
Uncertainty Quantification From Scaling Laws in Deep Neural Networks

Ibrahim Elsharkawy
Department of Physics
University of Illinois Urbana-Champaign
Urbana, IL 61801
ie4@illinois.edu

Benjamin Hooberman
Department of Physics
University of Illinois Urbana-Champaign
Urbana, IL 61801
benhoob@illinois.edu

Yonatan Kahn
Department of Physics and Center for Artificial Intelligence Innovation
University of Illinois Urbana-Champaign
Urbana, IL 61801
yfkahn@illinois.edu

Abstract

Quantifying the uncertainty from machine learning analyses is critical to their use in the physical sciences. In this work we focus on uncertainty inherited from the initialization distribution of neural networks. We compute the mean $\mu_{\mathcal{L}}$ and variance $\sigma_{\mathcal{L}}^2$ of the test loss \mathcal{L} for an ensemble of multi-layer perceptrons (MLPs) with neural tangent kernel (NTK) initialization in the infinite-width limit, and compare empirically to the results from finite-width networks for three example tasks: MNIST classification, CIFAR classification and calorimeter energy regression. We observe scaling laws as a function of training set size $N_{\mathcal{D}}$ for both $\mu_{\mathcal{L}}$ and $\sigma_{\mathcal{L}}$, but find that the coefficient of variation $\epsilon_{\mathcal{L}} \equiv \sigma_{\mathcal{L}}/\mu_{\mathcal{L}}$ becomes independent of $N_{\mathcal{D}}$ at both infinite and finite width for sufficiently large $N_{\mathcal{D}}$. This implies that the coefficient of variation of a finite-width network may be approximated by its infinite-width value, and may be calculable using finite-width perturbation theory.

1 Introduction

Deep learning techniques have improved performance beyond conventional methods in a wide variety of tasks. However, for neural networks in particular, it is not straightforward to assign network-induced uncertainty on their output as a function of network architecture, training algorithm, and initialization (1). One approach to uncertainty quantification (UQ) is to treat any individual network as a draw from an ensemble, and identify the systematic uncertainty with the variance in the neural network outputs over the ensemble (2; 3). This variance can certainly be measured empirically by training a large ensemble of networks, but it would be advantageous to be able to predict it from first principles. This is possible in the infinite-width limit of multi-layer perceptron (MLP) architectures, where the statistics of the network outputs after training are Gaussian with mean and variance determined by the neural tangent kernel (NTK) (4; 5; 6). For realistic MLPs with large but finite width n , one can compute corrections to this Gaussian distribution that are perturbative in $1/n$ (7). However, the finite-width computation depends on tensors of size $N_{\mathcal{D}}^4$, where $N_{\mathcal{D}}$ is the training set size, which rapidly becomes unwieldy as $N_{\mathcal{D}}$ grows large (1; 2; 3).

To sidestep the difficulties of large $N_{\mathcal{D}}$, we propose to exploit scaling laws (8; 9; 10), an apparently ubiquitous phenomenon in deep learning where test loss \mathcal{L} follows a power law $\mathcal{L} \propto N_{\mathcal{D}}^{-\alpha_{\mathcal{D}}}$ with some task-dependent scaling exponent $\alpha_{\mathcal{D}}$. Such scaling laws have been observed for both finite-

width and infinite-width networks, and have been argued to be relatively insensitive to the neural network architecture for a given task, notably including whether or not a network is at finite width. Without attempting to explain the origin of these scaling laws, we can nonetheless use their existence and the simplicity of the infinite-width Gaussian statistics to attempt to extrapolate the expected mean test loss $\mu_{\mathcal{L}}$ and its variance $\sigma_{\mathcal{L}}^2$ to large $N_{\mathcal{D}}$. The infinite-width prediction only depends on matrices of size $N_{\mathcal{D}} \times N_{\mathcal{D}}$, making a comparison to experiments more feasible.

In this work, we compute $\mu_{\mathcal{L}}$ and $\sigma_{\mathcal{L}}^2$ for infinite-width MLPs for regression tasks on three example datasets: the benchmark MNIST classification problem formulated as a regression to 1-hot labels, the benchmark CIFAR classification problem also formulated as a regression to 1-hot labels, and an example from high-energy physics (HEP), energy calibration in a calorimeter detector. For all examples, we find nontrivial scaling laws for both $\mu_{\mathcal{L}}$ and $\sigma_{\mathcal{L}}^2$, but with related scaling exponents such that the coefficient of variation

$$\epsilon_{\mathcal{L}} \equiv \frac{\sigma_{\mathcal{L}}}{\mu_{\mathcal{L}}} \quad (1)$$

asymptotes to a scaling exponent of $\alpha_D \approx 0$ at sufficiently large $N_{\mathcal{D}}$. We give a plausibility argument for this somewhat surprising ‘‘invariant’’ of infinite-width architectures, and we demonstrate with numerical experiments that very similar scaling exponents persist at finite width.

There is a large body of work investigating scaling laws for wide networks and linear models, including Refs. (11; 12; 13; 14; 9; 15; 16; 10; 17). To our knowledge, our work is the first to demonstrate the scaling exponent $\alpha_D \approx 0$ for the coefficient of variation statistic $\epsilon_{\mathcal{L}}$, in the NTK parameterization where finite-width corrections may in principle be computed perturbatively. In future work we intend to study how our results interplay with related work in mean-field or μ -parameterization (14), as well as to what extent feature learning is important for the effects of finite width on the coefficient of variation.

2 Problem setup

In order to maintain theoretical control and confine all stochasticity to the initialization distribution, we work with MLPs in critically initialized NTK parametrization, trained with full-batch gradient descent (GD). We have shown in A.3 that our results generalize to other optimizing schemes, such as ADAM. The MLP forward pass is given by

$$z_i^{(1)} = b_i^{(1)} + \sum_{j=1}^{n_0} W_{ij}^{(1)} x_j, \quad z_i^{(\ell+1)} = b_i^{(\ell+1)} + \sum_{j=1}^{n_{\ell}} W_{ij}^{(\ell+1)} \sigma(z_j^{(\ell)}), \quad (2)$$

