of radius  $\gamma$  around  $w^{*2}$ , and  $\varphi$  is a prior over model parameters w. If (6) is high, there is high posterior concentration around  $w^*$ . In this sense, (6) is a measure of the concentration of low-loss solutions near  $w^*$ . Next consider a log transformation of  $Z_n(B_{\gamma}(w^*))$  with a negative sign, i.e.,

$$F_n(B_\gamma(w^*)) = -\log Z_n(B_\gamma(w^*)).$$
 (7)

This quantity is sometimes called (negative) log marginal likelihood or free energy, depending on the discipline. Given a local minimum  $w^*$  of the population negative log likelihood L(w), it can be shown using the machinery of SLT that, asymptotically in n, we have

$$F_n(B_{\gamma}(w^*)) = n \underbrace{L_n(w^*)}_{\text{energy}} + \underbrace{\lambda(w^*)}_{\text{entropy}} \log n + O_p(\log \log n)$$
(8)

where  $\lambda(w^*)$  is the theoretical LLC expounded in Section 3. Remarkably, the asymptotic approximation in (8) holds even for singular models such as neural networks; further discussion on this can be found in Appendix B. The asymptotic approximation in (8) suggests that a reasonable estimator of  $\lambda(w^*)$  might come from re-arranging (8) to give what we call the idealized LLC estimator,

$$\hat{\lambda}^{\text{idealized}}(w^*) = \frac{F_n(B_\gamma(w^*)) - nL_n(w^*)}{\log n}.$$
 (9)

But as indicated by the name, the idealized LLC estimator cannot be easily implemented; computing or even MCMC sampling from the posterior to estimate  $F_n(B_{\gamma}(w^*))$  is made no less challenging by the need to confine sampling to the neighborhood  $B_{\gamma}(w^*)$ . In what follows, we use the idealized LLC estimator as inspiration for a practically-minded LLC estimator.

#### 4.2 Laplace no more

We take a moment to contrast the idealized LLC estimator with the estimator that would result if we had used the standard (and incorrect) Laplace approximation to derive the asymptotic expansion of the local free energy  $F_n(B_{\gamma}(w^*))$ .

Prior works including Zhang et al. (2018) and Le (2018) rely on the Laplace approximation to arrive at the result

$$F_n(B_{\gamma}(w^*)) = nL_n(w^*) + \frac{d}{2}\log n + \frac{1}{2}\log\det H(w^*) + O_P(1),$$
(10)

where H is the Hessian of the loss. However, a critical issue arises with (10) when the model is singular – the log determinant term becomes ill-defined. Both aforementioned works acknowledge this limitation. Zhang

<sup>&</sup>lt;sup>2</sup>Note that in contrast to the notation  $B(w^*, \epsilon)$  used in Section 2,  $B_{\gamma}(w^*)$  denotes the set of all parameters within distance  $\gamma$  of  $w^*$  instead of the set of all parameters within loss  $\epsilon$  of  $w^*$ .

et al. (2018) explicitly remark that "a more nuanced metric is needed to characterise flat minima with singular Hessian matrices," and Le (2018) state in their introduction that "to compute the (model) evidence, we must carefully account for this degeneracy".

Our work directly addresses these open questions. The correct asymptotics in (8) demonstrates that it is the log n term that is at play when the model is singular, indicating the energy-entropy competition in singular models is far more pronounced than in regular models. Furthermore, we show the correct notion of complexity is given by the LLC, not the parameter count d/2 nor the Hessian determinant.

# 4.3 Gaussian confinement instead of hard boundary

The first step towards a practically-minded LLC estimator is to circumvent the hard constraint posed by the neighborhood  $B_{\gamma}(w^*)$ . To this end, we introduce a localizing Gaussian prior as a practical alternative to hard enforcement of the domain of integration given by  $B_{\gamma}(w^*)$ . Specifically, let

$$\varphi_{\gamma}(w) \propto \exp\left\{-\frac{\gamma}{2}||w||_{2}^{2}\right\}$$

be a Gaussian prior centered at the origin with scale parameter  $\gamma > 0$ . We replace (6) with

$$Z_n(w^*,\gamma) = \int \exp\{-nL_n(w)\}\varphi_{\gamma}(w-w^*)\,dw$$

which, for  $\beta = 1$ , can also be recognized as the normalizing constant to the tempered posterior distribution given by

$$p(w|w^*,\beta,\gamma) \propto \exp\left\{-n\beta L_n(w) -\frac{\gamma}{2}||w-w^*||_2^2\right\},$$
(11)

where  $\beta > 0$  plays the role of an inverse temperature. Large values of  $\gamma$  force the posterior distribution in (11) to stay close to  $w^*$ . A word on the notation  $p(w|w^*, \beta, \gamma)$ : this is a distribution in w solely, the parameters  $w^*, \beta, \gamma$  are fixed, hence the normalizing constant to (11) is an integral over w only. As  $Z_n(w^*, \gamma)$ is to be viewed as a proxy to (6), we shall accordingly treat

$$F_n(w^*,\gamma) \coloneqq -\log Z_n(w^*,\gamma)$$

as a proxy for  $F_n(B_{\gamma}(w^*))$  in (7). Although it is tempting at this stage to simply drop  $F_n(w^*, \gamma)$  into the idealized LLC estimator in place of  $F_n(B_{\gamma}(w^*))$ , we have to address estimation of  $F_n(w^*, \gamma)$ , the subject of the next section.

#### 4.4 The LLC estimator

Let us denote the expectation of a function f(w)with respect to the posterior distribution in (11) as  $\mathbb{E}_{w|w^*,\beta,\gamma}f(w) \coloneqq \int f(w)p(w|w^*,\beta,\gamma) \, dw$ . Consider the quantity

$$\mathbf{E}_{w|w^*,\beta^*,\gamma}[nL_n(w)] \tag{12}$$

where the inverse temperature is deliberately set to  $\beta^* = 1/\log n$ . The quantity in (12) may be regarded as a localized version of the widely applicable Bayesian information criterion (WBIC) first introduced in (Watanabe, 2013). It can be shown that (12) is a good estimator of  $F_n(w^*, \gamma)$  in the following sense: the leading order terms of (12) match those of  $F_n(w^*, \gamma)$  when we perform an asymptotic expansion in sample size n. This justifies using (12) to estimate  $F_n(w^*, \gamma)$ . See Appendix D for further discussion.

Going back to (9), we approximate  $F_n(B_{\gamma}(w^*))$  first with  $F_n(w^*, \gamma)$ , which is further estimated by (12). We are finally ready to define the LLC estimator.

**Definition 2** (Local Learning Coefficient (LLC) estimator). Let  $w^*$  be a local minimum of L(w). Let  $\beta^* = 1/\log n$ . The associated local learning coefficient estimator is given by

$$\hat{\lambda}(w^*) \coloneqq n\beta^* \left[ \mathbb{E}_{w|w^*,\beta^*,\gamma} L_n(w) - L_n(w^*) \right].$$
(13)

Note that  $\hat{\lambda}(w^*)$  depends on  $\gamma$  but we have suppressed this in the notation. Let us ponder the pleasingly simple form that is the LLC estimator. The expectation term in (13) is a measure of the loss  $L_n$  under perturbation near  $w^*$ . If the perturbed loss, under this expectation, is very close to  $L_n(w^*)$ , then  $\hat{\lambda}(w^*)$  is small. This accords with our intuition that if  $w^*$  is simple, its loss should not change too much under reasonable perturbations. Finally we note that in applications, we use the empirical loss  $L_n(w)$  to determine a critical point of interest, i.e.,  $\hat{w}_n^* \coloneqq \arg\min_w L_n(w)$ . We lose something by plugging in  $\hat{w}_n^*$  to (13) directly since we end up using the dataset  $\mathcal{D}_n$  twice. However we do not observe adverse effects in our experiments, see Figure 3 for an example.

#### 4.5 The SGLD-based LLC estimator

The LLC estimator defined in (13) is not prescriptive as to how the expectation with respect to the posterior distribution should actually be approximated. A wide array of MCMC techniques are possible. However, to be able to estimate the LLC at the scale of modern deep learning, we must look at efficiency. In practice, the computational bottleneck to implementing (13) is the MCMC sampler. In particular, traditional MCMC samplers must compute log-likelihood gradients across the *entire* training dataset, which is prohibitively expensive at modern dataset sizes.

If one modifies these samplers to take minibatch gradients instead of full-batch gradients, this results in *stochastic-gradient MCMC*, the prototypical example of which is (Welling and Teh, 2011)'s Stochastic Gradient Langevin Dynamics (SGLD). The computational cost of this sampler is much lower: roughly the cost of a single SGD step times the number of samples required. The standard SGLD update applied to sampling (11) at the optimal temperature  $\beta^*$  required for the LLC estimator is given by

$$\Delta w_t = \frac{\epsilon}{2} \left( \frac{\beta^* n}{m} \sum_{(x,y) \in B_t} \nabla \log p(y|x, w_t) + \gamma(w^* - w_t) \right) + N(0, \epsilon)$$

where  $B_t = \{(x_i, y_i)\}_{i=1}^m$  is a randomly sampled minibatch of samples of size m for step t, and  $\epsilon$  controls both the variance of injected Gaussian noise  $N(0, \epsilon)$ and step size. Crucially, the log-likelihood gradient is evaluated using mini-batches. In practice, we choose to shuffle the dataset once and partition it into a sequence of size m segments as minibatches instead of drawing fresh random samples of size m.

Let us now suppose we have obtained T approximate samples  $\{w_1, w_2, \ldots, w_T\}$  of the tempered posterior distribution at inverse temperature  $\beta^*$  via SGLD. This is usually taken from the SGLD trajectory after burnin. We can then form what we call the **SGLD-based LLC estimator**,

$$\hat{\lambda}^{\text{SGLD}}(w^*) := n\beta^* \left[ \frac{1}{T} \sum_{t=1}^T L_n(w_t) - L_n(w^*) \right].$$
(14)

For further computation saving, we also recycle the forward passes that compute  $L_m(w_t)$ , which is required for computing  $\nabla_w L_m(w_t)$  via backpropagation, as unbiased estimate of  $L_n(w_t)$ . Here by  $L_m(w_t)$ , we mean  $-\frac{1}{m}\sum_{(x,y)\in B_t}\log p(y|x,w_t)$ , though the notation suppresses the dependence on  $B_t$  for brevity. Pseudocode for this minibatch version of the SGLD-based LLC estimator is provided in Appendix G. Henceforth, when we say the SGLD-based LLC estimator, we are referring to the minibatch version. In Appendix H, we give a comprehensive guide on best practices for implementing the SGLD-based LLC estimator including choices for  $\gamma, \epsilon$ , number of SGLD iterations, and required burn-in.

# 5 EXPERIMENTS

The goal of our experiments is to give evidence that the LLC estimator is accurate, scalable and can reveal insights on deep learning practice. Throughout the experiments, we implement the minibatch version of the SGLD-based LLC estimator presented in the pseudo-algorithm in Appendix  $G^3$ .

There are also a number of experiments we performed that are relegated to the appendices due to space constraints. They include

- deploying LLC estimation on transformer language models (Section L)
- an experiment on a small ReLU network verifying that SGLD sampler is just as accurate as one that uses full-batch gradients (Appendix M)
- an experiment comparing the optimizers SGD and entropy-SGD with the latter having intimate connection to our notion of complexity (Appendix N)
- an experiment verifying the scaling invariance property (Appendix O) set out in Appendix C

Every experiment described in the main text below has an accompanying section in the Appendix that offers full experimental details, further discussion, and possible additional figures/tables.

