Globally Gated Deep Linear Networks

Anonymous Author(s)
Affiliation
Address
email

Abstract

Recently proposed Gated Linear Networks (GLNs) present a tractable nonlinear network architecture, and exhibit interesting capabilities such as learning with local error signals and reduced forgetting in sequential learning. In this work, we introduce a novel gating architecture, named Globally Gated Deep Linear Networks (GGDLNs) where gating units are shared among all processing units in each layer, thereby decoupling the architectures of the nonlinear but unlearned gating and the learned linear processing motifs. We derive exact equations for the generalization properties of Bayesian Learning in these networks in the finite-width thermodynamic limit, defined by $N, P \to \infty$ while $P/N = O(1)$ where N and P are the hidden layers' width and size of training data sets respectfully. We find that the statistics of the network predictor can be expressed in terms of kernels that undergo shape renormalization through a data-dependent order parameter matrix compared to the infinite-width Gaussian Process (GP) kernels. Our theory accurately captures the behavior of finite width GGDLNs trained with gradient descent (GD) dynamics. We show that kernel shape renormalization gives rise to rich generalization properties w.r.t. network width, depth and L_2regularization amplitude. Interestingly, networks with a large number of gating units behave similarly to standard ReLU architectures. Although gating units in the model do not participate in the supervised learning, we show the utility of unsupervised learning of the gating parameters. Additionally, our theory allows the evaluation of the network's ability for learning multiple tasks by incorporating task-relevant information into the gating units. In summary, our work is the first exact theoretical solution of learning of a family of nonlinear networks with finite width. The rich and diverse behavior of the GGDLNs suggests that they are helpful analytically tractable models of learning single and multiple tasks, in finite-width nonlinear deep networks.

1 Introduction

Despite the recent advances in machine learning, theoretical understanding of how machine learning algorithms work is very limited. Many current theoretical approaches study neural network generalization in infinitely wide networks (3; 35; 34), where the input-output relation is equivalent to a Gaussian Process (GP) in function space with a covariance matrix defined by a GP kernel. However, this GP limit holds when the network width approaches infinity while the size of the training data remains finite, severely limiting its applicability to realistic conditions. Another line of work focuses on finite-width deep linear neural networks (DLNNs)(29; 32; 31), while applicable in a wider regime, the generalization behavior of linear networks are very limited, and the bias contribution always remains constant with network parameters (29), which fails to capture the behavior of generalization performance in general nonlinear networks. Therefore, a tractable nonlinear network architecture is in need for theoretically probing into the diverse generalization behavior of general nonlinear networks.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

Recently proposed Gated Linear Networks (GLNs) present a tractable nonlinear network architecture (25; 28; 26), with capabilities such as learning with local error signals and mitigating catastrophic forgetting in sequential learning. Inspired by these recent advances in GLNs, we propose Globally Gated Deep Linear Networks (GGDLNs) as a simplified GLN structure that preserves the nonlinear property of general GLNs, the decoupling of fixed nonlinear gating from learned linear processing units and the ability to separate the processing of multiple tasks using the gating units. Our GGDLN structure is different from previous GLNs in several ways. First, the gating units are shared across hidden layer units and different layers while in previous work each unit has its own set of gatings (13; 28; 26). Second, we define global learning objective instead of local errors (28; 26). These simplifications allow us to obtain direct analytical expressions of memory capacity and exact generalization error of these networks for arbitrary training and testing data, providing quantitative insight into the effect of learning in nonlinear networks, as opposed to studies of generalization bounds (13), expressivity estimates (26; 12), and indirect quantities relevant to generalization such as the implicit bias of the network (24). Furthermore, the kernel expression of the predictor statistics we propose in this work also allow us to make qualitative explanations of the generalization and how it's related to data structure and network representation for single and multiple tasks.

First, we introduce the architecture of our GGDLNs and analyze its memory capacity. We then derive our theory for generalization properties of GGDLNs, and make qualitative connections between the generalization behavior and the relation between the renormalization matrix and task structure. Second, we apply our theory to GGDLNs performing multiple tasks, focusing on two scenarios where tasks are defined by different input statistics and different output labels on the same inputs respectively. While the effect of kernel renormalization is different in the two cases, we find that for fixed gatings, de-correlation between tasks always improves generalization.

2 Globally gated deep linear networks

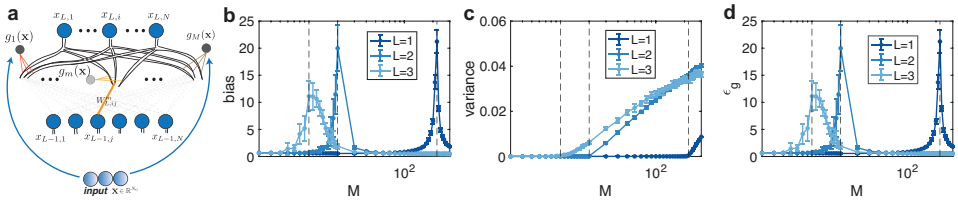

Figure 1: Globally gated deep linear network. (**a**)Structure of GGDLNs, each neuron in the hidden layer has M dendrites, each with a different input-dependent gating $g_m(\mathbf{x})$ which is fixed during training, the M gatings are shared across neurons in the hidden layer. The m-th dendritic branch of the i-th neuron in layer L connects to neuron j in the previous layer with weight W_{ij}^m(shown in orange). (**b**-**d**) Bias, variance and generalization error of networks with 1-3 hidden layers in the GP limit evaluated on MNIST (4). Bias and generalization error diverges, and variance becomes nonzero at network capacity (black dashed line). See Appendix C.2 for detailed parameters.