where x_j are the components of an input vector $\vec{x} \in \mathbb{R}^{n_0}$, $W_{ij}^{(\ell)}$ and $b_i^{(\ell)}$ are weights and biases at layer ℓ , $\sigma(z)$ is the activation function, all $L-1$ hidden layers have the same width n , and the network output is a vector $\vec{z}^{(L)} \in \mathbb{R}^{n_L}$. To facilitate analytic computations of the NTK (see Sec. 3 below), we choose $\sigma(z) = \text{erf}(z)$. For this choice of activation, initializing biases to zero and drawing weights in layer ℓ from zero-mean Gaussian distributions with variance $1/n_{\ell-1}$ enforces criticality, namely that the typical magnitude of preactivations does not grow or shrink exponentially with depth (7). Each network is then trained with full-batch gradient descent,

$$\theta_{\mu} \leftarrow \theta_{\mu} - \eta \sum_{\nu} \lambda_{\mu\nu} \frac{d\mathcal{L}_{\mathcal{A}}}{d\theta_{\nu}}, \quad \mathcal{L}_{\mathcal{A}} \equiv \frac{1}{n_L N_{\mathcal{D}}} \sum_{\alpha \in \mathcal{A}} \frac{1}{2} \|\vec{z}_{\alpha}^{(L)} - \vec{y}_{\alpha}\|^2, \quad (3)$$

where $\mathcal{L}_{\mathcal{A}}$ is the MSE loss on the training set $\mathcal{A} = \{\vec{x}_{\alpha}, \vec{y}_{\alpha}\}$, θ_{μ} indexes the weights and biases, η is the global learning rate, and (following (7)) the learning rate tensor $\lambda_{\mu\nu}$ is

$$\lambda_{b_{i_1}^{(\ell)} b_{i_2}^{(\ell)}} = \delta_{i_1 i_2} \lambda_b / \ell, \quad \lambda_{W_{i_1 j_1}^{(\ell)} W_{i_2 j_2}^{(\ell)}} = \delta_{i_1 i_2} \delta_{j_1 j_2} \lambda_W / n. \quad (4)$$

All of the above scalings permit a sensible infinite-width limit, $n \rightarrow \infty$. Our free hyperparameters at infinite width are therefore η , L , and λ_b / λ_W . We specify particular choices for these hyperparameters in Sec. 4, but have checked that our results do not appear to depend on these choices. We are interested in $\mu_{\mathcal{L}} \equiv \mathbb{E}[\mathcal{L}_{\mathcal{B}}]$ and $\sigma_{\mathcal{L}}^2 \equiv \text{Var}[\mathcal{L}_{\mathcal{B}}]$ as a function of $N_{\mathcal{D}}$, where \mathcal{B} is the test set and the expectation and variance are taken over the initialization distribution.

3 Infinite-width predictions

The predictions and statistics of an infinite-width MLP of depth L are fully characterized by two matrices: the kernel $K_{\alpha\beta}^{(L)}$ and the NTK $\Theta_{\alpha\beta}^{(L)}$, where α and β are inputs from the training or test sets. In the notation of Ref. (7), the components of these matrices can be defined through forward recursions

$$K_{\alpha\beta}^{(1)} = \frac{1}{n_0} \vec{x}_\alpha \cdot \vec{x}_\beta, \quad K_{\alpha\beta}^{(\ell+1)} = \langle \sigma(u_\alpha) \sigma(u_\beta) \rangle_{K^{(\ell)}}; \quad (5)$$

$$\Theta_{\alpha\beta}^{(1)} = \lambda_b + \frac{\lambda_W}{n_0} \vec{x}_\alpha \cdot \vec{x}_\beta; \quad \Theta_{\alpha\beta}^{(\ell+1)} = \frac{\lambda_b}{\ell} + \lambda_W \langle \sigma(u_\alpha) \sigma(u_\beta) \rangle_{K^{(\ell)}} + \langle \sigma'(u_\alpha) \sigma'(u_\beta) \rangle_{K^{(\ell)}} \Theta_{\alpha\beta}^{(\ell)}, \quad (6)$$

where

$$\langle F(u_{\alpha_1}, u_{\alpha_2}) \rangle_K \equiv \frac{1}{\sqrt{\det(2\pi \mathbf{K}_2)}} \int du_{\alpha_1} du_{\alpha_2} \exp\left(-\frac{1}{2} \mathbf{u}^T \mathbf{K}_2 \mathbf{u}\right) F(u_{\alpha_1}, u_{\alpha_2}) \quad (7)$$

is a two-dimensional Gaussian expectation over $\mathbf{u} \equiv (u_{\alpha_1}, u_{\alpha_2})$ with covariance $\mathbf{K}_2 \equiv K_{\alpha_1\alpha_2}$, the 2×2 submatrix of K . We choose $\sigma(z) = \text{erf}(z)$ in this work in order to analytically evaluate all Gaussian integrals in (5) and (6), as derived in Ref. (18; 19; 6), which minimizes numerical error.

For notational convenience, we now specialize to the case of one output neuron i.e. $n_L = 1$, though all of the steps below generalize to any n_L . At initialization, the output distribution $p(z^{(L)}|\mathcal{A})$ is Gaussian with mean zero and covariance $K^{(L)}$. Each subsequent GD step is a linear transformation,

$$z_\delta^{(L)} \leftarrow z_\delta^{(L)} - \eta \sum_{\alpha \in \mathcal{A}} \Theta_{\delta\alpha}^{(L)} (z_\alpha^{(L)} - y_\alpha), \quad (8)$$

which updates the mean and covariance of $p(z^{(L)}|\mathcal{A})$ but preserves the Gaussian nature of the distribution. If the NTK submatrix evaluated on the training set, $\Theta_{\mathcal{A}}$, is invertible, taking the limit of an infinite number of GD steps leads to a convergent geometric series. In this case, the end-of-training mean prediction m_β^∞ and variance $(\sigma_\beta^\infty)^2$ for a single test point \vec{x}_β is independent of η and is given by (6; 7):

$$m_\beta^\infty \equiv \mathbb{E}[z_\beta^{(L)}] = \Theta_\beta^T \Theta_{\mathcal{A}}^{-1} \mathbf{y}_{\mathcal{A}}, \quad (\sigma_\beta^\infty)^2 \equiv \text{Var}[z_\beta^{(L)}] = K_{\beta\beta}^{(L)} - 2\Theta_\beta^T \Theta_{\mathcal{A}}^{-1} \mathbf{K}_\beta + \Theta_\beta^T \Theta_{\mathcal{A}}^{-1} \mathbf{K}_{\mathcal{A}} \Theta_{\mathcal{A}}^{-1} \Theta_\beta, \quad (9)$$