#### 5.1 LLC for Deep Linear Networks (DLNs)

In this section, we verify the *accuracy* and *scalability* of our LLC estimator against theoretical LLC values in deep linear networks (DLNs) up to 100M parameters. Recall DLNs are fully-connected feedforward neural networks with identity activation function. The input-output behavior of a DLN is obviously trivial; it is equivalent to a single-layer linear network obtained by multiplying together the weight matrices. However, the *geometry* of such a model is highly non-trivial — in particular, the optimization dynamics and inductive biases of such networks have seen significant research interest (Saxe et al., 2013; Ji and Telgarsky, 2018; Arora et al., 2018).

Thus one reason we chose to study DLN model complexity is because they have long served as an important sandbox for deep learning theory. Another key factor is the recent clarification of theoretical LLCs in DLNs by Aoyagi (2024), making DLNs the most realistic setting where theoretical learning coefficients are available. The significance of Aoyagi (2024) lies in the substantial technical difficulty of deriving theoretical global learning coefficients (and, by extension, theoretical LLCs) which means that these coefficients

<sup>&</sup>lt;sup>3</sup>Code for the experiments can be found at this link: https://github.com/edmundlth/local\_learning\_ coefficient/tree/aistats2025\_submission.

are generally unavailable except in a few exceptional cases, with most of the research conducted decades ago (Yamazaki and Watanabe, 2005b; Aoyagi et al., 2005; Yamazaki and Watanabe, 2003, 2005a). Further details and discussion of the theoretical result in Aoyagi (2024) can be found in Appendix I.



Figure 3: Estimated LLC against true learning coefficient; model dimension shown in color. The estimated LLCs accurately measures the learning coefficient  $\lambda$  up to 100 million parameters in deep linear networks, as compared to known theoretical values (dashed line). See Figure J.1 for linear-scale plots.

The results are summarized in Figure 3. Our LLC estimator is able to accurately estimate the learning coefficient in DLNs up to 100M parameters. Further experimental details and results for this section are detailed in Appendix J. Overall it is quite remarkable that our LLC estimator, after the series of engineering steps described in Section 4, maintains such high levels of accuracy.

#### 5.2 LLC for ResNet

Here, we empirically test whether implicit regularization effectively induces a preference for simplicity as manifested by a *lower* LLC. We examine the effects of SGD learning rate, batch size, and momentum on the LLC. Note that we do so in isolation, i.e., we do not look at interactions between these factors. For instance, when we vary the learning rate, we hold everything else constant.

We perform experiments on ResNet18 trained on CI-FAR10 and show the results in Figure 1. We see that the training loss reaches zero in most instances and therefore on its own cannot distinguish between the effect of the implicit regularization. In contrast, we see that there is a consistent pattern of "stronger implicit regularization = higher test accuracy = lower LLC". Specifically, higher learning rate, lower batch size, higher momentum all apply stronger implicit regularization which is reflected in lower LLC. Full experimental details can be found in Appendix K. Note that in Figures 1 we employed SGD without momentum in the top two rows. We repeat the experiments in these top two rows for SGD with momentum; the associated results in Figure K.1 support very similar conclusions. We also conducted some explicit regularization experiments involving L2 regularization (Figure K.2 in Appendix K) and again conclude that stronger regularization is accompanied by lower LLC.

In contrast to the LLC estimates for DLN, the LLC estimates for ResNet18 cannot be calibrated as we do not know the true LLC values. With many models in practice, we find value in the LLC estimates by comparing their relative values between models with shared context. For instance, when comparing LLC values we might hold everything constant while varying one factor such as SGD batch size, learning rate, or momentum as we have done in this section.

# 6 RELATED WORK

We briefly cover related work here; more detail may be found in Appendix E. The primary reference for the theoretical foundation of this work, known as SLT, is Watanabe (2009). The global learning coefficient, first introduced by (Watanabe, 2001), provides the asymptotic expansion of the free energy, which is equivalent to the negative log Bayes marginal likelihood, an all-important important quantity in Bayesian analysis. Subsequent research has utilized algebraic-geometric tools to calculate the learning coefficient for various machine learning models (Yamazaki and Watanabe, 2005b; Aoyagi et al., 2005; Aoyagi, 2024; Yamazaki and Watanabe, 2003, 2005a). SLT has also enhanced the understanding of model selection criteria in Bayesian statistics. Of particular relevance to this work is Watanabe (2013), which introduced the WBIC estimator of free energy. This estimator has been applied in various practical settings (Endo et al., 2020; Fontanesi et al., 2019; Hooten and Hobbs, 2015; Sharma, 2017; Kafashan et al., 2021; Semenova et al., 2020).

The LLC can be seen as a singularity-aware version of basin broadness measures, which attempt to connect geometric "broadness" or "flatness" with model complexity (Hochreiter and Schmidhuber, 1997; Jiang et al., 2019). In particular, the LLC *estimator* bears resemblance to PAC-Bayes inspired flatness/sharpness measures (Neyshabur et al., 2017), but takes into account the non-Gaussian nature of the posterior distribution in singular models. This work also relates to the study of complex loss landscapes using singularity theory. As such, it relates to existing work such as Fukumizu and Amari (2000); Simsek et al. (2021) that studied the loss landscape of neural networks and the effect of network hierarchical structure and symmetries on the geometry of critical points.

The LLC is set apart from classic model complexity measures such as Rademacher complexity (Koltchinskii and Panchenko, 2000) and the VC dimension (Vapnik and Chervonenkis, 1971) because it measures the complexity of a specific model p(y|x, w) rather than the complexity over the function class  $\{f(x|w) : w\}$  where f is the neural network function. This makes the LLC an appealing tool for understanding the interplay between function class, data properties, and training heuristics.

Concurrent work by (Chen et al., 2023a) proposes a measure called the learning capacity, which can be viewed as a finite-n version of the learning coefficient, and investigates its behavior as a function of n.

# 7 OUTLOOK

An exciting direction for future research is to study the role of the LLC in detecting phase transitions and emergent abilities in deep learning models. A first step in this direction was undertaken in (Chen et al., 2023b) where the energy (training loss) and entropy (both estimated and theoretical LLC) were tracked as training progresses; it was observed that energy and entropy proceed along staircases in opposing directions. Further, Hoogland et al. (2024) showed how the estimated LLC can be used to detect phase transitions in the formation of in-context learning in transformer language models. It is natural to wonder if the LLC could shed light on the hypothesis that SGD-trained neural networks sequentially learn the target function with a "saddle-to-saddle" dynamic. Previous theoretical works had to devise complexity measures on a case-by-case basis (Abbe et al., 2023; Berthier, 2023). We posit that the free energy perspective could offer a more unified and general approach to understanding the intricate dynamics of learning in deep neural networks by accounting for competition between model fit and model complexity during training.

# Acknowledgments

We thank the reviewers for their constructive feedback that improved the quality and clarity of this manuscript. SW is the recipient of an Australian Research Council Discovery Early Career Researcher Award (project number DE200101253) funded by the Australian Government. SW is also partially funded by an unrestricted gift from Google. ZF is supported by the Melbourne Research Scholarship and the Rowden White Scholarship. The authors thank Matthew Farrugia-Roberts, Liam Carroll, Garrett Baker, Quintin Pope, Simon Pepin Lehalleur, and Jesse Hoogland for discussions and valuable feedback on initial drafts.

#### References

- Abbe, E., Adserà, E. B., and Misiakiewicz, T. (2023). SGD learning on neural networks: leap complexity and saddle-to-saddle dynamics. In Neu, G. and Rosasco, L., editors, *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pages 2552– 2623. PMLR.
- Aoyagi, M. (2024). Consideration on the learning efficiency of multiple-layered neural networks with linear units. *Neural Networks*, page 106132.
- Aoyagi, M., Watanabe, S., et al. (2005). Resolution of singularities and the generalization error with Bayesian estimation for layered neural network. *IE-ICE Trans*, 88(10):2112–2124.
- Arnold, V. I., Gusein-Zade, S. M., and Varchenko, A. N. (2012). Singularities of Differentiable Maps, volume 2.
- Arora, S., Cohen, N., Golowich, N., and Hu, W. (2018). A convergence analysis of gradient descent for deep linear neural networks. arXiv preprint arXiv:1810.02281.
- Berner, J., Elbrächter, D., and Grohs, P. (2019). How degenerate is the parametrization of neural networks with the ReLU activation function? In *Neural Information Processing Systems*.
- Berthier, R. (2023). Incremental learning in diagonal linear networks. *Journal of Machine Learning Re*search, 24(171):1–26.
- Bialek, W., Nemenman, I., and Tishby, N. (2001). Predictability, complexity, and learning. *Neural computation*, 13(11):2409–2463.
- Blalock, D., Gonzalez Ortiz, J. J., Frankle, J., and Guttag, J. (2020). What is the state of neural network pruning? *Proceedings of Machine Learning* and Systems, 2:129–146.
- Chaudhari, P., Choromanska, A., Soatto, S., LeCun, Y., Baldassi, C., Borgs, C., Chayes, J., Sagun, L., and Zecchina, R. (2019). Entropy-SGD: Biasing gradient descent into wide valleys. *Journal* of Statistical Mechanics: Theory and Experiment, 2019(12):124018.

- Chen, D., Chang, W., and Chaudhari, P. (2023a). Learning capacity: A measure of the effective dimensionality of a model.
- Chen, Z., Lau, E., Mendel, J., Wei, S., and Murfet, D. (2023b). Dynamical versus Bayesian phase transitions in a toy model of superposition.
- Deng, L. (2012). The MNIST database of handwritten digit images for machine learning research. *IEEE* Signal Processing Magazine, 29(6):141–142.
- Dherin, B., Munn, M., Rosca, M., and Barrett, D. G. T. (2022). Why neural networks find simple solutions: the many regularizers of geometric complexity.
- Dinh, L., Pascanu, R., Bengio, S., and Bengio, Y. (2017). Sharp minima can generalize for deep nets. In *International Conference on Machine Learning*, pages 1019–1028. PMLR.
- Endo, A., Abbott, S., Kucharski, A. J., Funk, S., et al. (2020). Estimating the overdispersion in COVID-19 transmission using outbreak sizes outside China. *Wellcome Open Research*, 5.
- Farrugia-Roberts, M. (2023). Functional equivalence and path connectivity of reducible hyperbolic tangent networks. In *Thirty-seventh Conference on Neural Information Processing Systems.*
- Fefferman, C. (1994). Reconstructing a neural net from its output. Revista Matemática Iberoamericana, 10(3):507–555.
- Fontanesi, L., Gluth, S., Spektor, M. S., and Rieskamp, J. (2019). A reinforcement learning diffusion decision model for value-based decisions. *Psychonomic Bulletin & Review*, 26(4):1099–1121.
- Fukumizu, K. (1996). A regularity condition of the information matrix of a multilayer perceptron network. *Neural Networks*, 9(5):871–879.
- Fukumizu, K. and Amari, S. (2000). Local minima and plateaus in hierarchical structures of multilayer perceptrons. *Neural Netw.*, 13(3):317–327.
- Gao, L., Biderman, S., Black, S., Golding, L., Hoppe, T., Foster, C., Phang, J., He, H., Thite, A., Nabeshima, N., et al. (2020). The Pile: an 800GB dataset of diverse text for language modeling. arXiv preprint arXiv:2101.00027.
- Grünwald, P. and Roos, T. (2019). Minimum description length revisited. *International journal of mathematics for industry*, 11(01):1930001.
- Haario, H., Saksman, E., and Tamminen, J. (1999). Adaptive proposal distribution for random walk Metropolis algorithm. *Computational Statistics*, 14:375–395.

