

On Neural Network Kernels and the Storage Capacity Problem

Jacob A. Zavatone-Veth

jzavatoneveth@g.harvard.edu

Department of Physics and Center for Brain Science, Harvard University, Cambridge, MA 02138, U.S.A.

Cengiz Pehlevan

cpehlevan@seas.harvard.edu

Center for Brain Science and John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, U.S.A.

In this short note, we reify the connection between work on the storage capacity problem in wide two-layer treelike neural networks and the rapidly growing body of literature on kernel limits of wide neural networks. Concretely, we observe that the “effective order parameter” studied in the statistical mechanics literature is exactly equivalent to the infinite-width neural network gaussian process kernel. This correspondence connects the expressivity and trainability of wide two-layer neural networks.

The study of two-layer neural networks in the limit of large hidden layer width has a long history in the statistical physics of learning (Barkai, Hansel, & Sompolinsky, 1992; Engel, Köhler, Tschepke, Vollmayr, & Zipelius, 1992). This work focuses on the Gardner storage capacity problem, which measures the computational power of a neural network by the largest random binary classification data set it can “memorize” (Gardner, 1988; Gardner & Derrida, 1988). In this problem, one performs Bayesian inference of network weights using a likelihood that is flat on the set of weights that yield zero classification error and vanishes uniformly otherwise. Memorization is possible if the support of the resulting Bayes posterior has nonzero volume, that is, if there exists a nonnegligible set of weights for which all examples are correctly classified. In the thermodynamic limit where the input dimensionality N and data set size P tend to infinity with fixed ratio $\alpha \equiv P/N$, there is a sharp transition between memorization with probability one and memorization with probability zero at some critical ratio α_c , referred to as the storage capacity.

These statistical mechanics calculations rely on studying functions of the overlap q between the hidden unit weight vectors of identical copies, known as replicas, of the network, which is termed the “order parameter.” Classic

work by Barkai et al. (1992) and Engel et al. (1992) for treelike networks with sign function hidden layer activations revealed a remarkable simplification in the limit where the number of hidden units K was taken to be large: the equations reduced to those of a perceptron (Gardner, 1988; Gardner & Derrida, 1988), with the order parameter replaced by an “effective” order parameter $q_{\text{eff}}(q)$. This result was recently extended to rectified linear unit activation functions by Baldassi et al. (2019) and general activation functions in our own work (Zavatone-Veth & Pehlevan, 2021, hereafter ZP).

Beginning with pioneering work by Neal (1996) and Williams (1997), a parallel line of research in the machine learning community has characterized the infinite-width limits of Bayesian neural networks with gaussian priors over their weights. This work also revealed a remarkable simplification: for a two-layer network with hidden layer activation function f , infinite-width inference is equivalent to shallow gaussian process (GP) inference with kernel¹

$$K_f(\mathbf{x}, \mathbf{y}) = \mathbb{E} [f(\mathbf{w} \cdot \mathbf{x})f(\mathbf{w} \cdot \mathbf{y})] : \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (1)$$

given by the deterministic infinite-width limit of the Gram matrix of hidden layer activations. Here, $\mathbf{x}, \mathbf{y} \in \mathbb{R}^N$ are interpreted as two inputs to the network, and $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is a random weight vector. In recent years, this neural network–gaussian process (NNGP) correspondence has proven to be an extremely fruitful tool in the theory of deep learning thanks to the observation that it can be extended to networks with more than a single hidden layer (Lee et al., 2018; Matthews et al., 2018). Importantly, this correspondence does not apply only at the level of the prior, but also at the level of the Bayes posterior. For sensible likelihoods that model the targets as being conditionally independent of the network weights given the network outputs, the function-space posterior distribution tends to that induced by the limiting gaussian process prior (Hron et al., 2020). Moreover, the connection between infinitely wide neural networks and kernel machines has been extended to study gradient-based training (Jacot, Gabriel, & Hongler, 2018).

However, a precise connection between these two lines of research on Bayesian inference in infinitely wide two-layer neural networks has yet to be drawn. In this brief note, we reinterpret results on the storage capacity problem in terms of the behavior of the NNGP kernel. This observation allows us to clarify the conceptual and technical connections between these parallel lines of work on wide and deep neural networks.