where $\mathbf{y}_{\mathcal{A}}$ is the vector of training set labels; $\mathbf{K}_{\mathcal{A}}$ is the kernel evaluated on the training set; and Θ_β and \mathbf{K}_β are columns of $\Theta_{\alpha\beta}^{(L)}$ and $K_{\alpha\beta}^{(L)}$, respectively. If $\Theta_{\mathcal{A}}$ is ill-conditioned, we may still compute the mean and variance of the prediction at each step of GD using Eq. (8), and implement early stopping to determine m_β^∞ and σ_β^∞ . In either case, our desired quantities $\mu_{\mathcal{L}}$ and $\sigma_{\mathcal{L}}^2$ are

$$\mu_{\mathcal{L}} \equiv \frac{1}{2|\mathcal{B}|} \sum_{\beta \in \mathcal{B}} \mathbb{E}[(z_\beta^{(L)} - y_\beta)^2], \quad \sigma_{\mathcal{L}}^2 \equiv \left(\frac{1}{4|\mathcal{B}|^2} \sum_{\beta_1, \beta_2 \in \mathcal{B}} \mathbb{E}[(z_{\beta_1}^{(L)} - y_{\beta_1})^2 (z_{\beta_2}^{(L)} - y_{\beta_2})^2] \right) - \mu_{\mathcal{L}}^2. \quad (10)$$

The Gaussian expectations are evaluated in Appendix A.1 and generalized for $n_L \neq 1$.

4 Finite-width Experiments

We compare the infinite-width predictions of Sec. 3 with experiments on finite-width networks. We use three example datasets: the standard MNIST benchmark (20), the standard CIFAR benchmark (21), and an example relevant for HEP, namely simulated calorimeter data generated for a prior study of electron and photon classification and energy measurement using regression (22). A common application of machine learning in HEP is to determine the energy of a high-energy particle based on the pattern of energy deposits (“hits”) in a segmented calorimeter detector. Because the spatial distribution of hits is important, an architecture encoding spatial correlations such as a convolutional neural network (CNN) is likely preferable to an MLP, but a persistent issue is quantifying uncertainties in the inferred energy from these analyses; currently, UQ is done by training large ensembles of

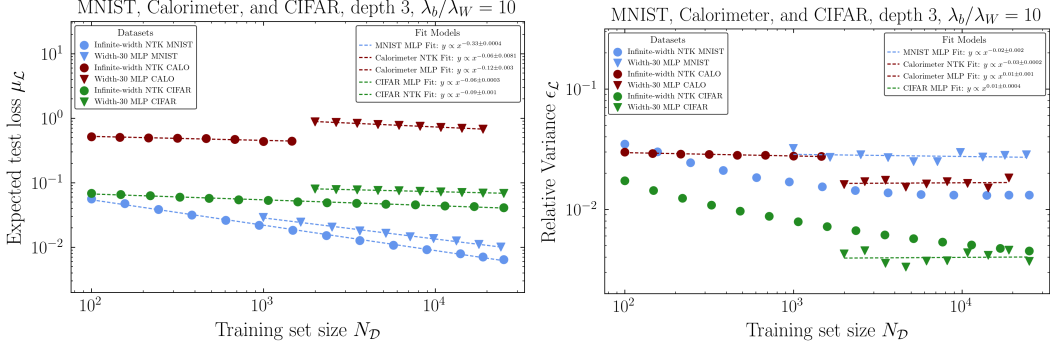


Figure 1: Mean loss (left) and coefficient of variation (right) scaling laws on our three example datasets. The fits show linear least-squares fits in log-log space along with 1σ errors on the power law index. Infinite-width MNIST and CIFAR appear to show a broken power law for $\epsilon_{\mathcal{L}}$ so we do not provide a fit.

neural networks, which is resource-intensive (22). To apply our MLP setup to this problem, we flatten the raw simulated calorimeter hits into a single vector (with the expectation of dramatic performance reduction) and regress to the true energy of the particle initiating the event. Details of the data generation procedure, as well as a link to our training set data, are given in Appendix A.2. In the calorimeter example, $\mu_{\mathcal{L}} \neq 0$ represents the magnitude of the bias away from the true energy of each event. If $\epsilon_{\mathcal{L}} \ll 1$, then $\mu_{\mathcal{L}}$ should be assigned as the systematic uncertainty on the energy calibration. However, if $\epsilon_{\mathcal{L}} \gtrsim 1$, the fluctuations of the predictions of the MLP ensemble are larger than the bias, and the systematic uncertainty is instead dominated by $\sigma_{\mathcal{L}}$.

4.1 Training and architecture details

In our experiments, we choose a random sample of size $N_{\mathcal{D}}$ from each of our example datasets as our training set, and train an ensemble 150 MLPs for each $N_{\mathcal{D}}$. Each MLP has $L = 3$ layers with a width of $n = 30$ for each hidden layer. The networks are initialized and trained as described in Sec. 2. Training is executed using TensorFlow (23), leveraging the NSF ACCESS program under an Explore allocation of ~ 4000 node-hours. We set $\lambda_b/\lambda_W = 10$ and use $\eta = 10$ for MNIST, $\eta = 5$ for CIFAR, and $\eta = 10^{-2}$ for the calorimeter dataset. For the calorimeter data at infinite width, we take $\eta = 10^{-4}$. Early stopping with a validation set of 10^3 and patience of 10^4 epochs is used to monitor for overfitting; each network trains in roughly 1 hour in $\approx 10^5 - 10^6$ epochs. Finally, the performance of each network is evaluated on a fixed test set of size 10^4 , which is consistent across all training set sizes. The resulting test loss of each network at each $N_{\mathcal{D}}$ is saved to compute the sample mean and variance offline.

4.2 Results

The results of our experiments are shown in Fig. 1. We see that the scaling of mean loss (left panel) and $\epsilon_{\mathcal{L}}$ (right panel) are quite similar between infinite and finite width. For our problems, $\mu_{\mathcal{L}}$ does not differ by much for the two architectures, and similarly, $\epsilon_{\mathcal{L}}$ is on the same order of magnitude. Further, $\alpha_{\mathcal{D}} \approx 0$ for $\epsilon_{\mathcal{L}}$ for infinite width and finite width for both datasets (or asymptotes to zero at large $N_{\mathcal{D}}$ for infinite-width MNIST and CIFAR), highlighting $\epsilon_{\mathcal{L}}$ as a potential MLP “invariant” meriting future study. We see similar results for the Adam optimizer (24) (Appendix A.3).