In GGDLNs, the network input-output relation is defined as follows,

$$f(\mathbf{x}) = \frac{1}{\sqrt{NM}} \sum_{i=1}^{N} \sum_{m=1}^{M} a_{m,i} x_{L,i} g_m(\mathbf{x}),\ x_{L,i} = \begin{cases} \frac{1}{\sqrt{N_0 M}} \sum_{j=1}^{N_0} \sum_{m=1}^{M} W_{L,ij}^m g_m(\mathbf{x}) x_{L-1,j} & L > 1 \\ \frac{1}{\sqrt{N_0}} \sum_{j=1}^{N_0} W_{L,ij} x_{L-1,j} & L = 1 \end{cases} \tag{1}$$

where $\mathbf{x}_0 = \mathbf{x}$ is the input, N is the hidden layer width, M is the number of gating units in each layer, and N_0 is the input dimension. Neurons in each layer has M dendrites, each with an input-dependent global gating $g_m(\mathbf{x})$ shared across all neurons. The m-th dendritic branch of neuron i in the L-th hidden layer connects to neurons in the previous layer with a dendrite-specific weight vector $\mathbf{W}_{L,i}^m$ (or with readout weight vector $\mathbf{a}_m$ for the output neuron), as shown in Fig.1 (a). Note that although the gatings are fixed during learning, changes in the weights affects how these gatings acts on the hidden layer activations, and it is interesting to understand how the learned task interacts with these gating operations. Since adding gatings at the input layer is equivalent to expanding the input dimension and replacing x_j by $x_j g_m(\mathbf{x})$, and learning does not affect how the gatings interact with the input, *we do not add gatings at the input layer for simplicity.*

Memory Capacity: Memory capacity refers to the number of random samples and labels for which there exists a set of parameters such that the network achieves zero training error (here we consider the mean squared error, MSE) and the capacity bounds of deep nonlinear networks has been extensively studied in many recent works (13; 27; 23). To calculate the capacity of GGDLNs, note that the input-output relation given by Eq.1 can be alternatively expressed as $f(\mathbf{x}) = \sum_{m_1,\cdots,m_L,j} W^{\text{eff}}_{m_1,\cdots,m_L,j} x^{\text{eff}}_{m_1,\cdots,m_L,j}$, which is a linear combination of the effective input $x^{\text{eff}}_{m_1,\cdots,m_L,j} = g_{m_1}(\mathbf{x})g_{m_2}(\mathbf{x})\cdots g_{m_L}(\mathbf{x})x_j$ ($m_l = 1,\cdots,M; j = 1,\cdots,N_0$), with some effective weights $\mathbf{W}^{\text{eff}}$ which is a complicated function of $\mathbf{a}$ and $\mathbf{W}$. Here m_l is the index of the gatings in the l-th layer. As the gating units are shared across layers, the effective input $\mathbf{x}^{\text{eff}}$ has $N_0\binom{M+L-1}{L}$ independent dimensions. This combinatorial term is the number of possible distinct products of L objects chosen from a pool of M different objects. Assuming $N \gg M^L$ such that the effective weight $W^{\text{eff}}_{m_1,\cdots,m_L,j}$ can take any desired value in the $N_0 M^L$ dimensional whole space, so the problem of finding $\mathbf{W}^{\text{eff}}$ with zero training error is equivalent to a linear regression problem with input $\mathbf{x}^{\text{eff}}$ and the target outputs. Therefore, the capacity is equivalent to the number of independent input dimensions, given by $P \leq N_0\binom{M+L-1}{L}$.

The above capacity is verified by Fig.1(c-e), showing that the behavior of generalization performance changes drastically at network capacity, where generalization error and its bias contribution diverge, and the variance contribution shrinks to 0 (see detailed calculation in the next paragraph and Appendix A.3). This double descent property of the generalization error is similar to previously studied in linear and nonlinear networks. Furthermore, although the output of the network is a linear function of the effective input $\mathbf{x}^{\text{eff}}$, due to the multiplicative nature of the network weights and the gatings, learning in GGDLNs is highly nonlinear and the space of solution for $\mathbf{W}$ and $\mathbf{a}$ is highly nontrivial, and the network exhibit properties unique to nonlinear networks, as we will show in the following sections.

Posterior distribution of network weights: We consider the posterior distribution of the network weights induced by learning with a Gaussian prior (11; 10; 30; 8), given by

$$P(\mathbf{\Theta}) = Z^{-1}\exp(-\frac{1}{2T}\sum_{\mu=1}^{P}(f(\mathbf{x}^\mu,\mathbf{\Theta}) - Y^\mu)^2 - \frac{1}{2\sigma^2}\mathbf{\Theta}^\top\mathbf{\Theta}) \tag{2}$$

where Z is the partition function $Z = \int d\mathbf{\Theta} P(\mathbf{\Theta})$. The first term in the exponent is the mean-squared deviation of the network outputs on a set of P training data points $\mathbf{x}^\mu$ from their target outputs Y^μ, and the second term is a Gaussian prior on the network parameters $\Theta = \{\mathbf{W}, \mathbf{a}\}$ with amplitude σ^{-2}. In this work we focus on the $T \to 0$ limit where the first data dependent term dominates. Below the network capacity, the distribution of Θ concentrates onto the space of solution that yields zero training error, the Gaussian prior then biases the solution space towards weights with smaller L_2 norms. The fundamental properties of the system can be derived from the partition function. As the distribution is quadratic in the readout weights $a_{m,i}$, it is straightforward to integrate them out, which yields

$$Z = \int d\mathbf{W}\exp[-\frac{1}{2\sigma^2}\text{Tr}(\mathbf{W}^\top\mathbf{W}) + \frac{1}{2}\mathbf{Y}^\top\mathbf{K}_L(\mathbf{W})^{-1}\mathbf{Y} + \frac{1}{2}\log\det(\mathbf{K}_L(\mathbf{W}))] \tag{3}$$

where $\mathbf{W}$ denotes all the remaining weights in the network, and $\mathbf{K}_L(\mathbf{W})$ is the $\mathbf{W}$ dependent $P \times P$ kernel on the training data, defined as $K^{\mu\nu}_L(\mathbf{W}) = (\frac{\sigma^2}{M} g(\mathbf{x}^\mu)^\top g(\mathbf{x}^\nu))(\frac{1}{N}\mathbf{x}^\mu_L(\mathbf{W})^\top \mathbf{x}^\nu_L(\mathbf{W}))$.