We start by observing that for unit-norm inputs $\mathbf{x}, \mathbf{y} \in \mathbb{S}^{N-1}$, the infinite-width NNGP kernel, equation 1, can be expressed as a function of the

¹For clarity of exposition, we make the simplifying assumptions that the network has no bias terms and the prior weight variance is unity.

overlap $q \equiv \mathbf{x} \cdot \mathbf{y} \in [-1, +1]$ (Cho & Saul, 2009; Lee et al., 2018; Matthews, Hron, Rowland, Turner, & Ghahramani, 2018; Neal, 1996; Williams, 1997):

$$K_f(q) = \begin{cases} \mathbb{E}[f(x)^2 : x \sim \mathcal{N}(0, 1)], & \text{if } q = 1 \\ \mathbb{E}\left[f(x)f(y) : \begin{pmatrix} x \\ y \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & q \\ q & 1 \end{pmatrix}\right)\right], & \text{if } -1 < q < 1 \\ \mathbb{E}[f(x)f(-x) : x \sim \mathcal{N}(0, 1)], & \text{if } q = -1. \end{cases} \quad (2)$$

For functions f that are square-integrable with respect to gaussian measure, $K_f(q)$ is a continuous function on the closed interval $[-1, 1]$, with a power series that converges throughout that interval (Bogachev, 1998; Daniely, Frostig, & Singer, 2016; Zavatone-Veth & Pehlevan, 2021).

The NNGP kernel coincides exactly with the effective order parameter studied in the statistical mechanics literature:²

$$q_{\text{eff}}(q) = K_f(q). \quad (3)$$

In particular, the result of ZP shows that the storage capacity remains finite in the infinite-width limit if the left derivative,

$$\partial_- K_f(q = 1) \equiv \lim_{q \uparrow 1} \frac{K_f(1) - K_f(q)}{1 - q}, \quad (4)$$

of the kernel at $q = 1$ is finite, and diverges otherwise. For the special cases of $f(x) = \text{sign}(x)$, which yields infinite capacity, and $f(x) = \text{ReLU}(x) = \max\{0, x\}$, which yields finite capacity, this was noted by Barkai et al. (1992) and Baldassi et al. (2019), respectively, based on direct computations of $K_{\text{sign}}(q)$ and $K_{\text{ReLU}}(q)$.³ Thus, the relationship of the limiting behavior of the NNGP kernel for sign activation functions to the expressivity of infinitely wide networks was implicitly studied 30 years ago.

Therefore, the storage capacity of a treelike committee machine is related to the behavior of the NNGP kernel for nearly colinear arguments. In ZP, we gave a general argument based on Fourier-Hermite expansions that $\partial_- K_f(q = 1)$ is finite if and only if the activation function f is in the Sobolev space of functions that are square-integrable with respect to gaussian measure and have weak derivatives that are also square-integrable with respect to gaussian measure (Bogachev, 1998). Roughly speaking, a kernel with a

²Possibly up to irrelevant constant offsets. In ZP we defined the effective order parameter as $q_{\text{eff}}(q) = K_f(q) - [\mathbb{E}_{x \sim \mathcal{N}(0,1)} f(x)]^2$.

³In the machine learning literature, kernels of this family were studied systematically by Cho and Saul (2009).

cusp at $q = 1$ will yield divergent capacity. Intuitively, $\partial_- K_f(q = 1)$ measures the ability of the kernel to discriminate between nearly colinear inputs, which has a natural relationship to expressivity (Paccolat, Spigler, & Wyart, 2021; Poole, Lahiri, Raghu, Sohl-Dickstein, & Ganguli, 2016).

The relationship between $\partial_- K_f(q = 1)$ and the expressivity of deep networks was noted in a different context by Poole et al. (2016). Those authors studied the expressivity of infinitely wide and deep, fully connected networks with random weights in terms of the fixed points of $K_f(q)$ under iteration. For networks without bias terms, their results have a close relationship to those of ZP, which we will illustrate in a simple setting. We assume that all inputs lie on the sphere and normalize the activation function such that $K_f(q = 1) = 1$. Then the result of Poole et al. (2016) shows that networks with $\partial_- K_f(q = 1) > 1$ display chaotic behavior in the sense that the fixed points of $K_f(q)$ are unstable, while those with $\partial_- K_f(q = 1) < 1$ display ordered behavior. Those authors proposed that the sharp enhancement of differences between inputs by networks in the chaotic regime is a signature of expressivity. We remark that capacity calculation is conceptually distinct from the settings of GP inference and Poole et al. (2016) in that the effective order parameter measures the similarity between the weight vectors of two different replicas for a random input example, while the kernel measures the similarity of two input examples for a random weight vector.