One can make the following heuristic plausibility argument for the infinite-width scaling of $\epsilon_{\mathcal{L}}$. Suppose the entries of the test-train NTK $\Theta_{\mathcal{B}}$, kernel \mathbf{K} , and inverse NTK $\Theta_{\mathcal{A}}^{-1}$ scale with $N_{\mathcal{D}}$ as $\Theta_{\mathcal{B}} \sim (N_{\mathcal{D}})^p$, $\mathbf{K} \sim (N_{\mathcal{D}})^r$, $\Theta_{\mathcal{A}}^{-1} \sim (N_{\mathcal{D}})^k$. Empirically measuring each of these scaling laws (Appendix A.4), we find that $p, r \approx 0$, but $k < 0$; the $N_{\mathcal{D}}$ scaling of the inverse NTK elements seems to be driving the $\mu_{\mathcal{L}}$ scaling law. Using Eq. 9 and assuming the last term in $(\sigma_{\beta}^{\infty})^2$ dominates, we find $\mu_{\mathcal{L}} \sim (N_{\mathcal{D}})^{2k+4}$ and $\sigma_{\mathcal{L}}^2 \sim (N_{\mathcal{D}})^{4k+8}$, where the extra powers of $N_{\mathcal{D}}$ come from the matrix multiplications over the training set. Since $\sigma_{\mathcal{L}}^2$ contains exactly twice as many powers of $\Theta_{\mathcal{A}}^{-1}$ and sums over \mathcal{A} as $\mu_{\mathcal{L}}$, the coefficient of variation scales as $\epsilon_{\mathcal{L}} \sim \frac{\sqrt{(N_{\mathcal{D}}^{4k+8})}}{(N_{\mathcal{D}})^{2k+4}} \sim (N_{\mathcal{D}})^0$. The broken power law in $\epsilon_{\mathcal{L}}$ for MNIST and CIFAR may result from the different scaling of other terms in the

variance in Eq. (9) for small $N_{\mathcal{D}}$. To connect with existing UQ approaches in physics using Bayesian networks (25; 26), we investigate a Bayesian interpretation of the infinite-width $\epsilon_{\mathcal{L}}$ in Appendix A.5.

5 Conclusion

In this work we have shown empirically that the mean and variance of the MLP test loss on three regression tasks exhibit scaling laws with training set size $N_{\mathcal{D}}$. The scaling exponents for the mean loss $\mu_{\mathcal{L}}$ resemble the corresponding scaling laws at infinite width, while the scaling exponent for the coefficient of variation $\epsilon_{\mathcal{L}}$ is approximately zero at finite width and asymptotes to zero at infinite width for large $N_{\mathcal{D}}$. These results already indicate the importance of finite-width effects for UQ, which can change the scaling laws for $\sigma_{\mathcal{L}}$ by order-1 factors. However, $\epsilon_{\mathcal{L}}$ may be well-approximated by the infinite-width value implying that systematic uncertainty is dominated by $\mu_{\mathcal{L}}$ alone. In future work it would be interesting to minimize $\epsilon_{\mathcal{L}}$ as a function of λ_b/λ_W at infinite width and see if this also minimizes $\epsilon_{\mathcal{L}}$ at finite width; understand how this result relates to existing work using μ -parametrization; and compute $\epsilon_{\mathcal{L}}$ at finite width by generalizing the results of Sec. 3 using finite-width perturbation theory. Indeed, the latter may be feasible even in light of the $N_{\mathcal{D}}^4$ scaling of the finite-width tensors because of the independence of $\epsilon_{\mathcal{L}}$ from $N_{\mathcal{D}}$, as long as the scaling laws persist to sufficiently small $N_{\mathcal{D}}$. Finally, it would be interesting to apply finite-width perturbation theory to Bayesian neural networks (27) to see if the same scaling laws appear.

Acknowledgments

We thank Tilman Plehn, Christoph Weniger, and the participants of the IAIFI 2024 Summer Workshop for stimulating conversations. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under Award Number DE-SC0023704. This work utilizes the computing resources of the HAL cluster (28) supported by the National Science Foundation’s Major Research Instrumentation program, grant #1725729, as well as the University of Illinois Urbana-Champaign. This work used the TAMU FASTER cluster at Texas A&M University through allocation 240449 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program (29), which is supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #213829.

References

- [1] T. Y. Chen, B. Dey, A. Ghosh, M. Kagan, B. Nord and N. Ramachandra, *Interpretable Uncertainty Quantification in AI for HEP*, in *Snowmass 2021*, 8, 2022. 2208.03284. DOI.
- [2] I. Osband, C. Blundell, A. Pritzel and B. Van Roy, *Deep exploration via bootstrapped dqn*, *Advances in neural information processing systems* **29** (2016) .
- [3] B. Lakshminarayanan, A. Pritzel and C. Blundell, *Simple and scalable predictive uncertainty estimation using deep ensembles*, *Advances in neural information processing systems* **30** (2017) , [1612.01474].
- [4] A. Jacot, F. Gabriel and C. Hongler, *Neural tangent kernel: Convergence and generalization in neural networks*, *Advances in Neural Information Processing Systems* **31** (2018) .
- [5] S. Arora, S. S. Du, W. Hu, Z. Li, R. R. Salakhutdinov and R. Wang, *On exact computation with an infinitely wide neural net*, *Advances in Neural Information Processing Systems* **32** (2019) .
- [6] J. Lee, L. Xiao, S. Schoenholz, Y. Bahri, R. Novak, J. Sohl-Dickstein et al., *Wide neural networks of any depth evolve as linear models under gradient descent*, *Advances in Neural Information Processing Systems* **32** (2019) .
- [7] D. A. Roberts, S. Yaida and B. Hanin, *The Principles of Deep Learning Theory*. Cambridge University Press, 5, 2022, 10.1017/9781009023405.
- [8] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child et al., *Scaling laws for neural language models*, 2001.08361.