- Haario, H., Saksman, E., and Tamminen, J. (2001). An adaptive metropolis algorithm. *Bernoulli*, pages 223–242.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 770–778. IEEE.
- Hinton, G., Vinyals, O., and Dean, J. (2015). Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531.
- Hironaka, H. (1964). Resolution of Singularities of an Algebraic Variety Over a Field of Characteristic Zero: I. Annals of Mathematics, 79(1):109–203.
- Hochreiter, S. and Schmidhuber, J. (1997). Flat Minima. Neural Computation, 9(1):1–42.
- Hoogland, J. and van Wingerden, S. (2023). You're measuring model complexity wrong.
- Hoogland, J., Wang, G., Farrugia-Roberts, M., Carroll, L., Wei, S., and Murfet, D. (2024). The developmental landscape of in-context learning.
- Hooten, M. B. and Hobbs, N. T. (2015). A guide to Bayesian model selection for ecologists. *Ecological Monographs*, 85(1):3–28.
- Imai, T. (2019a). Estimating real log canonical thresholds. arXiv preprint arXiv:1906.01341.
- Imai, T. (2019b). On the overestimation of widely applicable Bayesian information criterion. arXiv preprint arXiv:1908.10572.
- Iriguchi, R. and Watanabe, S. (2007). Estimation of poles of zeta function in learning theory using Padé approximation. In Artificial Neural Networks-ICANN 2007: 17th International Conference, Porto, Portugal, September 9-13, 2007, Proceedings, Part I 17, pages 88–97. Springer.
- Ji, Z. and Telgarsky, M. (2018). Gradient descent aligns the layers of deep linear networks. arXiv preprint arXiv:1810.02032.
- Jiang, Y., Neyshabur, B., Mobahi, H., Krishnan, D., and Bengio, S. (2019). Fantastic generalization measures and where to find them. In *International Conference on Learning Representations*.
- Jules, T., Brener, G., Kachman, T., Levi, N., and Bar-Sinai, Y. (2023). Charting the topography of the neural network landscape with thermal-like noise.
- Kafashan, M., Jaffe, A. W., Chettih, S. N., Nogueira, R., Arandia-Romero, I., Harvey, C. D., Moreno-Bote, R., and Drugowitsch, J. (2021). Scaling of sensory information in large neural populations shows signatures of information-limiting correlations. *Nature Communications*, 12(1):473.

- Koltchinskii, V. and Panchenko, D. (2000). Rademacher processes and bounding the risk of function learning. In Giné, E., Mason, D. M., and Wellner, J. A., editors, *High Dimensional Probability II*, pages 443–457, Boston, MA. Birkhäuser Boston.
- Krizhevsky, A. (2009). Learning multiple layers of features from tiny images.
- Kůrková, V. and Kainen, P. C. (1994). Functionally equivalent feedforward neural networks. *Neural Computation*, 6:543–558.
- LaMont, C. H. and Wiggins, P. A. (2019). Correspondence between thermodynamics and inference. *Physical Review E*, 99(5):052140.
- Le, S. L. S. a. Q. V. (2018). A Bayesian perspective on generalization and stochastic gradient descent. In International Conference on Learning Representations.

Nanda, N. and Bloom, J. (2022). TransformerLens.

- Neyshabur, B., Bhojanapalli, S., Mcallester, D., and Srebro, N. (2017). Exploring generalization in deep learning. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Petersen, P. C., Raslan, M., and Voigtlaender, F. (2020). Topological properties of the set of functions generated by neural networks of fixed size. *Foundations of Computational Mathematics*, 21:375 – 444.
- Phuong, M. and Lampert, C. H. (2020). Functional vs. parametric equivalence of ReLU networks. In International Conference on Learning Representations.
- Press, O., Smith, N. A., and Lewis, M. (2021). Shortformer: Better language modeling using shorter inputs.
- Roberts, G. O. and Rosenthal, J. S. (1998). Optimal scaling of discrete approximations to Langevin diffusions. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 60(1):255–268.
- Saxe, A. M., McClelland, J. L., and Ganguli, S. (2013). Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. arXiv preprint arXiv:1312.6120.
- Semenova, E., Williams, D. P., Afzal, A. M., and Lazic, S. E. (2020). A Bayesian neural network for toxicity prediction. *Computational Toxicology*, 16:100133.
- Sharma, S. (2017). Markov chain Monte Carlo methods for Bayesian data analysis in astronomy. Annual Review of Astronomy and Astrophysics, 55:213–259.

- Simsek, B., Ged, F., Jacot, A., Spadaro, F., Hongler, C., Gerstner, W., and Brea, J. (2021). Geometry of the loss landscape in overparameterized neural networks: Symmetries and invariances. In *International Conference on Machine Learning*, pages 9722–9732. PMLR.
- Skalse, J. (2023). My criticism of singular learning theory.
- Sussmann, H. J. (1992). Uniqueness of the weights for minimal feedforward nets with a given input-output map. Neural Networks, 5(4):589–593.
- Valle-Perez, G., Camargo, C. Q., and Louis, A. A. (2018). Deep learning generalizes because the parameter-function map is biased towards simple functions. arXiv preprint arXiv:1805.08522.
- Vapnik, V. N. and Chervonenkis, A. Y. (1971). On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability* & Its Applications, 16(2):264–280.
- Waskom, M. L. (2021). seaborn: statistical data visualization. Journal of Open Source Software, 6(60):3021.
- Watanabe, S. (2001). Algebraic analysis for nonidentifiable learning machines. Neural Computation, 13(4):899–933.
- Watanabe, S. (2009). Algebraic Geometry and Statistical Learning Theory. Cambridge University Press, USA.
- Watanabe, S. (2010). Asymptotic learning curve and renormalizable condition in statistical learning theory. *Journal of Physics: Conference Series*, 233:012014.
- Watanabe, S. (2013). A Widely Applicable Bayesian Information Criterion. Journal of Machine Learning Research, 14(Mar):867–897.
- Watanabe, S. (2018). Mathematical theory of Bayesian statistics. CRC Press.
- Wei, S., Murfet, D., Gong, M., Li, H., Gell-Redman, J., and Quella, T. (2022). Deep learning is singular, and that's good. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–14.
- Welling, M. and Teh, Y. W. (2011). Bayesian learning via stochastic gradient Langevin dynamics. In Proceedings of the 28th International Conference on Machine Learning.
- Xie, S. M., Santurkar, S., Ma, T., and Liang, P. (2023). Data selection for language models via importance resampling. arXiv preprint arXiv:2302.03169.
- Yamazaki, K. and Watanabe, S. (2003). Singularities in mixture models and upper bounds of stochastic complexity. *Neural Networks*, 16(7):1029–1038.

- Yamazaki, K. and Watanabe, S. (2005a). Algebraic geometry and stochastic complexity of hidden Markov models. *Neurocomputing*, 69(1-3):62–84.
- Yamazaki, K. and Watanabe, S. (2005b). Singularities in complete bipartite graph-type Boltzmann machines and upper bounds of stochastic complexities. *IEEE transactions on neural networks*, 16(2):312– 324.
- Yang, G., Hu, E. J., Babuschkin, I., Sidor, S., Liu, X., Farhi, D., Ryder, N., Pachocki, J., Chen, W., and Gao, J. (2022). Tensor programs V: Tuning large neural networks via zero-shot hyperparameter transfer. arXiv preprint arXiv:2203.03466.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2017). Understanding deep learning requires rethinking generalization.
- Zhang, Y., Saxe, A. M., Advani, M. S., and Lee, A. A. (2018). Energy–entropy competition and the effectiveness of stochastic gradient descent in machine learning. *Molecular Physics*, 116(21-22):3214–3223.

# Checklist

- 1. For all models and algorithms presented, check if you include:
  - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model.
    [Yes] Justification: The paper introduces the Local Learning Coefficient (LLC) as a novel complexity measure (Section 3) and develops a scalable estimator for it (Section 4).
  - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes] Justification: As discussed in Section 4.5, the computational efficiency of the algorithm is analyzed, noting that it is almost entirely determined by the computational efficiency of SGLD.
  - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes] Justification: The code for experiments in the main text is provided in an anonymized repository.
- 2. For any theoretical claim, check if you include:
  - (a) Statements of the full set of assumptions of all theoretical results. [Yes] Justification: Appendix A gives background assumptions from singular learning theory.
  - (b) Complete proofs of all theoretical results. [Yes] Justification: Appendices provide proofs for theoretical results, including welldefinedness of the LLC (Appendix B) and

reparameterization invariance of the LLC (Appendix C).

- (c) Clear explanations of any assumptions. [Yes] Justification: Assumptions are explained in the Appendices, particularly Appendix A which gives background assumptions from singular learning theory.
- 3. For all figures and tables that present empirical results, check if you include:
  - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes] Justification: Code and instructions for the experiments in the main text are provided in an anonymized repository.
  - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes] Justification: Each set of experiments has its own appendix that provides the required experimental details, in addition to the high-level details given in the main text.
  - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes] Justification: Error bars or histograms are shown and explained for every set of experiments with the factors of variability explained in the appendices that provide experimental details.
  - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes] Justification: The compute resources required for each set of experiments are provided in Appendix P.
- 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
  - (a) Citations of the creator If your work uses existing assets. [Yes] Justification: Standard datasets like the Pile, MNIST and CIFAR-10 are cited when used.
  - (b) The license information of the assets, if applicable. [Not applicable] Justification: These standard datasets are open source.
  - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes] Justification: Code has been open-sourced and is available in a repository.
  - (d) Information about consent from data providers/curators. [Not applicable] Justification: This project has no data providers or curators.

- (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable] Justification: The paper does not include sensitive content.
- 5. If you used crowdsourcing or conducted research with human subjects, check if you include:
  - (a) The full text of instructions given to participants and screenshots. [Not Applicable] Justification: The paper does not involve crowdsourcing or human subject research.
  - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable] Justification: The paper does not involve human subject research requiring IRB approval.
  - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable] Justification: The paper does not involve crowdsourcing or human subject research.

## A Background on Singular Learning Theory

Most models in machine learning are singular: they contain parameters where the Fisher information matrix is singular. While these parameters where the Fisher information is degenerate form a measure zero subset, their effect is far from negligible. Singular Learning Theory (SLT) shows that the geometry in the neighbourhood of these degenerate points determines the asymptotics of learning (Watanabe, 2009). The theory explains observable effects of degeneracy in common machine learning models under practical settings (Watanabe, 2018) and has the potential to account for important phenomena in deep learning (Wei et al., 2022). The central quantity of SLT is the learning coefficient  $\lambda$ . Many notable SLT results are conveyed through the learning coefficient which can be thought of as the complexity of the model class relative to the true distribution. In this section we carefully define the (global) learning coefficient  $\lambda$ , which we then contrast with the *local* learning coefficient  $\lambda(w^*)$ .

For notational simplicity, we consider here the unsupervised setting where we have the model p(x|w) parameterised by a compact parameter space  $W \subset \mathbb{R}^d$ . We assume fundamental conditions I and II of (Watanabe, 2009, §6.1, §6.2). In particular W is defined by a finite set of real analytic inequalities, at every parameter  $w \in W$  the distribution p(x|w) should have the same support as the true density q(x), and the prior density  $\varphi(w) = \varphi_1(w)\varphi_2(w)$  is a product of a positive smooth function  $\varphi_1(w)$  and non-negative real analytic  $\varphi_2(w)$ . We refer to

$$(p(x|w), q(x), \varphi(w)) \tag{15}$$

as the model-truth-prior triplet.