The connection between the storage capacity problem and the NNGP kernel also helps to clarify the relationship of ZP to Panigrahi, Shetty, and Goyal's (2020) study of gradient descent training in two-layer networks with fixed readout weights. In their setting, the speed of training is governed by the minimum eigenvalue of the "gradient Gram matrix" evaluated on the training examples, schematically given as $G_f(\mathbf{x}, \mathbf{y}) = \mathbf{x} \cdot \mathbf{y} K_f'(\mathbf{x}, \mathbf{y})$.⁴ Working under the assumption that all inputs lie on the sphere, they showed that activation functions with a discontinuity in their derivatives yield rapid gradient descent training under a weaker bound on the maximum overlap q between any two training examples than is required for smooth activation functions.

As we have $G_f(q) = qK_f'(q)$ for inputs on the sphere, we can apply the above intuition for the relationship between the left derivative $\partial_- K_f(q = 1)$ and discriminability. By the results noted, functions f with discontinuous first derivatives will result in divergent left derivatives $\partial_- G_f(q = 1)$ and thus sharp discrimination capabilities. This sharp discrimination capability will yield better separation between diagonal and off-diagonal elements of $G_f(q)$ for q near one. By the Gershgorin circle theorem, this should in turn yield better lower bounds on the minimum eigenvalue of the G_f as a function of q than would hold for $G_f(q)$ with finite $\partial_- G_f(q = 1)$, possibly

⁴Here, the activation function f is always assumed to be at least Lipschitz continuous and is assumed to have weak derivative f' that is square-integrable with respect to gaussian measure such that G_f is well defined.

decreasing training time (Horn & Johnson, 2012). Therefore, the link between trainability and the second weak derivative of the activation function noted by Panigrahi et al. (2020) and the link between storage capacity and the first weak derivative of the activation function noted in ZP bear a close conceptual relation: both depend on the discrimination capabilities of the appropriate kernel for nearly colinear inputs.

To conclude, the storage capacity of a wide two-layer treelike neural network is determined by the behavior of the corresponding NNGP kernel for nearly colinear inputs on the sphere. This connection yields an intuitive explanation for the varying behavior of the storage capacity in terms of input discriminability (Baldassi et al., 2019; Barkai et al., 1992; Engel et al., 1992; Gardner, 1988; Gardner & Derrida, 1988; Zavatone-Veth & Pehlevan, 2021), as well as a more precise description of the relationship of this line of research to work on kernel limits of neural networks (Cho & Saul, 2009; Hron et al., 2020; Jacot et al., 2018; Lee et al., 2018; Matthews et al., 2018; Neal, 1996; Panigrahi et al., 2020; Poole et al., 2016; Williams, 1997).

Further research will be required to fully understand the connections between studies of the storage capacity problem in statistical physics and results on the NNGP limit. Though both lines of research focus on Bayesian inference of network weights, they consider different settings. The statistical physics literature largely focuses on simple two-layer models in a limit where the input dimension tends to infinity with the hidden layer width (Baldassi et al., 2019; Barkai et al., 1992; Engel et al., 1992; Zavatone-Veth & Pehlevan, 2021), while most studies of the NNGP limit consider deeper fully connected networks with finite-dimensional inputs (Hron et al., 2020; Lee et al., 2018; Matthews et al., 2018; Neal, 1996; Williams, 1997). Given the growing interest in applying tools from statistical physics to study inference in wide neural networks,⁵ we hope that the connections noted in this work will spark more detailed investigation of possible commonalities between these seemingly disparate settings.

Acknowledgments

This work was supported by the Harvard Data Science Initiative Competitive Research Fund, the Harvard Dean's Competitive Fund for Promising Scholarship, and a Google Faculty Research Award.