- [9] Y. Bahri, E. Dyer, J. Kaplan, J. Lee and U. Sharma, *Explaining neural scaling laws*, 2102.06701.
- [10] B. Bordelon, A. Atanasov and C. Pehlevan, *A dynamical model of neural scaling laws*, arXiv preprint arXiv:2402.01092 (2024) .
- [11] B. Bordelon, A. Canatar and C. Pehlevan, *Spectrum dependent learning curves in kernel regression and wide neural networks*, in *International Conference on Machine Learning*, pp. 1024–1034, PMLR, 2020.
- [12] A. Canatar, B. Bordelon and C. Pehlevan, *Spectral bias and task-model alignment explain generalization in kernel regression and infinitely wide neural networks*, *Nature communications* **12** (2021) 2914.
- [13] A. Atanasov, B. Bordelon, S. Sainathan and C. Pehlevan, *The onset of variance-limited behavior for networks in the lazy and rich regimes*, 2212.12147.
- [14] B. Bordelon and C. Pehlevan, *Dynamics of finite width kernel and prediction fluctuations in mean field neural networks*, *Advances in Neural Information Processing Systems* **36** (2024) .
- [15] A. Maloney, D. A. Roberts and J. Sully, *A Solvable Model of Neural Scaling Laws*, 2210.16859.
- [16] B. Adlam, J. Lee, S. Padhy, Z. Nado and J. Snoek, *Kernel regression with infinite-width neural networks on millions of examples*, arXiv preprint arXiv:2303.05420 (2023) .
- [17] Z. Zhang, *Neural Scaling Laws From Large-N Field Theory: Solvable Model Beyond the Ridgeless Limit*, 2405.19398.
- [18] C. Williams, *Computing with infinite networks*, *Advances in neural information processing systems* **9** (1996) .
- [19] G. Pang, L. Yang and G. E. Karniadakis, *Neural-net-induced gaussian process regression for function approximation and pde solution*, *Journal of Computational Physics* **384** (2019) 270–288.
- [20] Y. LeCun and C. Cortes, *MNIST handwritten digit database*, 2010. <http://yann.lecun.com/exdb/mnist/>.
- [21] A. Krizhevsky, *Learning multiple layers of features from tiny images*, .
- [22] D. Belayneh, F. Carminati, A. Farbin, B. Hooberman, G. Khattak, M. Liu et al., *Calorimetry with deep learning: particle simulation and reconstruction for collider physics*, *The European Physical Journal C* **80** (2020) .
- [23] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro et al., *TensorFlow: Large-scale machine learning on heterogeneous systems*, 2015.
- [24] D. P. Kingma and J. Ba, *Adam: A method for stochastic optimization*, 1412.6980.
- [25] S. Bollweg, M. Haußmann, G. Kasieczka, M. Luchmann, T. Plehn and J. Thompson, *Deep-Learning Jets with Uncertainties and More*, *SciPost Phys.* **8** (2020) 006, [1904.10004].
- [26] G. Kasieczka, M. Luchmann, F. Otterpohl and T. Plehn, *Per-Object Systematics using Deep-Learned Calibration*, *SciPost Phys.* **9** (2020) 089, [2003.11099].
- [27] R. Pacelli, S. Ariosto, M. Pastore, F. Ginelli, M. Gherardi and P. Rotondo, *A statistical mechanics framework for bayesian deep neural networks beyond the infinite-width limit*, *Nature Machine Intelligence* **5** (2023) 1497–1507.
- [28] V. Kindratenko, D. Mu, Y. Zhan, J. Maloney, S. H. Hashemi, B. Rabe et al., *Hal: Computer system for scalable deep learning*, in *Practice and Experience in Advanced Research Computing*, PEARC '20, (New York, NY, USA), p. 41–48, Association for Computing Machinery, 2020. DOI.

- [29] T. J. Boerner, S. Deems, T. R. Furlani, S. L. Knuth and J. Towns, *Access: Advancing innovation: Nsf’s advanced cyberinfrastructure coordination ecosystem: Services & support*, in *Practice and Experience in Advanced Research Computing 2023: Computing for the Common Good*, PEARC ’23, (New York, NY, USA), p. 173–176, Association for Computing Machinery, 2023. DOI.
- [30] GEANT4 collaboration, S. Agostinelli et al., *GEANT4—a simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250–303.
- [31] R. M. Neal, *Bayesian learning for neural networks*, vol. 118. Springer Science & Business Media, 2012.
- [32] A. G. d. G. Matthews, J. Hron, R. E. Turner and Z. Ghahramani, *Sample-then-optimize posterior sampling for bayesian linear models*, in *NeurIPS Workshop on Advances in Approximate Bayesian Inference*, 2017.
- [33] J. Lee, Y. Bahri, R. Novak, S. S. Schoenholz, J. Pennington and J. Sohl-Dickstein, *Deep neural networks as gaussian processes*, in *International Conference on Learning Representations*, 2018.
- [34] A. de G, M. Rowland, J. Hron, R. Turner and Z. Ghahramani, *Gaussian process behaviour in wide deep neural networks*, in *Published as a conference paper at ICLR*, vol. 2018, p. 15, 2018.
- [35] B. Adlam, J. Lee, L. Xiao, J. Pennington and J. Snoek, *Exploring the uncertainty properties of neural networks’ implicit priors in the infinite-width limit*, in *International Conference on Learning Representations*.

A Appendix

A.1 Derivation of mean loss and loss variance at infinite width

Here we evaluate the Gaussian expectations to compute the mean test loss and its variance for infinite-width neural networks, given their mean predictions m_β^∞ and the covariance matrix $\Sigma_{\beta_1\beta_2}$.

We start with the case where $n_L = 1$. The diagonal entries $\Sigma_{\beta\beta} \equiv (\sigma_\beta^\infty)^2$ were given in Eq. (9); the full covariance matrix entries are (6)

$$\Sigma_{\beta_1\beta_2} = K_{\beta_1\beta_2}^{(L)} - \Theta_{\beta_1}^T \Theta_A^{-1} \mathbf{K}_{\beta_2} - \Theta_{\beta_2}^T \Theta_A^{-1} \mathbf{K}_{\beta_1} + \Theta_{\beta_1}^T \Theta_A^{-1} \mathbf{K}_A \Theta_A^{-1} \Theta_{\beta_2}. \quad (11)$$

It will be convenient to shift the output $z_\beta^{(L)}$ by the mean, so that $\tilde{z}_\beta^{(L)} \equiv z_\beta^{(L)} - m_\beta^\infty$ are zero-mean Gaussian variables with the same covariance $\Sigma_{\beta_1\beta_2}$ as the unshifted outputs. Defining $\Delta_\beta \equiv y_\beta^\infty - m_\beta$ as the mean prediction error for each test point \tilde{x}_β , we have (7)

$$\begin{aligned} \mu_{\mathcal{L}} \equiv \mathbb{E}[\mathcal{L}_B(T)] &= \mathbb{E} \left[\frac{1}{2|\mathcal{B}|} \sum_{\beta \in B} (z_\beta^{(L)} - y_\beta)^2 \right] \\ &= \frac{1}{2|\mathcal{B}|} \sum_{\beta \in B} \mathbb{E} \left[(\tilde{z}_\beta^{(L)} - \Delta_\beta)^2 \right] \\ &= \frac{1}{2|\mathcal{B}|} \sum_{\beta \in B} (\Delta_\beta^2 + \Sigma_{\beta\beta}), \end{aligned} \quad (12)$$

where in the last line we have used $\mathbb{E}[\tilde{z}_\beta^{(L)}] = 0$ and $\mathbb{E}[(\tilde{z}_\beta^{(L)})^2] = \Sigma_{\beta\beta}$.