Generalization in infinitely wide GGDLNs: It is well known that in infinitely wide networks where $N \to \infty$ while P remains finite (also referred to as the GP limit), $\mathbf{K}_L(\mathbf{W})$ is self-averaging and does not depend on the specific realization of $\mathbf{W}$. It can therefore be replaced by the GP kernel defined as $\langle \mathbf{K}_L(\mathbf{W})\rangle_{\mathbf{W}}$ where $\mathbf{W} \sim \mathcal{N}(0,\sigma^2)$ (35). For GGDLNs, the GP kernel for a pair of arbitrary data x and $\mathbf{y}$ is given by $K_{GP}(\mathbf{x},\mathbf{y}) = (\frac{\sigma^2}{M}\mathbf{g}(\mathbf{x})^\top\mathbf{g}(\mathbf{y}))^L K_0(\mathbf{x},\mathbf{y})$, where $K_0(\mathbf{x},\mathbf{y}) = \frac{\sigma^2}{N_0}\mathbf{x}^\top\mathbf{y}$. We denote the $P \times P$ kernel data $matrix$ as $\mathbf{K}_{GP}$ where $K^{\mu\nu}_{GP} = K_{GP}(\mathbf{x}^\mu,\mathbf{x}^\nu)$, and the input kernel matrix on training data as $\mathbf{K}_0$ where $K^{\mu\nu}_0 = K_0(\mathbf{x}^\mu,\mathbf{x}^\nu)$.

Generalization error measured by MSE including the bias and the mean, $\epsilon_g = (\langle f(\mathbf{x})\rangle_{\mathbf{\Theta}} - y(\mathbf{x}))^2 + \langle \delta f(\mathbf{x})^2\rangle_\Theta$, which depends on the first and second order statistics of the predictor. In the GP limit, we have $\langle f(\mathbf{x})\rangle = \mathbf{k}_{GP}(\mathbf{x})^\top\mathbf{K}^{-1}_{GP}\mathbf{Y}$, and $\langle \delta f(\mathbf{x})^2\rangle = \mathbf{K}_{GP}(\mathbf{x},\mathbf{x}) - \mathbf{k}_{GP}(\mathbf{x})^\top\mathbf{K}^{-1}_{GP}\mathbf{Y}$, where $k^\mu_{GP}(\mathbf{x}) = K_{GP}(\mathbf{x},\mathbf{x}^\mu)$. Note that the rank of $\mathbf{K}_{GP}$ is the same as the capacity of the network, and the kernel matrix becomes singular as P approaches its capacity (the interpolation threshold), which

results in a diverging bias and vanishing variance (Fig.1 (b-d)). The singularity of the kernel at the interpolation threshold holds even for finite width networks, and similar diverging bias and vanishing variance are seen in our finite width theory below (Section 3) and are confirmed by simulation of networks trained with GD (see Appendix B.1).

3 Kernel shape renormalization theory in finite-width GLNs

We now address the finite width thermodynamic limit, where $P, N \to \infty$ but $P/N \sim \mathcal{O}(1)$, $M, L \sim \mathcal{O}(1)$. In this limit, calculating the statistics of the network predictor requires integration over $\mathbf{W}$ in Eq.3 . To do so, we apply the previous method of Back-propagating Kernel Renormalization (29) (see Appendix A) to GGDLNs. The partition function for a single hidden layer network is given by $Z = \exp(-H_1)$, where the Hamiltonian H_1 is given by

$$H_1 = \frac{1}{2}\mathbf{Y}^\top \tilde{\mathbf{K}}_1^{-1}\mathbf{Y} + \frac{1}{2}\log\det(\tilde{\mathbf{K}}_1) - \frac{N}{2}\log\det \mathbf{U}_1 + \frac{1}{2\sigma^2}N\mathrm{Tr}(\mathbf{U}_1) \tag{4}$$

$$\tilde{K}_1^{\mu\nu} = (\frac{1}{M}\mathbf{g}(\mathbf{x}^\mu)^\top \mathbf{U}_1 \mathbf{g}(\mathbf{x}^\nu))K_0^{\mu\nu} \tag{5}$$

Comparing the matrix $\tilde{\mathbf{K}}_1$to $\mathbf{K}_{GP}$, we note that the GP kernel is renormalized by an an $M \times M$ matrix order parameter $\mathbf{U}_1$. This order paramter satisfies the self-consistent equation

$$\mathbf{U}_1 = I - N^{-1}M^{-1}\mathbf{U}_1^{1/2}\mathbf{g}^\top[\tilde{\mathbf{K}}_1^{-1} \circ \mathbf{K}_0]\mathbf{g}\mathbf{U}_1^{1/2} + N^{-1}M^{-1}\mathbf{U}_1^{1/2}\mathbf{g}^\top[\tilde{\mathbf{K}}_1^{-1}\mathbf{Y}\mathbf{Y}^\top\tilde{\mathbf{K}}_1^{-1} \circ \mathbf{K}_0]\mathbf{g}\mathbf{U}_1^{1/2} \tag{6}$$

where $\circ$ denotes element-wise multiplication. In the linear case (which corresponds to $M = 1$), the GP kernel is renormalized by a scalar factor. In the $M > 1$ case, the effect of renormalization is more drastic as it changes that not only the amplitude but also the shape of the kernel. The renormalization matrix has an interesting physical interpretation that relates it to the readout weights $\mathbf{a}$ of GGDLNs,

$$U_1^{mn} = \langle \frac{1}{N}\sum_{i=1}^{N} a_{m,i}a_{n,i}\rangle \tag{7}$$

The calculation can be extended to multiple layers with a new order parameter introduced for each layer(see Appendix A).

The predictor statistics for a input $\mathbf{x}$ can be expressed in terms of the renormalized kernels, for an $L = 1$ network

$$\langle f(\mathbf{x})\rangle_{\boldsymbol{\Theta}} = \tilde{\mathbf{k}}_1(\mathbf{x})^\top\tilde{\mathbf{K}}_1^{-1}\mathbf{Y}, \langle \delta f(\mathbf{x})^2\rangle_{\boldsymbol{\Theta}} = \tilde{K}_1(\mathbf{x},\mathbf{x}) - \tilde{\mathbf{k}}_1(\mathbf{x})^\top\tilde{\mathbf{K}}_1^{-1}\tilde{\mathbf{k}}_1(\mathbf{x}) \tag{8}$$

where $\tilde{K}_1(\mathbf{x},\mathbf{y}) = (\frac{1}{M}\mathbf{g}(\mathbf{x})^\top\mathbf{U}_1\mathbf{g}(\mathbf{y}))K_0(\mathbf{x},\mathbf{y})$ denotes the renormalized kernel for two arbitrary inputs $\mathbf{x}$ and $\mathbf{y}$, $\tilde{\mathbf{K}}_1$ denotes the $P \times P$ renormalized kernel on the training data, and $\tilde{\mathbf{k}}_1(\mathbf{x})$ is the $Px1$ vector, $\tilde{k}_{1\mu}(\mathbf{x}) = K(\mathbf{x},\mathbf{x}_\mu)$. The kernel renormalization in GGDLNs changes the shape of the kernel through the data dependent $\mathbf{U}_1$, reflecting the nonlinear property of the network, and resulting in more complex behavior of predictor statistics relative to the linear networks, as shown in Section 4. Our theory describes the properties of the posterior distribution of the network weights induced at equilibrium by Langevin dynamics with the MSE cost function and the Gaussian prior (29; 7; 6; 5). Simulating this dynamics agrees remarkably well with the simulation (see Appendix B.2). We will compare our theoretical results with numerics of the gradient descent (GD) dynamics widely used in machine learning. We consider the case where the network is initialized with Gaussian i.i.d. weights with variance σ^2, and the mean and variance of the predictors are evaluated across multiple initializations (see Appendix C.4 for details). As we will show, our theory makes accurate qualitative predictions for GD dynamics in all examples throughout this paper.