References

- Baldassi, C., Malatesta, E. M., & Zecchina, R. (2019). Properties of the geometry of solutions and capacity of multilayer neural networks with rectified linear unit activations. *Physical Review Letters*, 123(17), 170602. 10.1103/PhysRevLett.123.170602, PubMed: 31702271

⁵See our discussion in Zavatone-Veth, Canatar, Ruben, and Pehlevan (2021).

- Barkai, E., Hansel, D., & Sompolinsky, H. (1992). Broken symmetries in multilayered perceptrons. *Physical Review A*, 45(6), 4146. 10.1103/PhysRevA.45.4146
- Bogachev, V. I. (1998). *Gaussian measures*. Providence, RI: American Mathematical Society.
- Cho, Y., & Saul, L. (2009). Kernel methods for deep learning. In Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, & A. Culotta (Eds.), *Advances in neural information processing systems*, 22. Red Hook, NY: Curran.
- Daniely, A., Frostig, R., & Singer, Y. (2016). Toward deeper understanding of neural networks: The power of initialization and a dual view on expressivity. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in neural information processing systems*, 29 (pp. 2253–2261). Red Hook, NY: Curran.
- Engel, A., Köhler, H., Tschepke, F., Vollmayr, H., & Zippelius, A. (1992). Storage capacity and learning algorithms for two-layer neural networks. *Physical Review A*, 45(10), 7590. 10.1103/PhysRevA.45.7590
- Gardner, E. (1988). The space of interactions in neural network models. *Journal of Physics A: Mathematical and General*, 21(1), 257.
- Gardner, E., & Derrida, B. (1988). Optimal storage properties of neural network models. *Journal of Physics A: Mathematical and General*, 21(1), 271.
- Horn, R. A., & Johnson, C. R. (2012). *Matrix analysis*. Cambridge: Cambridge University Press.
- Hron, J., Bahri, Y., Novak, R., Pennington, J., & Sohl-Dickstein, J. (2020). *Exact posterior distributions of wide Bayesian neural networks*. arXiv:2006.10541.
- Jacot, A., Gabriel, F., & Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *Advances in neural information processing systems*, 31. Red Hook, NY: Curran.
- Lee, J., Sohl-Dickstein, J., Pennington, J., Novak, R., Schoenholz, S., & Bahri, Y. (2018). Deep neural networks as gaussian processes. In *Proceedings of the International Conference on Learning Representations*.
- Matthews, A. G. d. G., Hron, J., Rowland, M., Turner, R. E., & Ghahramani, Z. (2018). Gaussian process behaviour in wide deep neural networks. In *Proceedings of the International Conference on Learning Representations*.
- Neal, R. M. (1996). *Bayesian learning for neural networks* (pp. 29–53). Berlin: Springer.
- Paccolat, J., Spigler, S., & Wyart, M. (2021). How isotropic kernels perform on simple invariants. *Machine Learning: Science and Technology*, 2(2), 025020. 10.1088/2632-2153/abd485, PubMed: 34396127
- Panigrahi, A., Shetty, A., & Goyal, N. (2020). Effect of activation functions on the training of overparameterized neural nets. In *Proceedings of the International Conference on Learning Representations*.
- Poole, B., Lahiri, S., Raghu, M., Sohl-Dickstein, J., & Ganguli, S. (2016). Exponential expressivity in deep neural networks through transient chaos. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in neural information processing systems*, 29 (pp. 3360–3368). Red Hook, NY: Curran.
- Williams, C. K. (1997). Computing with infinite networks. In M. Jordan, M. Kearns, & S. Solla (Eds.), *Advances in neural information processing systems*, 10 (pp. 295–301). Cambridge, MA: MIT Press.

- Zavatone-Veth, J. A., Canatar, A., Ruben, B. S., & Pehlevan, C. (2021). Asymptotics of representation learning in finite Bayesian neural networks. In M. Ranzato, A. Beygelzimer, K. Nguyen, P. S. Liang, J. W. Vaughan, & Y. Dauphin (Eds.), *Advances in neural information processing systems*, 34. Red Hook, NY: Curran.
- Zavatone-Veth, J. A. & Pehlevan, C. (2021). Activation function dependence of the storage capacity of treelike neural networks. *Physical Review E*, 103, L020301. 10.1103/PhysRevE.103.L020301, PubMed: 33736039

Received November 14, 2021; accepted January 13, 2022.