For the variance, we additionally need to compute

$$\begin{aligned}
\mathbb{E}[\mathcal{L}_{\mathcal{B}}^2] &= \frac{1}{4|\mathcal{B}|^2} \sum_{\beta_1, \beta_2 \in \mathcal{B}} \mathbb{E} \left[(z_{\beta_1}^{(L)} - y_{\beta_1})^2 (z_{\beta_2}^{(L)} - y_{\beta_2})^2 \right] \\
&= \frac{1}{4|\mathcal{B}|^2} \sum_{\beta_1, \beta_2 \in \mathcal{B}} \left\{ \mathbb{E} \left[(\tilde{z}_{\beta_1}^{(L)})^2 (\tilde{z}_{\beta_2}^{(L)})^2 \right] + \Delta_{\beta_1}^2 \mathbb{E} \left[(\tilde{z}_{\beta_2}^{(L)})^2 \right] + \Delta_{\beta_2}^2 \mathbb{E} \left[(\tilde{z}_{\beta_1}^{(L)})^2 \right] \right. \\
&\quad \left. + 4\Delta_{\beta_1} \Delta_{\beta_2} \mathbb{E} \left[\tilde{z}_{\beta_1}^{(L)} \tilde{z}_{\beta_2}^{(L)} \right] + \Delta_{\beta_1}^2 \Delta_{\beta_2}^2 \right\} \\
&= \frac{1}{4|\mathcal{B}|^2} \sum_{\beta_1, \beta_2 \in \mathcal{B}} \left\{ \Sigma_{\beta_1 \beta_1} \Sigma_{\beta_2 \beta_2} + 2\Sigma_{\beta_1 \beta_2}^2 + \Delta_{\beta_1}^2 \Sigma_{\beta_2 \beta_2} + \Delta_{\beta_2}^2 \Sigma_{\beta_1 \beta_1} + 4\Delta_{\beta_1} \Delta_{\beta_2} \Sigma_{\beta_1 \beta_2} + \Delta_{\beta_1}^2 \Delta_{\beta_2}^2 \right\},
\end{aligned} \tag{13}$$

where in the second line we have used the fact that odd moments of $\tilde{z}_{\beta}^{(L)}$ vanish, and in the third line we have used Wick's theorem for the quartic expectation. The loss variance is then given by

$$\sigma_{\mathcal{L}}^2 = \mathbb{E}[\mathcal{L}_{\mathcal{B}}^2] - \mu_{\mathcal{L}}^2, \tag{14}$$

where the first term is given in Eq. (13) and the second term is given in Eq. (12).

For $n_L \neq 1$, we index the output neurons with Latin letters i, j : $\tilde{z}_{\beta, i}^{(L)} \equiv z_{\beta, i}^{(L)} - m_{\beta, i}^{\infty}$ and $\Delta_{\beta, i} \equiv y_{\beta, i}^{\infty} - m_{\beta, i}$. At infinite width, output neurons are perfectly uncorrelated so the covariance matrix is diagonal in neural indices,

$$\mathbb{E} \left[\tilde{z}_{\beta_1, i}^{(L)} \tilde{z}_{\beta_1, j}^{(L)} \right] = \delta_{ij} \Sigma_{\beta_1 \beta_2}. \tag{15}$$

The mean loss is

$$\begin{aligned}
\mu_{\mathcal{L}} &= \mathbb{E} \left[\frac{1}{2|\mathcal{B}|n_L} \sum_{\beta \in \mathcal{B}} \sum_{i=1}^{n_L} (z_{\beta, i}^{(L)} - y_{\beta, i})^2 \right] \\
&= \frac{1}{2|\mathcal{B}|n_L} \sum_{\beta \in \mathcal{B}} \sum_{i=1}^{n_L} \mathbb{E} \left[(\tilde{z}_{\beta, i}^{(L)} - \Delta_{\beta, i})^2 \right] \\
&= \frac{1}{2|\mathcal{B}|n_L} \sum_{\beta \in \mathcal{B}} \left(n_L \Sigma_{\beta \beta} + \sum_{i=1}^{n_L} \Delta_{\beta, i}^2 \right) \\
&= \frac{1}{2|\mathcal{B}|n_L} \sum_{\beta \in \mathcal{B}} \left(n_L \Sigma_{\beta \beta} + \|\vec{\Delta}_{\beta}\|^2 \right),
\end{aligned} \tag{16}$$

where the factor of n_L comes from summing over n_L identical copies of the covariance $\mathbb{E} \left[(\tilde{z}_{\beta, i}^{(L)})^2 \right]$, and in the last line we have changed to vector notation, $\sum_{i=1}^{n_L} \Delta_{\beta, i}^2 \equiv \|\vec{\Delta}_{\beta}\|^2 = \vec{\Delta}_{\beta} \cdot \vec{\Delta}_{\beta}$.