Let K(w) be the Kullback-Leibler divergence between the truth and the model

$$K(w) \coloneqq \mathrm{KL}\left(q(x) \| p(x|w)\right) = \int q(x) \log\left[q(x)/p(x|w)\right] dx$$

and define the (average) negative log likelihood to be

$$L(w) \coloneqq -\int q(x)\log p(x|w)dx = K(w) - S$$

where S is the entropy of the true distribution. Set  $K_0 = \inf_{w \in W} K(w)$ . Let

$$W_0 \coloneqq \{w \in W : K(w) = K_0\}$$

be the set of **optimal parameters**. We say the **truth is realizable by the model** if  $K_0 = 0$ . We do not assume that our models are regular or that the truth is realizable, but we do assume the model-truth-prior triplet satisfies the more general condition of **relative finite variance** (Watanabe, 2013). We also assume that there exists  $w_0^*$  in the interior of W satisfying  $K(w_0^*) = K_0$ .

Following Watanabe (2009) we define:

**Definition 3.** The zeta function of (15) is defined for  $\operatorname{Re}(z) > 0$  by

$$\zeta(z) = \int_{W} \left( K(w) - K_0 \right)^z \varphi(w) dw$$

and can be analytically continued to a meromorphic function on the complex plane with poles that are all real, negative and rational (Watanabe, 2009, Theorem 6.6). Let  $-\lambda \in \mathbb{R}$  be the largest pole of  $\zeta$  and m its multiplicity. Then, the **learning coefficient** and its **multiplicity** of the triple (15) are defined to be  $\lambda$  and m respectively.

When p(x|w) is a singular model,  $W_0$  is an analytic variety which is in general positive dimensional (not a collection of isolated points). As long as  $\varphi > 0$  on  $W_0$  the learning coefficient  $\lambda$  is equal to a birational invariant of  $W_0$  known in algebraic geometry as the **Real Log Canonical Threshold (RLCT)**. We will always assume this is the case, and now recall how  $\lambda$  is described geometrically.

With  $W_{\epsilon} = \{w \in W : K(w) - K_0 \leq \epsilon\}$  for sufficiently small  $\epsilon$ , resolution of singularities (Hironaka, 1964) gives us the existence of a birational proper map  $g : M \to W_{\epsilon}$  from an analytic manifold M which monomializes  $K(w) - K_0$  in the following sense, described precisely in (Watanabe, 2009, Theorem 6.5): there are local coordinate

charts  $M_{\alpha}$  covering  $g^{-1}(W_0)$  with coordinates u such that the reparameterisation w = g(u) puts  $K(w) - K_0$  and  $\varphi(w)dw$  into normal crossing form

$$K(g(u)) - K_0 = u_1^{2k_1} \dots u_d^{2k_d}$$
(16)

$$|g'(u)| = b(u)u_1^{h_1}\dots u_d^{h_d}$$
(17)

$$\varphi(w)dw = \varphi(g(u))|g'(u)|du \tag{18}$$

for some positive smooth function b(u). The RLCT of  $K - K_0$  is independent of the (non-unique) resolution map g, and may be computed as (Watanabe, 2009, Definition 6.4)

$$\lambda = \min_{\alpha} \min_{j=1\dots d} \frac{h_j + 1}{2k_j}.$$
(19)

The multiplicity is defined as

$$m = \max_{\alpha} \#\{j : \lambda_j = \lambda\}.$$

For  $P \in W_0$  there exist coordinate charts  $M_{\alpha^*}$  such that g(0) = P and (16), (17) hold. The RLCT of  $K - K_0$  at P is then (Watanabe, 2009, Definition 2.7)

$$\lambda(P) = \min_{\alpha^*} \min_{j=1\dots d} \frac{h_j + 1}{2k_j} \,. \tag{20}$$

We then have the RLCT  $\lambda = \inf_{P \in W_0} \lambda(P)$ . In regular models, all  $\lambda(P)$  are d/2 and hence the RLCT  $\lambda = d/2$ and m = 1, see (Watanabe, 2009, Remark 1.15).

#### A.1 Background assumptions in SLT

There are a few technical assumptions throughout SLT that we shall collect in this section. We should note that these are only sufficient but not necessary conditions for many of the results in SLT. Most of the assumptions can be relaxed on a case by case basis without invalidating the main conclusions of SLT. For in depth discussion see Watanabe (2009, 2010, 2018).

With the same hypotheses as above, the log likelihood ratio

$$r(x,w) := \log \frac{p(x|w_0)}{p(x|w)}$$

is assumed to be an  $L^s(q(x))$ -valued analytic function of w with  $s \ge 2$  that can be extended to a complex analytic function on  $W_{\mathbb{C}} \subset \mathbb{C}^d$ . A more conceptually significant assumption is the condition *relatively finite variance*, which consists of the following two requirements:

- 1. For any optimal parameters  $w_1, w_2 \in W_0$ , we have  $p(x|w_1) = p(x|w_2)$  almost everywhere. This is also known as essential uniqueness.
- 2. There exists c > 0 such that for all  $w \in W$ , we have

$$\mathbb{E}_{q(x)}\left[r(x,w)\right] \ge c\mathbb{E}_{q(x)}\left[r(x,w)^2\right]$$

Note that if the true density q(x) is *realizable* by the model, i.e. there exist  $w^* \in W$  such that  $q(x) = p(x|w^*)$  almost everywhere, then these conditions are automatically satisfied. This includes settings for which the training labels are synthetically generated by passing inputs through a target model (in which case realizability is satisfied by construction), such as our experiments involving DLNs and some of the experiments involving MLPs. For experiments involving real data, it is unclear how reasonable it is to assume that relative finite variance holds.

#### **B** Well-definedness of the theoretical LLC

Here we remark on how we adapt the (global) learning coefficient of Definition 3 to define the local learning coefficient in terms of the poles of a zeta function, and how this relates to the asymptotic volume formula in Definition 1 of the main text.

In addition to the setup described above, now suppose we have a local minimum  $w^*$  of the negative log likelihood L(w) and as in Section 3.1 we assume  $B(w^*)$  to be a closed ball centered on  $w^*$  such that  $L(w^*)$  is the minimum value of L on the ball. We define

$$V := \operatorname{Vol}(B(w^*)) = \int_{B(w^*)} \varphi(w) dw,$$
$$\bar{\varphi}(w) = \frac{1}{V} \varphi(w)$$

then we can form the **local triplet**  $(p, q, \bar{\varphi})$  with parameter space  $B(w^*)$ . Note that W is cut out by a finite number of inequalities between analytic functions, and hence so is  $B(w^*)$ .

Provided  $\varphi(w^*) > 0$  the prior does not contribute to the leading terms of the asymptotic expansions considered, and in particular does not appear in the asymptotic formula for the volume in Definition 1, and so we disregard it in the main text for simplicity.

Assuming relative finite variance of the local triplet, we can apply the discussion of Section A. In particular, we can define the LLC  $\lambda(w^*)$  and the local multiplicity  $m(w^*)$  in terms of the poles of the "local" zeta function

$$\zeta(z, w^*) = \int_{B(w^*)} \left( K(w) - K(w^*) \right)^z \bar{\varphi}(w) dw$$

Note that  $L(w) - L(w^*) = K(w) - K(w^*)$  since K, L differ by S which does not depend on w.

In order for the LLC to be a learning coefficient in its own right, it must be related, asymptotically, to the local free energy in the manner stipulated in (8) in the main text. We verify this next. To derive the asymptotic expansion of the local free energy  $F_n(B(w^*))$ , we assume in addition that  $\lambda(w^*) \leq \lambda(P)$  if  $P \in B(w^*)$  and  $L(P) = L(w^*)$ . That is,  $w^*$  is at least as degenerate as any nearby minimiser. Note that the KL divergence for this triplet is just the restriction of  $K : W \to \mathbb{R}$  to  $B(w^*)$ , but the local triplet has its own set of optimal parameters

$$W_0(w^*) = \{ w \in B(w^*) : L(w) = L(w^*) \}.$$
(21)

Borrowing the proof in Watanabe  $(2009, \S3)$ , we can show that

$$F_n(B(w^*)) = nL_n(w^*) + \lambda(w^*)\log n - (m(w^*) - 1)\log\log n + O_P(1)$$
(22)

where the difference between  $\varphi$  and  $\bar{\varphi}$  contributes a summand log V to the constant term. Note that the condition of relative finite variance is used to establish (22). This explains why we can consider  $\lambda(w^*)$  as a local learning coefficient, following the ideas sketched in (Watanabe, 2009, Section 7.6). The presentation of the LLC in terms of volume scaling given in the main text now follows from (Watanabe, 2009, Theorem 7.1).

## C Reparameterization invariance of the LLC

The LLC is invariant to *local diffeomorphism* of the parameter space - roughly a locally smooth and invertible change of variables.

Note that this automatically implies several weaker notions of invariance, such as rescaling invariance in feedforward neural networks; other e.g. Hessian-based complexity measures have been undermined by their failure to stay invariant to this symmetry (Dinh et al., 2017). We now show how the LLC is invariant to local diffeomorphism. The fact that the LLC is an asymptotic quantity is crucial; this property would not hold for the volume  $V(\epsilon)$  itself, for any value of  $\epsilon$ .

Let  $U \subset W$  and  $\widetilde{U} \subset \widetilde{W}$  be open subsets of parameter spaces W and  $\widetilde{W}$ . A *local diffeomorphism* is an invertible map  $\phi : U \to \widetilde{U}$ , such that both  $\phi$  and  $\phi^{-1}$  are infinitely differentiable. We further require that  $\phi$  respect the loss function: that is, if  $L : U \to \mathbb{R}$  and  $\widetilde{L} : \widetilde{U} \to \mathbb{R}$  are loss functions on each space, we insist that  $L(u) = \widetilde{L}(\phi(u))$ for all  $u \in U$ .

Supposing such a  $\phi$  exists, the statement to be proved is that the LLC at  $u^* \in U$  under L(u) is equal to the LLC at  $\tilde{u}^* = \phi(u^*) \in \tilde{U}$  under  $\tilde{L}(\tilde{u})$ . Define

$$V(\epsilon) = \int_{L(u) - L(u^*) < \epsilon} du, \quad \widetilde{V}(\epsilon) = \int_{\widetilde{L}(\widetilde{u}) - \widetilde{L}(\widetilde{u}^*) < \epsilon} d\widetilde{u}$$

Now note that by the change of variables formula for diffeomorphisms, we have

$$\widetilde{V}(\epsilon) = \int_{L(u)-L(u^*)<\epsilon} |\det D\phi(u)| du$$

where det  $D\phi(u)$  is the Jacobian determinant of  $\phi$  at u.

The fact that  $\phi$  is a local diffeomorphism implies that there exists constants  $c_1, c_2$  such that  $c_1 \leq |\det D\phi(u)| \leq c_2$  for all  $u \in U$ . This means that

$$c_1 V(\epsilon) \le V(\epsilon) \le c_2 V(\epsilon)$$

Finally, applying the definition of the LLC  $\lambda$  and its multiplicity m, and leveraging the fact that this definition is asymptotic as  $\epsilon \to 0$ , we can conclude that

$$V(\epsilon) \propto \widetilde{V}(\epsilon) \propto \epsilon^{\lambda} (-\log(\epsilon))^{m-1}$$

which demonstrates that the LLC is preserved by the local diffeomorphism  $\phi$ .

#### D Consistency of local WBIC

In the main text, we introduced (12) as an estimator of  $F_n(w^*, \gamma)$ . Our motivation for this estimator comes directly from the well-known widely applicable Bayesian information criterion (WBIC) (Watanabe, 2013). In this section, we refer to (12) as the local WBIC and denote it WBIC( $w^*$ )

It is a direct extension of the proofs in Watanabe (2013) to show that the first two terms in the asymptotic expansion of the local WBIC match those of  $F_n(w^*, \gamma)$ . By this we mean that it can be shown that

$$F_n(w^*, \gamma) = nL_n(w^*) + \tilde{\lambda}(w^*)\log n - (m-1)\log\log n + R_n$$

and

WBIC
$$(w^*) = nL_n(w^*) + \tilde{\lambda}(w^*)\log n + U_n\sqrt{\tilde{\lambda}(w^*)\log n/2} + O_P(1).$$
 (23)

This firmly establishes that (12) is a good estimator of  $F_n(w^*, \gamma)$ . However, it is important to understand that we cannot immediately conclude that it is a good estimator of  $F_n(B_{\gamma}(w^*))$ .