4 Generalization

For linear networks the generalization error depends on N, σ^2 and L through the variance only, while the mean predictor always assumes the same value as the mean in the GP limit (29). This is because the scalar kernel renormalization of $\tilde{\mathbf{k}}_1(\mathbf{x})$ is cancelled out in the mean predictor by the renormalization of the inverse kernel $\tilde{\mathbf{K}}_1^{-1}$. In contrast, for GGDLNs the mean predictor and hence the error bias also change with these network parameters due to the matrix nature of the kernel renormalization(Eq.8) . Below we investigate in detail how matrix renormalization of the kernel affects the generalization behavior (especially the bias term) of the network.

4.1 Networks with single hidden layer

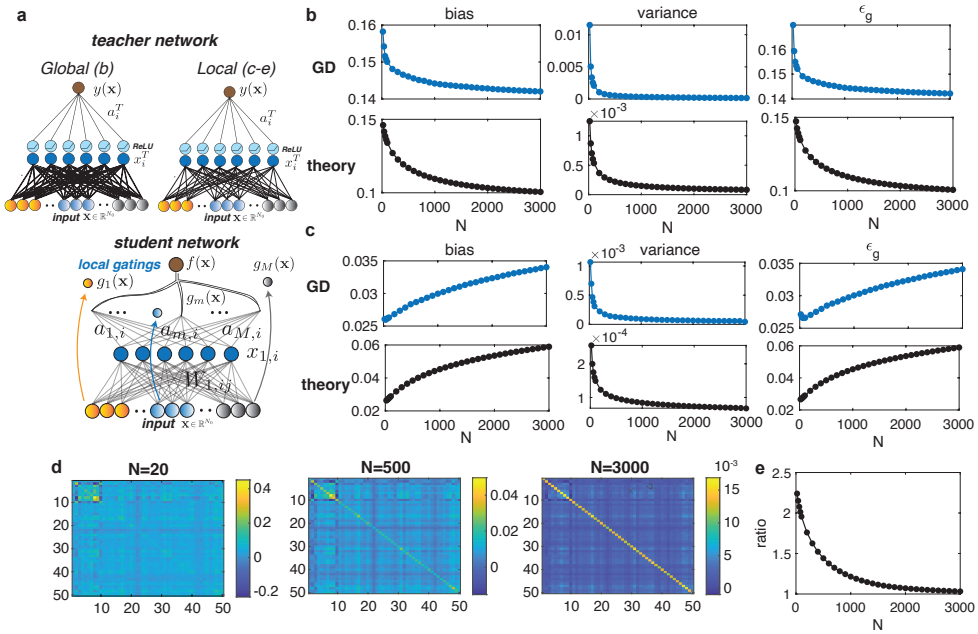

Figure 2: Dependence of generalization error on network width for a ReLU teacher task. (**a**)*Top:* The ReLU teacher network, the input $\mathbf{x}$ is divided into 5 subsets of input dimensions, the input layer weights either assume same order of magnitude across different input dimensions (left, (b)), or assume larger amplitudes for one subset of input dimensions-the preferred inputs (right, bold connections to a subset of input neurons,(c-e)). *Bottom:* The student network is a GGDLN with one hidden layer and gatings with localized receptive fields: each gating is connected to only a subset of input dimensions. (**b**) Bias, variance and generalization error decreases as a function of N for a regular ReLU teacher, theory agrees qualitatively well with GD dynamics. (**c**) Bias and generalization error increases as a function of N for ReLU teacher with preferred inputs. (**d**)The renormalization matrix $\mathbf{U}_1$ for different network widths for the teacher with preferred inputs. The first 10×10 block corresponds to the gatings with the same receptive field as the teacher's preferred inputs, and is amplified for small N. (**e**)The ratio of the average amplitude of the first 10×10 block relative to the average amplitude of the other four 10×10 diagonal blocks decreases as a function of N.

Dependence on network width N: Unlike in DLNs, the bias term in GGDLNs depends on N, exhibiting different dependence in different parameter regimes. This dependence also varies with the choice of the gating functions. In Fig.2we consider a student-teacher learning task (3; 2; 2; 9). We present results of learning a ReLU teacher task in GGDLNs with gatings that have *localized receptive fields* (Fig.2(a) bottom), where the student GGDLN is required to learn the input-output relation of a given ReLU teacher. For a ReLU teacher with a single fully connected hidden layer (Fig.2(a) top left), gatings with different receptive fields are of equal importance, hence the renormalization does not play a beneficial functional role, and the infinitely wide network performs better than finite N. As shown in Fig.2(b), bias, variance and generalization error all decrease with N. For a 'local' ReLU teacher with larger input weights for one subset of input components (the preferred inputs, Fig.2(a) top right), renormalization improves task performance by the selective increase of the elements of $\mathbf{U}_1$ that correspond to gating units whose receptive fields overlap the teacher preferred inputs (Fig.2(d&e)). Hence, narrower networks (with a stronger renormalization) generalize better, and both the bias and the generalization error increase with N (Fig.2(c)). Our theory not only captures the more complex behavior of generalization (especially bias) as a function of network width, but also provides qualitative explanation of how generalization is affected by the structure of the renormalization matrix in different tasks.