Similarly, for the variance we need

$$\begin{aligned}
\mathbb{E}[\mathcal{L}_{\mathcal{B}}^2] &= \frac{1}{(2|\mathcal{B}|n_L)^2} \sum_{\beta_1, \beta_2 \in \mathcal{B}} \sum_{i, j=1}^{n_L} \mathbb{E} \left[(z_{\beta_1, i}^{(L)} - y_{\beta_1, i})^2 (z_{\beta_2, j}^{(L)} - y_{\beta_2, j})^2 \right] \\
&= \frac{1}{(2|\mathcal{B}|n_L)^2} \sum_{\beta_1, \beta_2 \in \mathcal{B}} \sum_{i, j=1}^{n_L} \left\{ \mathbb{E} \left[(\tilde{z}_{\beta_1, i}^{(L)})^2 (\tilde{z}_{\beta_2, j}^{(L)})^2 \right] + \Delta_{\beta_1, i}^2 \mathbb{E} \left[(\tilde{z}_{\beta_2, j}^{(L)})^2 \right] + \Delta_{\beta_2, j}^2 \mathbb{E} \left[(\tilde{z}_{\beta_1, i}^{(L)})^2 \right] \right. \\
&\quad \left. + 4\Delta_{\beta_1, i} \Delta_{\beta_2, j} \mathbb{E} \left[\tilde{z}_{\beta_1, i}^{(L)} \tilde{z}_{\beta_2, j}^{(L)} \right] + \Delta_{\beta_1, i}^2 \Delta_{\beta_2, j}^2 \right\} \\
&= \frac{1}{(2|\mathcal{B}|n_L)^2} \sum_{\beta_1, \beta_2 \in \mathcal{B}} \left[n_L^2 \Sigma_{\beta_1 \beta_1} \Sigma_{\beta_2 \beta_2} + 2n_L \Sigma_{\beta_1 \beta_2}^2 + n_L \sum_{i=1}^{n_L} \left\{ \Delta_{\beta_1, i}^2 \Sigma_{\beta_2 \beta_2} + \Delta_{\beta_2, i}^2 \Sigma_{\beta_1 \beta_1} \right\} \right. \\
&\quad \left. + 4\Sigma_{\beta_1 \beta_2} \sum_{i=1}^{n_L} \left\{ \Delta_{\beta_1, i} \Delta_{\beta_2, i} \right\} + \sum_{i, j=1}^{n_L} \Delta_{\beta_1, i}^2 \Delta_{\beta_2, j}^2 \right] \\
&= \frac{1}{(2|\mathcal{B}|n_L)^2} \sum_{\beta_1, \beta_2 \in \mathcal{B}} \left[n_L^2 \Sigma_{\beta_1 \beta_1} \Sigma_{\beta_2 \beta_2} + n_L \left(2\Sigma_{\beta_1 \beta_2}^2 + \|\vec{\Delta}_{\beta_1}\|^2 \Sigma_{\beta_2 \beta_2} + \|\vec{\Delta}_{\beta_2}\|^2 \Sigma_{\beta_1 \beta_1} \right) \right. \\
&\quad \left. + 4\Sigma_{\beta_1 \beta_2} \vec{\Delta}_{\beta_1} \cdot \vec{\Delta}_{\beta_2} + \|\vec{\Delta}_{\beta_1}\|^2 \|\vec{\Delta}_{\beta_2}\|^2 \right]. \tag{17}
\end{aligned}$$

As before, the variance $\sigma_{\mathcal{L}}^2$ is computed by subtracting Eq. (16) from Eq. (17).

A.2 Calorimeter Data Details

Particle colliders and their detectors are the prototypical tools in HEP. These detectors include calorimeters, which measure particles' energies and are composed of an array of detecting elements, or "cells," where particles deposit energy. This results in a spatial distribution of energy deposition that is used to both identify entering particles and quantify their energies. Thus, machine learning can be used for both (particle) classification and (energy) regression; in this work we focus on the latter.

Following Ref. (22), we generate our dataset using GEANT4 (30), simulating the proposed Linear Collider Detector (LCD) at the CLIC accelerator. In the LCD, we simulate an electromagnetic calorimeter (ECAL) and a hadronic calorimeter (HCAL). The ECAL is structured with 25 silicon sensor planes interspersed with tungsten absorber layers and arranged in a cylindrical geometry with square cells measuring 5.1 mm on each side. The HCAL, positioned behind the ECAL, comprises 60 layers of polystyrene scintillators with steel absorbers, segmented into larger 3×3 cm² square cells. The dataset is formed by outputting 3D arrays of energy deposits in each of the ECAL and HCAL cells, which due to their geometry gives us, for each event, an $11 \times 11 \times 60$ array for HCAL and a $25 \times 25 \times 25$ array for ECAL. These arrays are then flattened into a large 22885-dimensional vector.

While many particles can in principle be simulated, we specialize to a sample of 10^5 electrons entering the HCAL and ECAL (where one electron entering corresponds to an event and thus a datapoint) with a range of incident angles close to the beam direction and a range of 10 – 100 GeV of incoming energy. This incoming energy is set as the label corresponding to our 22885-dimensional input vector. We take into account the effects of solenoidal magnetic field and other materials preceding the calorimeters.

Our dataset of 10^5 electron events can be found here: <https://zenodo.org/records/13715377>.

A.3 Variation of training algorithm

While our infinite-width results apply, strictly speaking, only for full-batch gradient descent, we can investigate how robust the $N_{\mathcal{D}}$ independence of $\epsilon_{\mathcal{L}}$ is when we vary the training algorithm. In Figs. 2 and 3 we show the mean test loss and coefficient of variation as in Fig. 1, but using networks trained with the Adam optimizer with $\eta = 10^{-3}$ for MNIST and $\eta = 10^{-6}$ for calorimeter data. For both datasets, we use $\beta_1 = .9$, $\beta_2 = .999$ and with a minibatch size of 1000. The results are somewhat noisier, but the overall trend is for a roughly flat scaling law for $\epsilon_{\mathcal{L}}$ (right panels). However, while the $\mu_{\mathcal{L}}$ scaling laws for MNIST agree almost perfectly between the two optimizers, the $\mu_{\mathcal{L}}$ scaling laws

for the calorimeters are markedly different, and the values for $\epsilon_{\mathcal{L}}$ differ by more than a factor of 2. It would be interesting to investigate in future work to what extent the different behavior of MNIST and the calorimeter data is a consequence of finite-width feature learning.

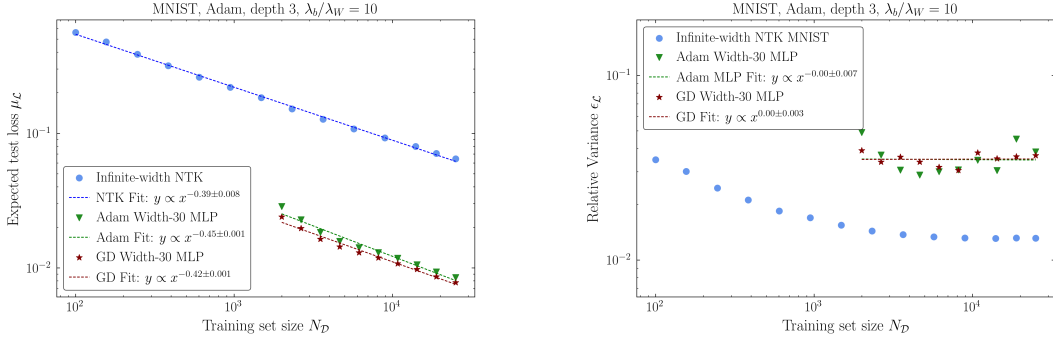


Figure 2: Mean loss and coefficient of variation scaling laws using the Adam optimizer in MNIST. For this example, Adam and full-batch GD show very similar scaling laws for both $\mu_{\mathcal{L}}$ and $\epsilon_{\mathcal{L}}$.