We conjecture that  $\tilde{\lambda}(w^*)$  is equal to the LLC  $\lambda(w^*)$  given certain conditions on  $\gamma$ . So far, all our empirical findings suggest this. Detailed proof of this conjecture, in particular ascertaining the exact conditions on  $\gamma$ , will be left as future theoretical work.

#### E Related work

We review both the singular learning theory literature directly involving the learning coefficient, as well as research from other areas of machine learning that may be relevant. This is an expanded version of the discussion in Section 6.

Singular learning theory. Our work builds upon the singular learning theory (SLT) of Bayesian statistics: good references are Watanabe (2009, 2018). The global learning coefficient, first introduced by (Watanabe, 2001), provides the asymptotic expansion of the free energy, which is equivalent to the negative log Bayes marginal likelihood, an all-important important quantity in Bayesian analysis. Later work used algebro-geometric tools to bound or exactly calculate the learning coefficient for a wide range of machine learning models, including Boltzmann machines Yamazaki and Watanabe (2005b), single-hidden-layer neural networks Aoyagi et al. (2005), DLNs Aoyagi (2024), Gaussian mixture models Yamazaki and Watanabe (2003), and hidden Markov models Yamazaki and Watanabe (2005a).

SLT has also enhanced the understanding of model selection criteria in Bayesian statistics. Of particular relevance to this work is Watanabe (2013), which introduced the WBIC estimator of free energy. This estimator has been applied in various practical settings (e.g. Endo et al., 2020; Fontanesi et al., 2019; Hooten and Hobbs, 2015; Sharma, 2017; Kafashan et al., 2021; Semenova et al., 2020). Some of the estimation methodology in this

paper can be seen as a localized extension of the WBIC. Several other papers have explored improvements or alternatives to this estimator Iriguchi and Watanabe (2007); Imai (2019a,b).

**Basin broadness**. The learning coefficient can be seen as a Bayesian version of *basin broadness* measures, which typically attempt to empirically connect notions of geometric "broadness" or "flatness" with model complexity Hochreiter and Schmidhuber (1997); Jiang et al. (2019). However, the evidence supporting the (global) learning coefficient (in the Bayesian setting) is significantly stronger: the learning coefficient provably determines the Bayesian free energy to leading order Watanabe (2009). We expect the utility of the learning coefficient as a geometric measure to apply beyond the Bayesian setting, but whether the connection with generalization will continue to hold is unknown.

Neural network identifiability. A core observation leading to singular learning theory is that the map  $w \mapsto p(x|w)$  from parameters w to statistical models p(x|w) may not be one-to-one (in which case, the model is singular). This observation has been made in parallel by researchers studying *neural network identifiability* Sussmann (1992); Fefferman (1994); Kůrková and Kainen (1994); Phuong and Lampert (2020). Recent work<sup>4</sup> has shown that the degree to which a network is identifiable (or inverse stable) is not uniform across parameter space Berner et al. (2019); Petersen et al. (2020); Farrugia-Roberts (2023). From this perspective, the LLC can be viewed as a quantitative measure of "how identifiable" the network is near a particular parameter.

Statistical mechanics of the loss landscape. A handful of papers have explored related ideas from a statistical mechanics perspective. Jules et al. (2023) use Langevin dynamics to probe the geometry of the loss landscape. Zhang et al. (2018) show how a bias towards "wide minima" may be explained by free energy minimization. These observations may be formalized using singular learning theory LaMont and Wiggins (2019). In particular, the learning coefficient may be viewed as a heat capacity LaMont and Wiggins (2019), and learning coefficient estimation corresponds to measuring the heat capacity by molecular dynamics sampling.

**Other model complexity measures.** The LLC is set apart from a number of classic model complexity measures such as Rademacher complexity (Koltchinskii and Panchenko, 2000), the VC dimension (Vapnik and Chervonenkis, 1971) because the latter measures act on an *entire class of functions* while the LLC measures the complexity of a specific individual function within the context of the function class carved out by the model (e.g. via DNN architecture). This affords the LLC a better position for unraveling the theoretical mysteries of deep learning, which cannot be disentangled from the way in which DNNs are trained or the data that they are trained on.

In the context studied here, our proposed LLC measures the complexity of a trained neural network rather than complexity over the entire function class of neural networks. It is also sensitive to the data distribution, making it ideal for understanding the intricate dance between function class, data properties, and implicit biases baked into different training heuristics.

Similarity to PAC-Bayes. We have just described how the theoretical LLC is the sought-after notion of model complexity coming from earlier works who adopt the energy-entropy competition perspective. Interestingly, the actual LLC estimator also has connections to another familiar notion of model complexity. Among the diverse cast of complexity measures, see e.g., Jiang et al. (2019) for a comprehensive overview of over forty complexity measures in modern deep learning, the LLC estimator bears the most resemblance to PAC-Bayes inspired flatness/sharpness measures (Neyshabur et al., 2017). Indeed, it may be immediately obvious that, other than the scaling of  $n\beta^*$ ,  $\hat{\lambda}(w^*)$  can be viewed as a PAC-Bayes flatness measure which utilises a very specific posterior distribution localised to  $w^*$ . Recall the canonical PAC-Bayes flatness measure is based on

$$\lambda_{\text{PAC-bayes}}(w^*) = \mathcal{E}_{q(w|w^*)}\ell_n(w) - \ell_n(w^*), \qquad (24)$$

where  $\ell_n$  is a general empirical loss function (which in our case is the sample negative log likelihood) and the "posterior" distribution q is often taken to be Gaussian, i.e.,  $q(w|w^*) = \mathcal{N}(w^*, \sigma^2 I)$ . A simple derivation shows us that the quantity in (24), if we use a Gaussian q around  $w^*$ , reduces approximately to  $\frac{1}{2}\sigma^2 Tr(H(w^*))$ , where H is the Hessian of the loss, i.e.,  $H(w^*) = \nabla_w^2 \ell_n(w)|_{w^*}$ . However, for singular models, the posterior distribution around  $w^*$ , e.g., (11), is decidedly not Gaussian. This calls into question the standard choice of the Gaussian posterior in (24).

Learning capacity. Finally, we briefly discuss concurrent work that measures a quantity related to the learning

 $<sup>^{4}</sup>$ Within the context of singular learning theory, this fact was known at least as early as Fukumizu (1996).

coefficient. In Chen et al. (2023a), a measure called the learning capacity is proposed to estimate the complexity of a hypothesis class. The learning capacity can be viewed as a finite-*n* version of the learning coefficient; the latter only appears in the  $n \to \infty$  limit. Chen et al. (2023a) is largely interested in the learning capacity as a function of training size *n*. They discover the learning capacity saturates at very small and large *n* with a sharp transition in between.

**Applications.** Recently, the LLC estimation method we introduce here has been used to empirically detect "phase transitions" in toy ReLU networks Chen et al. (2023a), and the development of in-context learning in transformers Hoogland et al. (2024).

## F Model complexity vs model-independent complexity

In this paper, we have described the LLC as a measure of "model complexity." It is worth clarifying what we mean here — or rather, what we *do not* mean. This clarification is in part a response to Skalse (2023).

We distinguish measures of "model complexity," such as those traditionally found in statistical learning theory, from measures of "model-independent complexity," such as those found in algorithmic information theory. Measures of model complexity, like the parameter count, describe the expressivity or degrees of freedom available to a *particular model*. Measures of model-independent complexity, like Kolmogorov complexity, describe the complexity inherent to the *task itself*.

In particular, we emphasize that — a priori — the LLC is a measure of model complexity, not model-independent complexity. It can be seen as the amount of information required to nudge a model towards  $w^*$  and away from other parameters. Parameters with higher LLC are more complex for that particular model to implement.

Alternatively, the model is *inductively biased*<sup>5</sup> towards parameters with lower LLC — but a different model could have different inductive biases, and thus different LLC for the same task. This is why it is not sensible to conclude that a bias towards low LLC, would, on its own, explain observed "simplicity bias" in neural networks Valle-Perez et al. (2018) — this is tautological, as Skalse (2023) noted.

To highlight this distinction, we construct a statistical model where these two notions of complexity diverge. Let  $f_1(x)$  be a Kolmogorov-simple function, like the identity function. Let  $f_2(x)$  be a Kolmogorov-complex function, like a random lookup table. Then consider the following regression model with a single parameter  $w \in [0, 1]$ :

$$f(x,w) = w^8 f_1(x) + (1-w^8) f_2(x)$$

For this model,  $f_1(x)$  has a learning coefficient of  $\lambda = \frac{1}{2}$ , whereas  $f_2(x)$  has a learning coefficient of  $\lambda = \frac{1}{16}$ . Therefore, despite  $f_1(x)$  being more Kolmogorov-simple, it is more complex for f(x, w) to implement — the model is biased towards  $f_2(x)$  instead of  $f_1(x)$ , and so  $f_1(x)$  requires relatively more information to learn.

Yet, this example feels contrived: in realistic deep learning settings, the parameters w do not merely interpolate between handpicked possible algorithms, but themselves *define* an internal algorithm based on their values. That is, it seems intuitively like the parameters play a role closer to "source code" than "tuning constants."

Thus, while in general LLC is not a model-independent complexity measure, it seems distinctly possible that for neural networks (perhaps even models in some broader "universality class"), the LLC could be model-independent in some way. This would theoretically establish the inductive biases of neural networks. We believe this to be an intriguing direction for future work.

# G SGLD-based LLC estimator: minibatch version pseudocode

<sup>&</sup>lt;sup>5</sup>The role of the LLC in inductive biases is only rigorously established for Bayesian learning, but we suspect it also applies for learning with SGD.

# Algorithm 1 computing $\hat{\lambda}(w^*)$ Input:

- initialization point:  $w^*$
- scale:  $\gamma$
- step size:  $\epsilon$
- number of iterations: SGLD\_iters
- $\bullet\,$  batch size: m
- dataset of size n:  $\mathcal{D}_n = \{(x_i, y_i)\}_{i=1,\dots,n}$
- averaged log-likelihood function for  $w \in \mathbb{R}^d$  and arbitrary subset D of data:

$$\log \mathcal{L}(D, w) = \frac{1}{|D|} \sum_{(x_i, y_i) \in D} \log p(y_i | x_i, w)$$

**Output**:  $\hat{\lambda}(w^*)$ 1:  $\beta^* \leftarrow \frac{1}{\log n}$  {Optimal sampling temperature.} 2:  $w \leftarrow w^*$  {Initialize at the given parameter} 3: arrayLogL  $\leftarrow$  [] 4: for  $t = 1 \dots \text{SGLD}_{\text{iters}} \text{ do}$ 5: $B \leftarrow$  random minibatch of size mappend  $\log L(B, w)$  to arrayLogL 6:  $\eta \sim N(0, \epsilon)$  {d-dimensional Gaussian, variance  $\epsilon$ } 7:  $\Delta w \leftarrow \frac{\epsilon}{2} \left[ \gamma(w^* - w) + n\beta^* \nabla_w \text{logL}(B, w) \right] + \eta$ 8: 9:  $w \leftarrow w + \Delta w$ 10: **end for** 11:  $\widehat{WBIC} \leftarrow -n \cdot Mean(arrayLogL)$ 12:  $nL_n(w^*) \leftarrow -n \cdot \log L(\mathcal{D}_n, w^*)$ 13:  $\hat{\lambda}(w^*) \leftarrow \frac{\widehat{\text{WBIC}} - nL_n(w^*)}{\log n}$ 14: return  $\hat{\lambda}(w^*)$ 

## H Recommendations for accurate LLC estimation and troubleshooting

In this section, we collect several recommendations for estimating the LLC accurately in practice. Note that these recommendations are largely based on our experience with LLC estimation for DLN (see Section 5.1) as it is the only realistic model (it being wide and deep) where the theoretical learning coefficient is available.