Dependence on weight regularization strength σ^2: Similar to the dependence on N, generalization also exhibits different behavior as a function of the regularization parameter σ in different parameter

regimes, with contributions from both the bias and the variance. The dependence of error bias on σ also arises due to the matrix nature of the renormalization. In Fig.3 , we show parameter regimes where the bias can increase (Fig.3 (a-c)) or decrease (Fig.3 (d-f)) with σ on MNIST (4) (Appendix C.2). Although the dependence on σ is complicated and diverse, and we lack a general rule for when the qualitative behavior changes, we found that our theory accurately captures the qualitative behavior of results obtained from GD (Appendix B.3 Fig.3). In both regimes the variance increases with σ as the solution space expands for a weaker regularization. Specifically in Fig.3 (d-f), due to the increasing variance (e) and decreasing bias (d), there is a minimum error rate ((f), Appendix A.3 Eq.49 for how error rate is calculated from the mean and variance of the predictor) at intermediate σ, indicating an optimal level of regularization strength as opposed to linear networks (29), where strong regularization ($\sigma = 0$) always results in optimal generalization

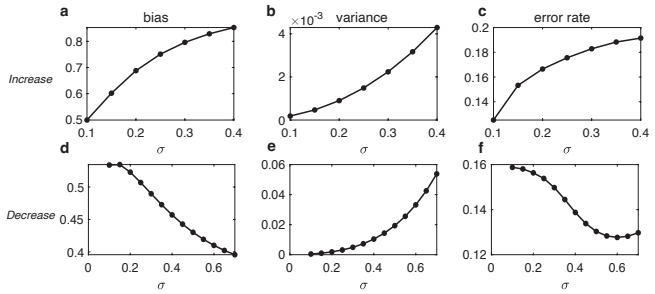

Figure 3: Generalization as a function of σ for GGDLN trained on MNIST dataset predicted by our theory. (**a-c**) Bias (a), variance (b) and error rate (c) increase as a function of σ. (**d-f**) Bias decreases as a function of σ while variance increases, leading to an optimal σ with minimum error rate.

Comparable performance to ReLU networks: The nonlinear operation of the gatings enables the network to learn nonlinear tasks. In Fig.4, we show that although the gatings are fixed during training, the network achieves comparable performance as a fully trained nonlinear (ReLU) network with the same hidden layer width for classifying even and odd digits in MNIST data when M is sufficiently large.

Furthermore, although the gatings are fixed during the supervised training of the GGDLN, they can be cleverly chosen to improve generalization performance. To demonstrate this strategy, we compared two different choices of gatings. *Random gatings* take the form $g_m(\mathbf{x}) = \Theta(\frac{1}{\sqrt{N_0}}\mathbf{V}_m^\top \mathbf{x} - b)$, where $\mathbf{V}_m$ is a $1 \times N_0$ random vector with standard Gaussian i.i.d. elements, b is a scalar threshold, and $\Theta(x)$ is the heaviside step function. The *pretrained gatings* are trained on the *unlabelled* training dataset with unsupervised soft k-means clustering, such that the m-th gating $g_m(x)$ outputs the probability of assigning data x to the m-th cluster (Appendix C.2). As shown in Fig.4, for pretrained gatings, generalization performance improves with M much faster compared to random gatings, and approaches the performance of ReLU network at a smaller M. Our theory (Fig.4) and numerical results of GD dynamics (Appendix B.3 Fig.4) agree qualitatively well. The result shows that GGDLNs can still achieve competitive performance on nonlinear tasks while remaining theoretically amenable.

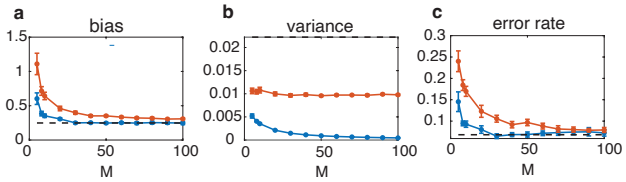

Figure 4: Dependence of generalization on M for GGDLNs trained on MNIST dataset predicted by our theory. Bias (a), variance (b) and error rate (c) as a function of M for random (red lines) and pretrained gatings (blue lines), and ReLU network with the same width.

4.2 Kernel shape renormalization in deeper networks

We now consider the effect of the matrix renormalization on GGDLN with more layers. We begin by analyzing the effect of $\mathbf{U}_1$ on the shape of the kernel in deep architectures. It is well known that the GP kernel of many nonlinear networks flattens as network depth increases (35), ultimately losing information about the input and degrading generalization performance. Here we show that kernel shape renormalization slows down flattening of kernels by incorporating data relevant information into the learned weights.

To study the *shape* of the kernel independent of kernel magnitude, we define the normalized kernel $K^L_{norm}(\mathbf{x},\mathbf{y}) = \frac{\tilde{K}_L(\mathbf{x},\mathbf{y})}{\tilde{K}_L(\mathbf{x},\mathbf{x})^{1/2}\tilde{K}_L(\mathbf{y},\mathbf{y})^{1/2}}$, where $\tilde{K}_L(\mathbf{x},\mathbf{y})$ denotes the renormalized kernel for GGDLN with L hidden layers. This normalized kernel measures the cosine of the vectors $\mathbf{x}$ and $\mathbf{y}$ with generalized inner product defined by the kernel $\tilde{K}_L(\mathbf{x},\mathbf{y})$, and therefore $K^L_{norm}(\mathbf{x},\mathbf{y}) \in [-1,1]$. For the GP kernel of GGDLNs, we have $K^L_{norm}(\mathbf{x},\mathbf{y}) = \cos(\mathbf{g}(\mathbf{x}),\mathbf{g}(\mathbf{y}))^L \cos(\mathbf{x},\mathbf{y})$. While K^L_{norm} depends on the specific choice of gatings in general, in the special case of *random gatings* with zero threshold $g_m(\mathbf{x}) = \Theta(\frac{1}{\sqrt{N_0}}\mathbf{V}_m^\top \mathbf{x})$ and the number of gatings $M \to \infty$, we can write K^L_{norm} analytically as a function of the angle θ between input vectors x and y, given by $K^L_{norm}(\theta) = (\frac{\pi-\theta}{\pi})^L \cos(\theta)$. Thus, $K^L_{norm}(\theta)$ converges to a δ-function as L goes to infinity, shrinking to zero except for $\theta = 0$. This 'flattening' effect also empirically holds true for networks with finite M (see Appendix B.4).