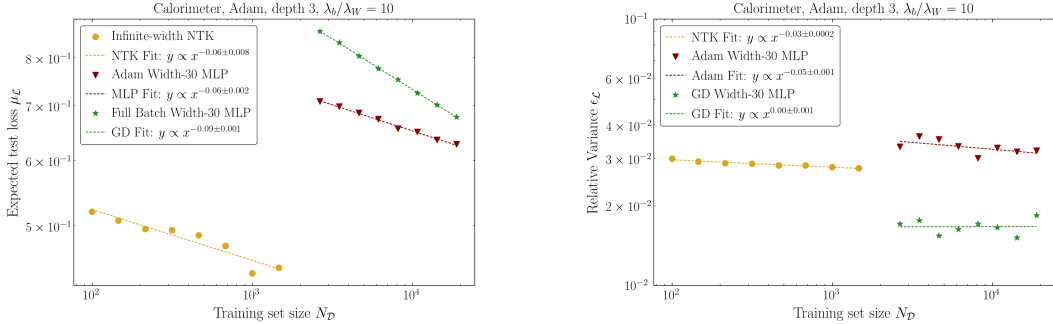


Figure 3: Mean loss and coefficient of variation scaling laws using the Adam optimizer in calorimeter data. For this example, the $\epsilon_{\mathcal{L}}$ scaling appears to be flat for Adam as it was for full-batch GD, but the $\mu_{\mathcal{L}}$ scaling law is markedly different.

A.4 Measurements of kernel and NTK element scalings with dataset size

For the case where the NTK is invertible, we can attempt to gain some intuition for the $N_{\mathcal{D}}$ independence of the coefficient of variation statistic $\epsilon_{\mathcal{L}}$ by ignoring the matrix structure of the infinite-width predictions (9) and simply studying the scaling behavior of the kernel and NTK elements. For the MNIST dataset where the NTK is invertible, in Fig. 4 we show the mean matrix elements for the test-train NTK $\Theta_{\mathcal{B}}$, the training set kernel $\mathbf{K}_{\mathcal{A}}$, and the inverse of the training NTK, $\Theta_{\mathcal{A}}^{-1}$. Only the latter scales nontrivially with $N_{\mathcal{D}}$ (left panel). In the right panel, we show the scalings of the mean diagonal and off-diagonal elements, $\Theta_{\alpha\alpha}^{-1}$ and $\Theta_{\alpha\beta}^{-1}$, respectively, of the inverse NTK. The scalings are different, with the diagonal elements growing with $N_{\mathcal{D}}$ while the off-diagonal elements decrease. It would be interesting to investigate further how these scaling laws relate to generalization behavior in the infinite-width limit.

A.5 Bayesian interpretation of infinite-width ensemble

In physics, model uncertainty is typically expressed as the width of a Bayesian posterior, which quantifies the degree of confidence in the prediction given the training data. While neural networks trained under gradient descent do not perform exact Bayesian model fitting at finite width (7), infinite-width MLPs *can* perform Bayesian inference with a modified gradient descent algorithm where only the last layer parameters are updated (18; 31; 32; 33; 34; 6). The Bayesian posterior is still Gaussian, with mean and covariance given by Eq. (9) with all instances of Θ replaced by the kernel \mathbf{K} . It is therefore interesting to investigate whether the infinite-width NTK scaling laws resemble those

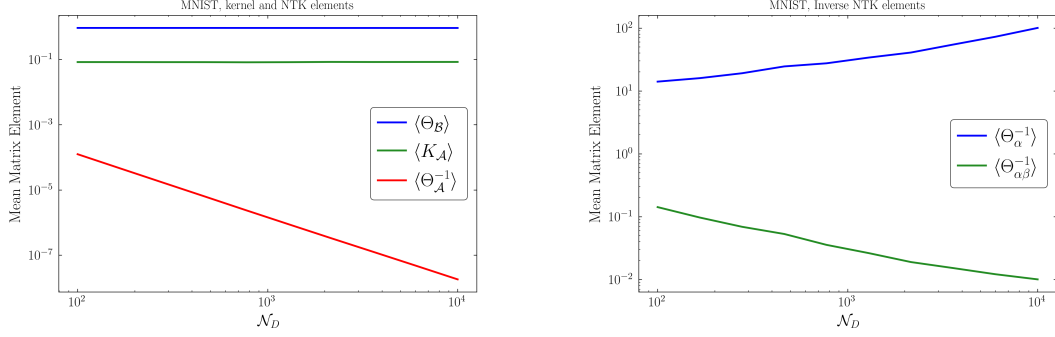


Figure 4: Scaling of mean matrix elements with training set size N_D . Left: The only matrix contributing to the mean and variance of the infinite-width prediction that appears to scale nontrivially with N_D is the training NTK inverse $\Theta_{\mathcal{A}}^{-1}$. Right: the diagonal and off-diagonal elements of the NTK scale differently.




Figure 5: Comparison of Bayesian inference and gradient descent in infinite-width networks trained on MNIST. Left: scaling of mean loss $\mu_{\mathcal{L}}$. Right: scaling of coefficient of variation $\epsilon_{\mathcal{L}}$.

for Bayesian inference, which would imply that the infinite-width variance $\sigma_{\mathcal{L}}$ (and possibly the finite-width variance as well) can be interpreted as a true Bayesian credible interval (35). In Fig. 5, we show the results of this comparison for the MNIST data. In order to obtain the Bayesian prediction, the kernel has to be (numerically) invertible, which in our examples is only true for MNIST. We see that the Bayesian and NTK results for the mean agree extremely well. Furthermore, while the coefficient of variation differs somewhat between the two training algorithms at small N_D , both Bayesian inference and GD training asymptote to the same value of $\epsilon_{\mathcal{L}}$ at large N_D , with variations only at the percent level. The similarity of infinite-width Bayesian inference (also called neural network Gaussian processes) to infinite-width GD (or NTK Gaussian processes) was noted in the first papers to study the infinite-width NTK dynamics for deep networks (6; 35), but to our knowledge the agreement for $\epsilon_{\mathcal{L}}$ at large N_D is a new result. In future work it would be interesting to investigate how robust this result is to both the depth of the network L and the choice of learning rate hyperparameter λ_b/λ_W .