#### H.1 Step size

From experience, the most important hyperparameter to the performance and accuracy of the method is the step size  $\epsilon$ . If the step size is too low, the sampler may not equilibriate, leading to underestimation. If the step size is too high, the sampler can become numerically unstable, causing overestimation or even "blowing up" to NaN values.

Manual tuning of the step size is possible. However we strongly recommend a particular diagnostic based on the acceptance criterion for Metropolis-adjusted Langevin dynamics (MALA). This is used to *correct* numerical errors in traditional MCMC, but here we use it only to *detect* them.

In traditional (full-gradient) MCMC, numerical errors caused by the step size are completely corrected by a secondary step in the algorithm, the *acceptance check* or *Metropolis correction*, which accepts or rejects steps with some probability roughly<sup>6</sup> based on the likelihood of numerical error. The proportion of steps accepted additionally becomes an important diagnostic as to the health of the algorithm: a low acceptance ratio indicates that the acceptance check is having to compensate for high levels of numerical error.

The acceptance probability between step  $X_k$  and proposed step  $X_{k+1}$  is calculated as:

$$\min\left(1, \frac{\pi(X_k) q(X_k | X_{k+1})}{\pi(X_{k+1}) q(X_{k+1} | X_k)}\right)$$

where  $\pi(x)$  is the probability density at x (in our case,  $\log \pi(x) = \beta n L_n(x)$ ), and q(x'|x) is the probability of our sampler transitioning from x to x'.

In the case of MALA,  $q(x'|x) \neq q(x|x')$  and so we must explicitly calculate this term. For MALA, it is:

$$q(x'|x) \propto \exp(-\frac{1}{4\epsilon}||x'-x-\epsilon\nabla\log\pi(x)||^2)$$

We choose to use MALA's formula because we are using SGLD, and both MALA and SGLD propose steps using Langevin dynamics. MALA's formula is the correct one to use when attempting to apply Metropolis correction to Langevin dynamics.

For various reasons, directly implementing such an acceptance check for stochastic-gradient MCMC (while possible) is typically either ineffective or inefficient. Instead we use the acceptance probability merely as a diagnostic.

We recommend tuning the step size such that the average acceptance probability is in the range of 0.9-0.95. Below this range, increase step size to avoid numerical error. Above this range, consider decreasing step size for computational efficiency (to save on the number of steps required). For efficiency, we recommend calculating the acceptance probability for only a fraction of steps — say, one out of every twenty.

Note that since we are not actually using an acceptance check, these acceptance "probabilities" are not really probabilities, but merely diagnostic values.

#### H.2 Step count and burn-in

The step count for sampling should be chosen such that the sampler has time to equilibriate or "burn in". Insufficient step count may lead to underestimating the LLC. Excessive step count will not degrade accuracy, but is unnecessarily time-consuming.

<sup>&</sup>lt;sup>6</sup>Technically, the acceptance probability is based on maintaining detailed balance, not necessarily numerical error, as can be seen in the case of e.g. Metropolis-Hastings. But this is a fine intuition for gradient-based algorithms like MALA or HMC.

We recommend increasing the number of steps until the loss,  $L_m(w_t)$ , stops increasing after some period of time. This can be done with manual inspection of the loss trace. See Figure H.1 for some examples of loss traces and MALA acceptance probabilities over SGLD trajectories for DLN models at different scales.

It is worth noting that the loss trace should truly be flat — a slow upwards slope can still be indicative of significant underestimation.

We also recommend that samples during this burn-in period are discarded. That is, loss values should only be tallied once they have flattened out. This avoids underestimation.

### H.3 Other issues and troubleshooting

We note some miscellaneous other issues and some troubleshooting recommendations:

- Negative LLC estimates: This can happen when  $w^*$  fails to be near a local minimum of L(w). However even when  $w^*$  is a local minimum, we still might get negative LLC estimates if we are not careful. This can happen when the SGLD trajectory wanders to an area with lower loss than the initialization, causing the numerator in (13) to be negative. This could be alleviated by smaller step size or shorter chain length. This, however, risks under-exploration. This can also be alleviated by having larger restoring force  $\gamma$ . This risks the issue discussed below.
- Large  $\gamma$ : An overly concentrated localizing prior ( $\gamma$  too large) can overwhelm the gradient signal coming from the log-likelihood. This can result in samples that are different from the posterior, destroying SGLD's sensitivity to the local geometry.

To sum up, in pathological cases like the SGLD trajectory falling to lower loss regions or blowing up beyond floating point limits, we recommend keeping  $\gamma$  small (1.0 to 10.0), gradually lowering the step-size  $\epsilon$  while lengthening the sampling chain so that the loss trace still equilibrates.



Figure H.1: Sample loss trace (blue, left axis) and MALA acceptance probability (red, right axis) over DLN training trajectories at different model sizes.

## I Learning coefficient of DLN (Aoyagi, 2024)

A DLN is a feedforward neural network without nonlinear activation. Specifically, a biasless DLN with M hidden layers, layer sizes  $H_1, H_2, \ldots, H_M$  and input dimension  $H_0$  is given by:

$$y = f(x, w) = W_M \dots W_2 W_1 x \tag{25}$$

where  $x \in \mathbb{R}^{H_0}$  is the input vector and the model parameter w consist of the weight matrices  $W_j$  of shape  $H_j \times H_{j-1}$  for  $j = 1, \ldots, M$ . Given a DLN, f(x, w), (c.f. Equation 25) with M hidden layers, layer sizes  $H_1, H_2, \ldots, H_M$  and input dimension  $H_0$ , the associated regression model with additive Gaussian noise is given by

$$p(x, y|w) = \frac{q(x)}{\sqrt{2\pi\sigma^2}^{H_M}} e^{-\frac{1}{2\sigma^2} \|y - W_M \dots W_2 W_1 x\|^2}$$
(26)

where q(x) is some distribution on the input  $x \in \mathbb{R}^{H_0}$ ,  $w = (W_1, \ldots, W_M)$  is the parameter consisting of the weight matrices  $W_j$  of shape  $H_j \times H_{j-1}$  for  $j = 1, \ldots, M$  and  $\sigma^2$  is the variance of the additive Gaussian noise. Let q(x, y) be the density of the true data generating process and  $w^* = (W_1^*, \ldots, W_M^*)$  be an optimal parameter that minimizes the KL-divergence between q(x, y) and p(x, y|w).

Here we shall pause and emphasize that this result gives us the (global) learning coefficient, which is conceptually distinct from the LLC. They are related: the learning coefficient is the minimum of the LLCs of the *global* minima of the population loss. In our experiments, we will be measuring the LLC at a randomly chosen global minimum of the population loss. While we don't expect LLCs to differ much among global minima for DLN, we do not know that for certain and it is of independent interest that the estimated LLC can tell us about the learning coefficient.

**Theorem 1** (DLN learning coefficient, Aoyagi, 2024). Let  $r := \operatorname{rank}(W_M^* \dots W_2^* W_1^*)$  be the rank of the linear transformation implemented by the true DLN, f(x, w) and set  $\Delta_j := H_j - r$ , for  $j = 0, \dots, M$ . There exist a subset  $\Sigma \subset \{0, 1, \dots, M\}$  of indices,  $\Sigma = \{\sigma_1, \dots, \sigma_{\ell+1}\}$  with cardinality  $\ell+1$  that satisfy the following conditions:

$$\max\{\Delta_{\sigma} \mid \sigma \in \Sigma\} < \min\{\Delta_{k} \mid k \notin \Sigma\}$$
$$\sum_{\sigma \in \Sigma} \Delta_{\sigma} \ge \ell \cdot \max\{\Delta_{\sigma} \mid \sigma \in \Sigma\}$$
$$\sum_{\sigma \in \Sigma} \Delta_{\sigma} < \ell \cdot \min\{\Delta_{\sigma} \mid \sigma \notin \Sigma\}.$$

Assuming that the DLN truth-model pair (q(x, y), p(x, y|w)) satisfies the relatively finite variance condition (Appendix A.1), their learning coefficient is then given by

$$\lambda = \frac{-r^2 + r(H_0 + H_L)}{2} + \frac{a(\ell - a)}{4\ell} - \frac{\ell(\ell - 1)}{4} \left(\frac{1}{\ell} \sum_{j=1}^{\ell+1} \Delta_{\sigma_j}\right)^2 + \frac{1}{2} \sum_{1 \le i < j \le \ell+1} \Delta_{\sigma_i} \Delta_{\sigma_j}.$$

As we mention in the introduction, trained neural networks are less complex than they seem. It is natural to expect that this, if true, is reflected in deep networks having more degenerate (good parameters has higher volume) loss landscape. With the theorem above, we are given a window into the volume scaling behaviour in the case of DLNs, allowing us to investigate an aspect of this hypothesis.

Figure I.1 shows the true learning coefficient,  $\lambda$ , and the multiplicity, m, of many randomly drawn DLNs with different numbers of hidden layers. Observe that  $\lambda$  decreases with network depth.

This plot is generated by creating networks with 2-800 hidden layers, with width randomly drawn from 100-2000 including the input dimension. The overall rank of the DLN is randomly drawn from the range of zero to the maximum allowed rank, which is the minimum of the layer widths. See also Figure 1 in Aoyagi (2024) for more theoretical examples of this phenomenon.



Figure I.1: The top graph shows  $\lambda$  decreasing as the DLN becomes deeper, even though model parameter count increases with number of layers. The bottom graph shows the true multiplicities, m. Since regular models can only have m = 1, the graph shows that most of these randomly generated DLNs are singular.

# J Experiment: DLN

We compare the estimated  $\hat{\lambda}(w^*)$  against theoretical  $\lambda$  (with  $w^*$  being a randomly generated true parameter), by randomly generating many DLNs with different architectures and model sizes that span several orders of magnitude (OOM). Each DLN is constructed randomly, as follows. Draw an integer  $M \sim U(M_{low}, \ldots, M_{high})$ as the number of hidden layers, where  $U(a, \ldots, b)$  denotes the discrete uniform distribution on the finite set  $\{a, a + 1, \ldots, b\}$ . Then, draw layer size  $H_j \sim U(H_{low}, \ldots, H_{high})$  for each  $j = 0, \ldots, M$  where  $H_0$  denotes the input dimension. The weight matrix  $W_j$  for layer j is then a  $H_j \times H_{j-1}$  matrix with each matrix element independently sampled from N(0, 1) (random initialization). To obtain a more realistic true parameter, with probability 0.5, each matrix  $W_j^*$  is modified to have a random rank of  $r \sim U(0, \ldots, \min(H_{j-1}, H_j))$ . For each DLN generated, a corresponding synthetic training dataset of size n is generated to be used in SGLD sampling.

The configuration values  $M_{low}, M_{high}, H_{low}, H_{high}$  are chosen differently for separate set of experiments with model size targeting DLN size of different order of magnitude. See Table 1 for the values used in the experiments. SGLD hyperparameters  $\epsilon$ ,  $\gamma$ , and number of steps are chosen to suit each set of experiments according to our recommendations outlined in Appendix H. The exact hyperparameter values are given in the following section.

We emphasize that hyperparameters must be tuned independently for different model sizes. In particular, the required step size for numerical stability tends to decrease for larger models, forcing a compensatory increase in the step count. See Table 1 for an example of this tuning with scale.