In Fig.5 ,we study the effect of kernel renormalization on the 'flattening' effect for deep networks trained to separate even and odd digits in MNIST dataset. As shown in Fig.5 (a)-(c), the elements of the renormalized kernel shrink to zero at a much slower rate compared to the GP kernel. (Note that unlike the variance, that bias is affected only by shape changes, but not by changes in the amplitude of the kernel, in Fig.5(d) we plot only the bias contribution to the generalization.) While mitigating the flattening of the GP kernel is a general feature of our renormalized kernel for different parameters, its effect on the generalization performance (especially the bias) may be different for different network parameters. In the specific example in Fig.B.4, finite width networks with a less 'flattened' renormalized kernel achieve better performance than the GP limit. Both the GP limit and the finite width networks have optimal performance at $L = 2$ in this example.

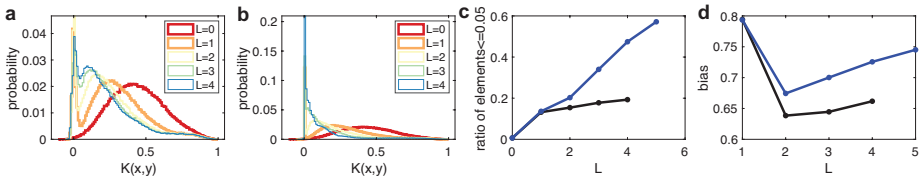

Figure 5: Shape renormalization slows down flattening of kernels in deep networks. (**a**-**b**) Distribution of kernel elements $K^L_{norm}(\mathbf{x},\mathbf{y})$ for the renormalized kernel (a) and GP kernel (b) for different network depth L. (**c**) Ratio of kernel elements smaller or equal to 0.05 increases faster for GP kernel (blue line) compared to the renormalized kernel (black line), the renormalization slows down the rate at which elements in the GP kernel shrink to zero as a function of L. (**d**) The bias contribution to the generalization first decreases then increases as a function of L due to the flattening of the kernel (blue line). Finite width network with renormalized kernel performs better for $L > 1$ in this parameter regime (black line). See Appendix C.2 for detailed parameters.

5 GGDLNs for multiple tasks

In this section, we apply our theory to investigate the ability of GGDLNs to perform multiple tasks. We consider two different scenarios below. First, different tasks require the network to learn input-output mappings on input data with different statistics. This scenario corresponds to real life situations where the training data distribution is non-stationary. The tasks can be separated based on difference in the inputs alone, without any additional top-down information. In this case, the gatings are bottom-up, and are functions of the input data only. In the second case, different tasks give conflicting labels for the same inputs, corresponding to the situation where performing the two tasks require additional top-down contextual information, and the information can be incorporated into the gating units in GGDLNs. In both scenarios, when the gatings are fixed and we modulate the de-correlation by changing network width and thus the strength of the kernel renormalization, we demonstrate that de-correlation between tasks always leads to better generalization performance.

5.1 Bottom-up gating units

First we consider learning different tasks defined by vastly different input statistics with bottom-up gatings, using permuted MNIST as an example. Previous works have shown that GLNs mitigate

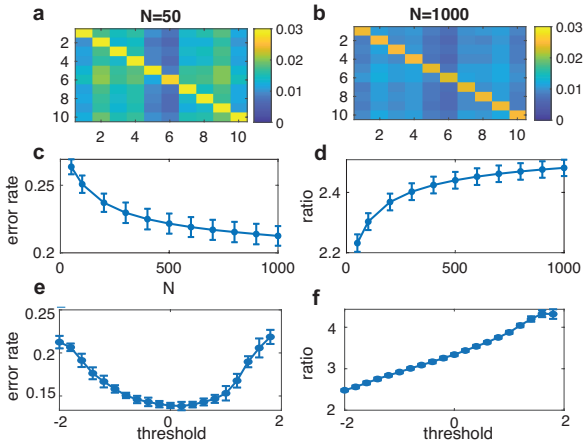

Figure 6: GGDLNs with bottom-up gating units learning multiple tasks trained on permuted MNIST. (**a**-**b**) Coefficient matrix C for $N = 50$ and $N = 1000$, different permutations are more decorrelated for larger N. (**c**) Error rate decreases as a function of N due to the decorrelation. (**d**) Ratio of the average amplitude of diagonal elements versus off-diagonal elements in C increases as a function of N. (**e**) Error rate first decreases then increases as a function of gating threshold. (**f**) Decorrelation increases as a function of gating threshold.

catastrophic forgetting when sequentially trained on permuted MNIST (28). While our theory does not address directly the dynamics of sequential learning, we aim to shed light on this question indirectly by asking how the two tasks interfere with each other when they are learned simultaneously.

We introduce a measure of *inter-task interference* by noting that after learning the mean predictor on a new data $\mathbf{x}$, Eq.8, is a linear combination of the output labels Y^μ of all the training data, and the coefficient of this linear combination, is given by the μ-th coefficient of $\tilde{\mathbf{k}}(\mathbf{x})^\top \tilde{\mathbf{K}}^{-1}$. Thus, we define a task-task correlation matrix, via $C_{pq} = \sum_{\mu=1}^{P} \sum_{\gamma=1}^{P_t} |\tilde{\mathbf{k}}^T \tilde{\mathbf{K}}^{-1}|_{p\gamma,q\mu} (p, q = 1, \cdots, n)$, where we assume there are P training examples and P_t test data for each task, with a total of n tasks. The amplitude of each element C_{pq} measures how much *training data of task* q contribute to the prediction on the *test data of task* p. Stronger diagonal elements indicates that the network separates the processing of data of different tasks ((Fig.6(a)-(b))). As we show in Fig.6, we can tune the relative strength of the diagonal elements of $\mathbf{C}$ smoothly by changing the network width (Fig.6(a)-(d)) or by changing the threshold of the gating (Fig.6 (e)-(f)). In the case where the gatings are fixed and the network width is changed, an increase in the strength of the diagonal elements (Fig.6(d)) results in better generalization (Fig.6(c)), indicating that the network generalizes better by processing data of different tasks more separately through the gating units. However, in the case where we change the activation of the gatings by adjusting the threshold, although different tasks are more de-correlated when the threshold is large due to a set of less overlapping gatings activated for each task, generalization error first decreases and then increases again. This is because for large threshold the sparsity of the gatings activated for each task limits the nonlinearity of the network, and therefore the generalization performance on this nonlinear task.

5.2 Combined top-down and bottom-up gating units

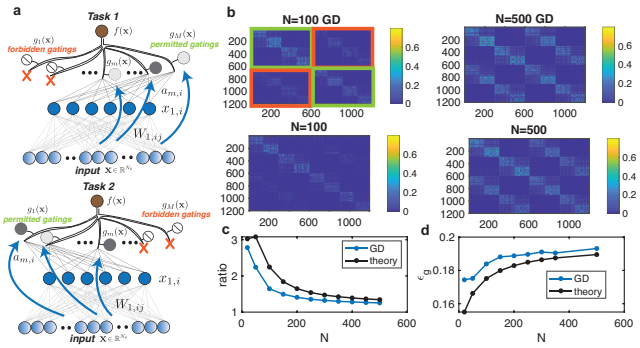

Figure 7: Kernel renormalization decorrelates different tasks defined by different labels on the same inputs. (**a**) GGDLNs performing two tasks using combined top-down and bottom-up task signal. (**b**) Top: Renormalized kernel calculated with Eq. 7 from GD dynamics. Bottom: Renormalized kernel theory. (**c**) Ratio of the magnitude of diagonal (red) versus off diagonal blocks (green) decreases as a function of (**d**) Generalization error increases with N.

We now consider learning two tasks that provide conflicting labels on the same input data. The gating units combine both top-down task specific signal which informs the system of which task to perform

for a given input, and bottom-up signals which, as before, depend on the input. In different tasks, different sets of gatings are *permitted or forbidden* depending on the top-down signal, then the states of the permitted gatings are further determined as a function of the input $\mathbf{x}$, while the forbidden gatings are set to 0 , meaning that the corresponding dendritic branches do not connect to the previous layer neurons (Fig.7(a)) in this task. For a single hidden layer network, with a similar argument as in Section 2, it is straightforward to show that the number of different tasks that can be memorized is given by $n \leq M$ and the number of training examples for each task needs to satisfy $P \leq N_0 M_p$, where M_p is the number of permitted gating units in each task. In the limiting case where a set of non-overlapping gating units are permitted in each of the n tasks, the network is equivalent to n sub-networks, each independently performing one task. In this case M_p is limited by M/n, which in turn limits the capacity and the effective input-output nonlinearity for each independent task. We consider the case where the permitted gatings are chosen randomly for each task and are therefore in general overlapping across tasks. We then investigate, through the data-dependent renormalization matrix, how learning modifies the correlation induced by the overlapping gatings. As an example we consider training on permuted and un-permuted MNIST digits of 0 and 1's. One task is to perform binary classification of the two digits in both permuted and un-permuted data, and the second task is to separate the permuted digits (both 0 and 1) from the un-permuted digits. The labels of the two tasks are entirely uncorrelated, while the permitted gatings of the two tasks are partially overlapping. In this case the renormalized kernel $\tilde{\mathbf{K}}_1$ can be written as $\tilde{K}_{p\mu,q\nu} = (\frac{1}{M}\mathbf{g}^p(\mathbf{x}^\mu)^\top \mathbf{U}_1 \mathbf{g}^q(\mathbf{x}^\nu))\frac{\sigma^2}{N_0}\mathbf{x}^{\mu\top}\mathbf{x}^\nu$. Here $p, q \in \{1, 2\}$ is the task index, and μ,$\nu = 1, \cdots, P$ are the input index. The kernel is therefore $2P \times 2P$ as shown in Fig.7(b); the diagonal blocks (green) correspond to kernels of task 1 and task 2, while the off diagonal blocks (red) correspond to the cross kernels.

In Fig.7(b) bottom, we show the renormalized kernel with the renormalization matrix $\mathbf{U}_1$ calculated by solving Eq.6. Similar results are achieved by by numerically estimating the readout weights from results obtained from GD dynamics by using Eq. 7 . The results demonstrate that stronger kernel renormalization achieved in narrower networks suppresses more strongly the correlation between tasks, reflected by the weaker off-diagonal blocks in Fig.7(b) A decreasing ratio between the average amplitudes of the diagonal and off-diagonal blocks show that the de-correlation effect diminishes for large N, leading to increasing generalization error with N(Fig.7(c & d))

6 Discussion

In this work, we proposed a novel gating network architecture, the GGDLN, amenable to rigorous theoretical analysis of the network expressivity and generalization performance. The predictor statistics of GGDLNs can be expressed in terms of kernels that undergo shape renormalization, resulting diverse behavior of the bias as a function of various network parameters. This renormalization slows down the flattening of the GP kernel in deep networks, suggesting that the loss of input information as L increases may be prevented in finite-width nonlinear networks. We also investigate the capability of GGDLNs to perform multiple tasks. While our theory is an exact description of the posterior of weight distribution induced by Langevin dynamics in Bayesian learning, it provides surprisingly well qualitative agreement with results obtained with GD dynamics for not only the generalization but also the kernel representation with matrix renormalization (Fig.7 (b)), largely extending its applicability.

While our theory currently addresses learning in GGDLNs using a *global* cost function, exploring the probability of extending the formalization of the equilibrium distribution to characterize local learning dynamics by using the properly normalized conditional probability of $\mathbf{W}_l$ on $\boldsymbol{W}_k (k < l)$ is an ongoing work. Recent works have shown that multilayer perceptrons (MLPs) with learned gatings that implements spatial attention have surprisingly good performance on Natural Language Processing (NLP) and computer vision (22). Extension of our theory to learnable gatings that implements attention mechanisms remains to be explored. Furthermore, our theory applies to fully connected linear networks with gating units, while other types of neural network architectures such as convolutional neural networks are commonly used for image processing (19; 17; 21; 18; 20). Incorporating convolutional architecture into our GGDLNs and using the gating units to encode context-dependent modification of different feature maps is an interesting direction related to the fast-developing research topic of visual question-answering (VQA) (16; 15; 14) , where answering different questions about the same image is similar to performing multiple tasks in different contexts with different labels on the same dataset, as we discussed in Section 5.2. We leave these exciting research directions for future work.