Future work using e.g.  $\mu$ -parameterization Yang et al. (2022) may be able to alleviate this issue.

# J.1 Hyperparameters and further details

As described in the prior section, the experiments shown in Figure 3 consist of randomly constructed DLNs. For each target order of magnitude of DLN parameter count, we randomly sample the number of layers and their widths from a different range. We also use a different set of SGLD hyperparameters chosen according to the recommendation made in Section H. Configuration values that vary across OOM are shown in Table 1 and other configurations are as follows:

- Batch size used in SGLD is 500.
- The amount of burn-in steps used for SGLD samples is set to 90% of total SGLD chain length, i.e. only the last 10% of SGLD samples are used in estimating the LLC.
- The parameter  $\gamma$  is set to 1.0.
- For each DLN,  $f(x, w^*)$  with a chosen true parameter  $w^*$ , a synthetic dataset,  $\{(x_i, y_i)\}_{i=1,...,n}$  is generated by randomly sampling each element of the input vector x uniformly from the interval [-10, 10] and by setting the output as  $y = f(x, w^*)$ , which effectively means we are setting a very small noise variance  $\sigma^2$ .
- For LLC estimation done at a trained parameter instead of the true parameter (shown in Fig J.3), the network is first trained using SGD with learning rate 0.01 and momentum 0.9 for 50000 steps.

For each target OOM, a number of different experiments are run with different random seeds. The number of such experiments is determined by our compute resources and is reported in Table 1 with some experiment failing due to SGLD chains "blowing up" (See discussion in Appendix H) for the SGLD hyperparameters used. Figure J.4 shows that there is a left tail to the mean MALA acceptance rate distribution that hints at instability in SGLD chains encountered in some  $\hat{\lambda}(w^*)$  estimation runs.

## J.2 Additional plots for DLN experiments

- Figure J.1 is a linear scale version of Figure 3 in the main text. This shows the estimated LLC against the true learning coefficients for experiments at different model size range without log-scale distortion.
- Figure J.2 shows the relative error  $(\lambda \hat{\lambda}(w^*))/\lambda$  across multiple orders of magnitude of DLN model size.
- Figure J.3 shows the results of an additional experiment, where the network is trained with SGD instead of initialized at a global minimum.

OOM	$\begin{vmatrix} \mathbf{Num} & \mathbf{la} \\ \mathbf{ers}, & M_l \\ M_{high} \end{vmatrix}$	$\begin{array}{c c} \mathbf{ay-} & \mathbf{Widths}, \\ & & \\$	$\epsilon$	Num SGLD steps	n	Num experiments
	norght	1	1 7		~	
1k	2-5	5-50	$5 \times 10^{-7}$	10k	$10^{5}$	99
10k	2-10	5-100	$5 \times 10^{-7}$	10k	$10^{5}$	100
100k	2-10	50-500	$1 \times 10^{-7}$	50k	$10^{6}$	100
1M	5-20	100-1000	$5 \times 10^{-8}$	50k	$10^{6}$	99
10M	2-20	500-2000	$2 \times 10^{-8}$	50k	$10^{6}$	93
100M	2-40	500-3000	$2 \times 10^{-8}$	50k	$10^{6}$	54

Table 1: Table of experimental configuration for each batch of experiment at different order of magnitudes (OOM) in DLN model size. n denotes the training dataset size and  $\epsilon$  denotes SGLD step size.



Figure J.1: Supplementary plot to Figure 3. Each plot shows a single batch of DLN experiments with model size at different orders of magnitude. The SGLD hyperparameter is tuned once for each batch. Their values are listed in Table 1 In contrast to Figure 3 which is in log scale, all plots here are in linear scale.



Figure J.2: Relative error of estimated LLC compared to the theoretical learning coefficient, for DLNs across different orders of magnitude of model size.



Figure J.3: Results of LLC estimation when (as is typical in practice) one does not evaluate the LLC at a local minimum of the population loss, but is instead forced to use SGD to first find a local minimum  $\hat{w}_n^*$ . Compare to the results of Figure 3, and note that performance is similar.



Figure J.4: Mean MALA acceptance probability over the entire SGLD trajectory for every DLN experiment. Model size is not the only factor affecting the correct scale for SGLD step size. Local geometry varies significantly among different models, and among different neighbourhoods in the parameter space. Without tuning SGLD hyperparameters individually for each experiment, we get a spread of (mean) MALA acceptance probability over all experiments. Those with low acceptance probability may indicate poor  $\hat{\lambda}(w^*)$  estimation quality.

# K Experiment: LLC for ResNet

This section provides details for the experiments where we aimed to investigate how LLC vary over neural network training with different training configuration. These experiments parallel those performed in Dherin et al. (2022) who propose the *geometric complexity* measure.

We train ResNet18 (He et al., 2016) on the CIFAR10 dataset (Krizhevsky, 2009) using SGD with cross-entropy loss. We vary SGD hyperparameters such as learning rate, batch size, momentum and  $L^2$ -regularization rate and track the resulting LLC estimates over evenly spaced checkpoints of SGD iterations.

For the LLC estimates to be comparable, we need to ensure that the SGLD hyperparameters used in the LLC algorithm is the same. To this end, for every set of experiments where we vary a single optimizer hyperparameter, we first perform a calibration run to select SGLD hyperparameters according to the recommendation out lined in Appendix H. Once selected, this set of SGLD hyperparameter is then used for all LLC estimation within the set of experiments. That include the LLC estimation for every checkpoint of every ResNet18 training run for different optimizer configuration. Also following the same recommendation, we also burn away 90% of the sample trajectory and only use last 10% of samples. We also note that, since the SGLD hyperparameters are not tuned for all experiments within a single set, there is a possibility of negative LLC estimates or divergent SGLD chains. See Appendix H.3 for discussion on such cases and how to troubleshoot them. We manually remove these cases. They are rare enough that they do not change the LLC curves, just widen the error bars (confidence intervals) due to having less repeated experiments.

Each experiment is repeated with 5 different random seeds. While the model architecture and dataset (including the train-test split) is fixed, other factors like the network initialization, training trajectories and SGLD samples are randomized. In each plot, the error bars show the 95% confidence intervals over 5 repeated experiments of the statistics plotted. The error bars were calculated using the inbuilt error bar function to Python Seaborn library (Waskom, 2021) plotting function, seaborn.lineplot(..., errorbar=("ci", 95), ...).

# K.1 Details for main text Figures 1

For experiments that vary the learning rate in Figure 1 (top), for each learning rate value in [0.005, 0.05, 0.01, 0.1, 0.2] we run SGD *without* momentum with a fixed batch size of 512 for 30000 iterations. LLC estimations were performed every 1000 iterations with SGLD hyperparameters as follows: step size  $\epsilon = 2 \times 10^{-7}$ , chain length of 3000 iterations, batch size of 2048 and  $\gamma = 1.0$ .

For experiments that vary the batch size in Figure 1 (middle), for each batch size value in [16, 32, 64, 128, 256, 512, 1024] we run SGD without momentum with a fixed learning rate of 0.01 for 100000 iterations. LLC estimations were performed every 2500 iterations with SGLD hyperparameters as follows: step size  $\epsilon = 2 \times 10^{-7}$ , chain length of 2500 iterations, batch size of 2048 and  $\gamma = 1.0$ .

For experiments that vary the SGD momentum in Figure 1 (bottom), for each momentum value in [0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9] we run SGD with momentum with a fixed learning rate of 0.05 and a fixed batch size of 512 for 20000 iterations. LLC estimations were performed every 1000 iterations with SGLD hyperparameters as follows: step size  $\epsilon = 2 \times 10^{-7}$ , chain length of 3000 iterations, batch size of 2048 and  $\gamma = 1.0$ .

# K.2 Additional ResNet18 + CIFAR10 LLC experiments

**Experiments using SGD** with momentum We repeat the experiments varying learning rate and batch size shown in Figure 1 (top and middle) in the main text, but this time we train ResNet18 on CIFAR10 using SGD with momentum instead. The results are shown in Figure K.1 and the experimental details are as follows:

- For experiments that varies the learning rate (top), for each learning rate value in [0.005, 0.05, 0.01, 0.1, 0.2] we run SGD with momentum of 0.9 with a fixed batch size of 512 for 30000 iterations. LLC estimations were performed every 1000 iterations with SGLD hyperparameters as follows: step size  $\epsilon = 2 \times 10^{-7}$ , chain length of 3000 iterations, batch size of 2048 and  $\gamma = 1.0$ .
- For experiments that varies the batch size (bottom), for each batch size value in [16, 32, 64, 128, 256, 512, 1024] we run SGD with momentum of 0.9 with a fixed learning rate of 0.01 for 100000 iterations. LLC estimations

were performed every 2500 iterations with SGLD hyperparameters as follows: step size  $\epsilon = 2 \times 10^{-7}$ , chain length of 2500 iterations, batch size of 2048 and  $\gamma = 1.0$ .

**Explicit**  $L^2$ -regularization We also run analogous experiments using explicit  $L^2$ -regularization. We trained ResNet18 on CIFAR10 dataset using SGD both with and without momentum using the usual cross-entropy loss but with an added  $L^2$ -regularization term,  $\alpha ||w||_2^2$  where w denotes the network weight vector and  $\alpha$  is the regularization rate hyperparameter that we vary in this set of experiments. Similar to other experiments in this section, we track LLC estimates over evenly-spaced training checkpoints. However, it is noting that the loss function used for SGLD sampling as part of the LLC estimation is the original cross-entropy loss function without the added  $L^2$ -regularization term.

The results are shown in Figure K.2 and the details are as follows:

- For the experiment with momentum (Figure K.2 top), for each regularization rate  $\alpha \in [0.0, 0.01, 0.025, 0.05, 0.075, 0.1]$ , we run SGD with momentum of 0.9 with a fixed learning rate of 0.0005 and batch size of 512 for 15000 iterations. LLC estimations were performed every 500 iterations with SGLD hyperparameters as follows: step size  $\epsilon = 5 \times 10^{-8}$ , chain length of 2000 iterations, batch size of 2048 and  $\gamma = 1.0$ .
- For the experiment without momentum (Figure K.2 bottom), for each regularization rate  $\alpha \in [0.0, 0.01, 0.025, 0.05, 0.075, 0.1]$ , we run SGD without momentum with a fixed learning rate of 0.001 and batch size of 512 for 50000 iterations. LLC estimations were performed every 1000 iterations with SGLD hyperparameters as follows: step size  $\epsilon = 5 \times 10^{-8}$ , chain length of 2000 iterations, batch size of 2048 and  $\gamma = 1.0$ .



Figure K.1: Impact of varying different training configuration believed to exert implicit regularization pressure when training ResNet18 on CIFAR10 data using SGD *with momentum* (contrast with those without momentum reported in Figure 1 in the main text). Top: varying learning rate. Bottom: varying batch size.



Figure K.2: Impact of varying explicit  $L^2$ -regularization rate when training ResNet18 on CIFAR10 data using SGD with (top) and without (bottom) momentum.

# L Additional Experiment: LLC for language model

The purpose of this experiment is simply to verify that the LLC can be consistently estimated for a transformer trained on language data. In Figure L.1, we show the LLC estimates for  $\hat{w}_n^*$  at the end of training over a few training runs. We see the LLC estimates are a small fraction of the 3.3m total parameters. We also notice that the value of the LLC estimates are remarkably stable over multiple training runs. See Appendix L for experimental details.



Figure L.1: SGLD-based LLC estimates for  $\hat{w}_n^*$  at the end of training an attention-only transformer on subset of the Pile dataset. The distribution reports LLC estimates over 10 training repetitions. Again the LLC is a tiny fraction of the 3.3m parameters in the transformer.