References

[1] Madhu S Advani, Andrew M Saxe, and Haim Sompolinsky. High-dimensional dynamics of generalization error in neural networks. *Neural Networks*, 132:428–446, 2020. 4.1, C.4

[2] Lin Wang and Kuk-Jin Yoon. Knowledge distillation and student-teacher learning for visual intelligence: A review and new outlooks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 4.1

[3] Shinji Watanabe, Takaaki Hori, Jonathan Le Roux, and John R Hershey. Student-teacher network learning with enhanced features. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5275–5279. IEEE, 2017. 4.1

[4] Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010. 1, 4.1

[5] H. Risken and H. Haken. *The Fokker-Planck Equation: Methods of Solution and Applications Second Edition*. Springer, 1989. 3

[6] Jean Zinn-Justin. *Quantum field theory and critical phenomena*, volume 171. Oxford university press, 2021. 3

[7] Gadi Naveh, Oded Ben David, Haim Sompolinsky, and Zohar Ringel. Predicting the outputs of finite deep neural networks trained with noisy gradients. *Physical Review E*, 104(6):064301, 2021. 3

[8] Andreas Engel and Christian Van den Broeck. *Statistical mechanics of learning*. Cambridge University Press, 2001. 2

[9] Hyunjune Sebastian Seung, Haim Sompolinsky, and Naftali Tishby. Statistical mechanics of learning from examples. *Physical review A*, 45(8):6056, 1992. 4.1

[10] Madhu Advani, Subhaneil Lahiri, and Surya Ganguli. Statistical mechanics of complex neural systems and high dimensional data. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(03):P03014, 2013. 2

[11] Daniel J Amit, Hanoch Gutfreund, and Haim Sompolinsky. Statistical mechanics of neural networks near saturation. *Annals of physics*, 173(1):30–67, 1987. 2

[12] Joel Veness, Tor Lattimore, Avishkar Bhoopchand, Agnieszka Grabska-Barwinska, Christopher Mattern, and Peter Toth. Online learning with gated linear networks. *arXiv preprint arXiv:1712.01897*, 2017. 1

[13] Jonathan Fiat, Eran Malach, and Shai Shalev-Shwartz. Decoupling gating from linearity. *arXiv preprint arXiv:1906.05032*, 2019. 1, 2

[14] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. *Advances in neural information processing systems*, 29, 2016. 6

[15] Kushal Kafle and Christopher Kanan. Visual question answering: Datasets, algorithms, and future challenges. *Computer Vision and Image Understanding*, 163:3–20, 2017. 6

[16] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015. 6

[17] Gadi Naveh and Zohar Ringel. A self consistent theory of gaussian processes captures feature learning effects in finite cnns. *Advances in Neural Information Processing Systems*, 34, 2021. 6

[18] Adrià Garriga-Alonso, Carl Edward Rasmussen, and Laurence Aitchison. Deep convolutional networks as shallow gaussian processes. *arXiv preprint arXiv:1808.05587*, 2018. 6

[19] Mark Van der Wilk, Carl Edward Rasmussen, and James Hensman. Convolutional gaussian processes. *Advances in Neural Information Processing Systems*, 30, 2017. 6

[20] Saad Albawi, Tareq Abed Mohammed, and Saad Al-Zawi. Understanding of a convolutional neural network. In *2017 international conference on engineering and technology (ICET)*, pages 1–6. Ieee, 2017. 6

[21] Jiuxiang Gu, Zhenhua Wang, Jason Kuen, Lianyang Ma, Amir Shahroudy, Bing Shuai, Ting Liu, Xingxing Wang, Gang Wang, Jianfei Cai, et al. Recent advances in convolutional neural networks. *Pattern Recognition*, 77:354–377, 2018. 6

[22] Hanxiao Liu, Zihang Dai, David So, and Quoc V Le. Pay attention to mlps. *Advances in Neural Information Processing Systems*, 34:9204–9215, 2021. 6

[23] Masami Yamasaki. The lower bound of the capacity for a neural network with multiple hidden layers. In *International Conference on Artificial Neural Networks*, pages 546–549. Springer, 1993. 2

[24] Samuel Lippl, LF Abbott, and SueYeon Chung. The implicit bias of gradient descent on generalized gated linear networks. *arXiv preprint arXiv:2202.02649*, 2022. 1

[25] David Budden, Adam Marblestone, Eren Sezener, Tor Lattimore, Gregory Wayne, and Joel Veness. Gaussian gated linear networks. *Advances in Neural Information Processing Systems*, 33:16508–16519, 2020. 1

[26] Joel Veness, Tor Lattimore, David Budden, Avishkar Bhoopchand, Christopher Mattern, Agnieszka Grabska-Barwinska, Eren Sezener, Jianan Wang, Peter Toth, Simon Schmitt, et al. Gated linear networks. *arXiv preprint arXiv:1910.01526*, 2019. 1

[27] Roman Vershynin. Memory capacity of neural networks with threshold and rectified linear unit activations. *SIAM Journal on Mathematics of Data Science*, 2(4):1004–1033, 2020. 2

[28] Eren Sezener, Agnieszka Grabska-Barwińska, Dimitar Kostadinov, Maxime Beau, Sanjukta Krishnagopal, David Budden, Marcus Hutter, Joel Veness, Matthew Botvinick, Claudia Clopath, et al. A rapid and efficient learning rule for biological neural circuits. *BioRxiv*, 2021. 1, 5.1

[29] Qianyi Li and Haim Sompolinsky. Statistical mechanics of deep linear neural networks: The backpropagating kernel renormalization. *Physical Review X*, 11(3):031059, 2021. 1, 3, 3, 4, 4.1

[30] Yasaman Bahri, Jonathan Kadmon, Jeffrey Pennington, Sam S Schoenholz, Jascha Sohl-Dickstein, and Surya Ganguli. Statistical mechanics of deep learning. *Annual Review of Condensed Matter Physics*, 11:501–528, 2020. 2

[31] Andrew K Lampinen and Surya Ganguli. An analytic theory of generalization dynamics and transfer learning in deep linear networks. *arXiv preprint arXiv:1809.10374*, 2018. 1

[32] Andrew M Saxe, James L McClelland, and Surya Ganguli. A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences*, 116(23):11537–11546, 2019. 1

[33] Youngmin Cho and Lawrence Saul. Kernel methods for deep learning. *Advances in neural information processing systems*, 22, 2009. 1, B.4

[34] Alexander G de G Matthews, Mark Rowland, Jiri Hron, Richard E Turner, and Zoubin Ghahramani. Gaussian process behaviour in wide deep neural networks. *arXiv preprint arXiv:1804.11271*, 2018. 1

[35] Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. Deep neural networks as gaussian processes. *arXiv preprint arXiv:1711.00165*, 2017. 1, 2, 4.2

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section **??**.
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 3 and 6
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 3
   (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix A
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] In supplementary material
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Appendix C
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Fig.6
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix C
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes] Appendix C
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]