We trained a two-layer attention-only (no MLP layers) transformer architecture on a resampled subset of the Pile dataset Gao et al. (2020); Xie et al. (2023) with a context length of 1024, a residual stream dimension of  $d_{\text{model}} = 256$ , and 8 attention heads per layer. The architecture also uses a learnable Shortformer embedding Press et al. (2021) and includes layer norm layers. Additionally, we truncated the full GPT-2 tokenizer, which has a vocabulary of around 50,000 tokens, down to the first 5,000 tokens in the vocabulary to reduce the size of the model. The resulting model has a total parameter count of d = 3,355,016. We instantiated these models using an implementation provided by TransformerLens Nanda and Bloom (2022).

We trained 10 different seeds over a single epoch of 50,000 steps with a minibatch size of 100, resulting in about 5 billion tokens used during training for each model. We used the AdamW optimizer with a weight decay value of 0.05 and a learning rate of 0.001 with no scheduler.

We run SGLD-based LLC estimation once at the end of training for each seed at a temperature of  $\beta = 1/\log(100)$ . We set  $\gamma = 100$  and  $\epsilon = 0.001$ . We take samples over 20 SGLD chains with 200 draws per chain using a validation set.

# M Additional Experiment: MALA versus SGLD

Here we verify empirically that our SGLD-based LLC estimator (Algorithm 1) does not suffer from using minibatch loss for both SGLD sampling and LLC calculation. Specifically, we compare to LLC estimation via Equation 13 and the Metropolis-adjusted Langevin Algorithm (MALA), a standard gradient-based MCMC algorithm Roberts and Rosenthal (1998). Notice that this comparison rather stacks the odds against the minibatch-version SGLD-based LLC estimator in Algorithm 1 so that it is all the more surprising we see such good results below.

We test with a two-hidden-layer ReLU network with ten inputs, ten outputs, and twenty neurons per hidden layer. Denote the inputs by x, the parameters by w, and the output of this network f(x, w). The data are generated to create a "realizable" data generating process, with "true parameter"  $w^*$ : inputs X are generated from a uniform distribution, and labels Y are (noiselessly) generated based on the true network, so that  $Y_i = f(X_i, w^*)$ .

Hyperparameters and experimental details are as follows. We sweep dataset size from 100 to 100000, and compare our LLC estimator and the MALA-based one. For all dataset sizes, SGLD batch size was set to 32 and  $\gamma = 1.0$ , and MALA and SGLD shared the same true parameter  $w^*$  (set at random according to a normal distribution). Both MALA and SGLD used a step size of 1e-5 and the asymptotically optimal inverse temperature  $\beta^* = 1/\log n$ . Experiments were run on CPU.



Figure M.1: We compare LLC estimation using SGLD with LLC estimation using MALA. They make similar estimates (top) but the SGLD-based method is significantly faster (bottom), especially for large dataset sizes.

The results are summarized in Figure M.1. We find that across all dataset sizes, the SGLD and MALA estimates of the LLC agree (Figure M.1, top), but the SGLD-based estimate has far lower computational cost, especially as the dataset size grows (Figure M.1, bottom).

# N Additional Experiment: SGD versus eSGD

We fit a two hidden-layer feedforward ReLU network with 1.9m parameters to MNIST using two stochastic optimizers: SGD and entropy-SGD (Chaudhari et al., 2019). We choose entropy-SGD because its objective is to minimize  $F_n(w^*, \gamma)$  over  $w^*$ , so we expect that the local minima found by entropy-SGD will have lower  $\hat{\lambda}(w^*)$ .

Figure N.1 shows the LLC estimates  $\hat{\lambda}(\hat{w}_n^*)$  for  $\hat{w}_n^*$  at the end of training, optimized by either entropy-SGD or standard SGD. Notice the LLC estimates, for both stochastic optimizers, are on the order of 1000, much lower than the 1.9m parameters in the ReLU network.



Figure N.1: LLC estimates for  $\hat{w}_n^*$  at the end of training a feedforward ReLU network on MNIST. The distribution reports  $\hat{\lambda}(\hat{w}_n^*)$  over 80 training repetitions where the training data remains fixed in this repetition, only the randomness in the stochastic optimizer is being modded out. We compare two stochastic optimizers – SGD and entropy-SGD. Note all  $\hat{\lambda}(\hat{w}_n^*)$  are on the order of 1000 while the parameter count in the ReLU network is 1.9**m**.

Figure N.1 confirms our expectation that entropy-SGD finds local minima with lower LLC, i.e., entropy-SGD is attracted to more degenerate (simpler) critical points than SGD. Interestingly Figure N.1 also reveals that the LLC estimate has remarkably low variance over the randomness of the stochastic optimizer. Finally it is noteworthy that the LLC estimate of a learned NN model for both stochastic optimizers is, on average, a tiny percentage of the total number of weights in the NN model:  $\hat{\lambda}(w^*) \approx 1000$ .

In this experiment, we trained a feedforward ReLU network on the MNIST dataset Deng (2012). The dataset consists of 60000 training samples and 10000 testing samples. The network is designed with 2 hidden layers having sizes [1024, 1024], and it contains a total of 1863690 parameters.

For training, we employed two different optimizers, SGD and entropy-SGD, minimizing cross-entropy loss. Both optimizers are set to have learning rate of 0.01, momentum parameter at 0.9 and batch size of 512. SGD is trained with Nesterov look-ahead gradient estimator. The number of samples L used by entropy-SGD for local free energy estimation is set to 5. The network is trained for 200 epochs. The number of epochs is chosen so that the classification error rate on the training set falls below  $10^{-4}$ .

The hyperparameters used for SGLD are as follows:  $\epsilon$  is set to  $10^{-5}$ , chain length to 400 and the minibatch size 512, and  $\gamma = 100$ . We repeat each SGLD chain 4 times to compute the variance of estimated quantities and also as a diagnostic tool, a proxy for estimation stability.

## **O** Additional Experiment: scaling invariance

In Appendix C, we show theoretically that the theoretical LLC is invariant to local diffeomorphism. Here we empirically verify that our LLC *estimator* (with preconditioning) is capable of satisfying this property in a specific easy-to-test case (though we do not test it in general).

The function implemented by a feedforward ReLU networks is invariant to a type of reparameterization known as *rescaling symmetry*. Invariance of other measures to this symmetry is not trivial or automatic, and other geometric measures like Hessian-based basin broadness have been undermined by their failure to stay invariant to these symmetries Dinh et al. (2017).

For simplicity, suppose we have a two-layer ReLU network, with weights  $W_1, W_2$  and biases  $b_1, b_2$ . Then rescaling symmetry is captured by the following fact, for some arbitrary scalar  $\alpha$ :

$$W_2 \operatorname{ReLU}(W_1 x + b_1) + b_2 = \alpha W_2 \operatorname{ReLU}\left(\frac{1}{\alpha} W_1 x + \frac{1}{\alpha} b_1\right) + b_2$$

That is, we may choose new parameters  $W'_1 = \frac{1}{\alpha}W_1, b'_1 = \frac{1}{\alpha}b_1, W'_2 = \alpha W_2, b'_2 = b_2$  without affecting the inputoutput behavior of the network in any way. This symmetry generalizes to any two adjacent layers in ReLU networks of arbitrary depth.

Given that these symmetries do not affect the function implemented by the network, and are present globally throughout all of parameter space, it seems like these degrees of freedom are "superfluous", and should ideally not affect our tools. Importantly, this is the case for the LLC.

We verify empirically that this property also appears to hold for our LLC estimator, when proper preconditioning is used.

#### 0.1 Preconditioned SGLD

We must slightly modify our SGLD sampler to perform this experiment tractably. This is because applying rescaling symmetry makes the loss landscape significantly anisotropic, forcing prohibitively small step sizes with the original algorithm.

Thus we add preconditioning to the SGLD sampler from Section 4.5, with the only modification being a fixed preconditioning matrix A:

$$\Delta w_t = A \frac{\epsilon}{2} \left( \frac{\beta^* n}{m} \sum_{(x,y) \in B_t} \nabla \log p(y|x, w_t) + \gamma(w^* - w_t) \right) + N(0, \epsilon)$$
(27)

In the experiment to follow, this preconditioning matrix is hardcoded for convenience, but if this algorithm is to be used in practice, the preconditioning matrix must be learned adaptively using standard methods for adaptive preconditioning (Haario et al., 1999).

#### **O.2** Experiment

We take a small feedforward ReLU network, and rescale two adjacent layers in the network by  $\alpha$  in the fashion described above. We vary the value of  $\alpha$  across eight orders of magnitude and measure the estimated LLC using Eq 27 for SGLD sampling and Eq 13 for calculating LLC from samples.<sup>7</sup>

Crucially, we must use preconditioning here so as to avoid prohibitively small step size requirements. In this case, the preconditioning matrix A is set manually, to be a diagonal matrix with entries  $\alpha^2$  for parameters corresponding to  $W_1$  and  $b_1$ , entries  $\frac{1}{\alpha^2}$  for  $W_2$ , and entries 1 otherwise.<sup>8</sup>

<sup>&</sup>lt;sup>7</sup>Note that this means that we are not using Algorithm 1 here, both because of preconditioned SGLD and because we are calculating LLC using full-batch loss instead of mini-batch loss.

<sup>&</sup>lt;sup>8</sup>In practical situations, the preconditioning matrix cannot be set manually, and must be learned adaptively. Standard methods for adaptive preconditioning exist in the MCMC literature Haario et al. (2001, 1999).



Figure O.1: LLC estimation is invariant to rescaling symmetries in ReLU networks. As the rescaling parameter  $\alpha$  is varied over eight orders of magnitude, the estimated value of the LLC remains invariant (up to statistical error). The small error bar across multiple SGLD runs illustrates the stability of the estimation method. Model layer sizes, including input dimension is shown in the legend.

The results can be found in Figure O.1. We conclude that LLC estimation appears invariant to ReLU network rescaling symmetries.

# P Compute resources disclosure

Specific details about each type of experiment we carried out are listed below. Additional compute resources were required for finding suitable SGLD hyperparameters for LLC estimation, but they did not significantly impact the overall resource requirement.

**ResNet18 + CIFAR10 experiments.** Each experiment, i.e. training a ResNet18 for the reported number of iteration and performing LLC estimation for the stated number of checkpoints, is run on a node with either a single V100 or A100 NVIDIA GPU depending on availability hosted on internal HPC cluster with 2 CPU cores and 16GB memory allocated. No significant storage required. Each experiment took between a few minutes to 3 hours depending on the configuration (mean: 87.6 minutes, median: 76.7 minutes).

Estimated total compute is 287 GPU hours spread across 208 experiments.

**DLN experiments.** Each experiment, either estimating LLC at the true parameter or at a trained SGD parameter (thus require training), is run on a single A100 NVIDIA GPU node hosted on internal HPC cluster with 2 CPU cores and no more than 8GB memory allocated. No significant storage required. Each experiment took less than 15 minutes.

Estimated total compute is 150 GPU hours: 600 experiments (including failed ones) each around 15 GPU minutes.

**MNIST experiments.** Each repetition of training a feedforward ReLU network using SGD or eSGD optimizer on MNIST data and estimating the LLC at the end of training is run on a single A100 NVIDIA GPU node hosted on internal HPC cluster with 8 CPU cores and no more than 16GB memory allocated. No significant storage required. Each experiment took less than 10 minutes.

Estimated total compute is 27 GPU hours: 2 sets of 80 repetitions, 10 GPU minutes each.

Language model experiments. Training the language model and estimating its LLC at the last checkpoint took around 30 minutes on Google Colab on a single A100 NVIDIA GPU with 84GB memory and 1 CPU allocated. The storage space used for the training data is around 27GB.

Estimated total compute is 5 GPU hours: 10 repetitions of 30 GPU minutes each.