
Why bigger is not always better: on finite and infinite neural networks

Laurence Aitchison¹

Abstract

Recent work has argued that neural networks can be understood theoretically by taking the number of channels to infinity, at which point the outputs become Gaussian process (GP) distributed. However, we note that infinite Bayesian neural networks lack a key facet of the behaviour of real neural networks: the fixed kernel, determined only by network hyperparameters, implies that they cannot do any form of representation learning. The lack of representation or equivalently kernel learning leads to less flexibility and hence worse performance, giving a potential explanation for the inferior performance of infinite networks observed in the literature (e.g. Novak et al. 2019). We give analytic results characterising the prior over representations and representation learning in finite deep linear networks. We show empirically that the representations in SOTA architectures such as ResNets trained with SGD are much closer to those suggested by our deep linear results than by the corresponding infinite network. This motivates the introduction of a new class of network: infinite networks with bottlenecks, which inherit the theoretical tractability of infinite networks while at the same time allowing representation learning.

One approach to understanding and improving neural networks is to perform Bayesian inference in an infinitely wide network (Lee et al., 2018; Matthews et al., 2018; Garriga-Alonso et al., 2019; Novak et al., 2019). In this limit the outputs become Gaussian process distributed, enabling efficient and exact reasoning about uncertainty, and giving a means of interpretation using the parameter-free kernel function (which depends only on network hyperparameters such as depth). However, the performance of Bayesian infinite networks lags considerably behind state-of-the-art finite

networks trained using SGD (e.g. compare performance in Garriga-Alonso et al. (2019), Novak et al. (2019) and Arora et al. (2019) against He et al. (2016) and Chen et al. (2018)). This seems surprising, because, to our knowledge, there are no reports of wider networks degrading classification performance (indeed, the opposite is sometimes argued; see Zagoruyko & Komodakis, 2016), and because exact Bayesian inference is provably optimal, if the prior accurately describes our beliefs (Ramsey, 1926). Indeed, recent work on the Neural Tangent Kernel (NTK) (Li et al., 2019) has suggested that deterministic gradient descent in an infinite network gives slightly lower performance than Bayesian inference in the same network.

Our hypothesis is that the poor performance of Bayesian infinite networks arises because the top-layer representation (equivalent to the kernel), is fixed by the network hyperparameters, and thus cannot be learned from data. This breaks many of our key intuitions about why deep networks are effective. For instance in transfer learning (Huh et al., 2016) we use a large-scale dataset such as ImageNet to learn a good high-level representation, then apply this representation to other tasks where less data is available. However, transfer learning is impossible in infinite Bayesian neural networks, because the top-layer representation is fixed by the network hyperparameters and so cannot be learned using e.g. ImageNet.

To understand these issues, we analysed finite networks using tools from the infinite network literature (Lee et al., 2018; Matthews et al., 2018; Garriga-Alonso et al., 2019; Novak et al., 2019). We begin by giving a toy, two-layer example, contrasting the flexibility of finite networks with the inflexibility of infinite networks, showing that flexible finite networks offer benefits under conditions of model-mismatch. We then introduce infinite networks with bottlenecks, which combine the theoretical tractability of infinite networks with the flexibility of finite networks. To obtain an analytic understanding of kernel/representation flexibility and learning in such networks, we consider linear infinite networks with bottlenecks, which are equivalent to finite deep linear networks. We took two approaches to characterising these networks. First, we considered the prior viewpoint, i.e. the covariance in the top-layer kernel induced by randomness in the lower-layer weights. In particular, we showed that narrower, deeper networks offer more flexibility, and that

¹University of Bristol, Bristol, UK. Correspondence to: Laurence Aitchison <laurence.aitchison@gmail.com>.

CNNs offer more flexibility than locally connected networks (LCNs) when the input is spatially structured. Second, we considered the posterior viewpoint, showing that under both MAP inference and posterior sampling, the representations in learned neural networks slowly transition from being similar to the input kernel (i.e. the inner product of the inputs) to being similar to the output kernel (i.e. the inner product of one-hot vectors representing the labels). We found an important difference between MAP inference and sampling: for MAP inference, the learned representations transition from the input to the output kernel, irrespective of the network width. Bayesian networks behave similarly when the network width and the number of output channels are equal, but as the network width increases, the learned representations become increasingly dominated by the prior, and insensitive to the outputs. Remarkably, we find that in a ResNet trained using SGD on CIFAR-10, the representation differs dramatically from the corresponding infinite network and is instead very close to the output kernel, as suggested by our deep linear results. This confirms the importance of working with a theoretical model, such as infinite networks with bottlenecks, that is capable of capturing representation learning.

1. Toy Example

In the introduction, we noted that infinite Bayesian networks perform worse than standard neural networks trained using stochastic gradient descent. Thus, as we make finite neural networks wider, there should be some point at which performance begins to degrade. We considered a simple, two-layer, fully-connected linear network with the full set of 20 4-dimensional inputs denoted \mathbf{X} , hidden unit activations denoted \mathbf{H} , and 10-dimensional outputs denoted \mathbf{Y} ,

$$\mathbf{H} = \mathbf{X}\mathbf{W} \quad \mathbf{Y} = \mathbf{H}\mathbf{V} + \sigma\boldsymbol{\Xi} \quad (1)$$

where $\boldsymbol{\Xi}$ is IID standard Gaussian noise, \mathbf{W} is the input-to-hidden weight matrix and \mathbf{V} is the hidden-to-output weight matrix, whose columns, \mathbf{w}_μ and \mathbf{v}_ν are generated IID from,

$$P(\mathbf{w}_\mu) = \mathcal{N}(\mathbf{w}_\mu; \mathbf{0}, \frac{1}{X}\mathbf{I}) \quad P(\mathbf{v}_\nu) = \mathcal{N}(\mathbf{v}_\nu; \mathbf{0}, \frac{1}{H}\mathbf{I}), \quad (2)$$

and where the variance of the weights is normalised by the number of inputs to that layer, $X = 4$ for the 4-dimensional input, and H for the width of the hidden layer.

In the first example (Fig. 1 left), we generated targets for supervised learning using a second neural network with weights generated as described above, with $H_{\text{gen}} \in \{1, 2, 4\}$ hidden units. We evaluated the Bayesian model-evidence for networks with many different numbers of hidden units (x-axis). Bayesian reasoning would suggest that the model evidence for the true model (i.e. with a matched number of hidden units) should be higher than the model evidence

for any other model, as indeed we found (Fig. 1 top left), and these patterns held true for the predictive probability, or equivalently test performance (Fig. 1 bottom left). While these results give an example where smaller networks perform better, they do not necessarily help us to understand the behaviour of neural networks on real datasets, where the true generative process for the data is not known, and is not in our model class. As such, we considered two further examples where the neural network generating the targets lay outside of our model class. In particular, we again generated target outputs by sampling a “true” network from the prior, but we modified the inputs to this network, first by multiplying those inputs by 100 (Fig. 1 middle), then by zeroing-out all but the first input unit (Fig. 1 right). Critically, we ensured model-mismatch by putting the original, unmodified inputs into the trained networks. In both of these experiments, there was an optimum number of hidden units, after which performance degraded as more hidden units were included.

To understand why this might be the case, it is insightful to consider the methods we used to evaluate the model evidence and generate these results. In particular, note that conditioned on \mathbf{H} , the output for any given channel, \mathbf{y}_ν , is IID and depends only on the corresponding column of the output weights, \mathbf{v}_ν ,

$$\begin{aligned} P(\mathbf{Y}|\mathbf{V}, \mathbf{H}) &= \prod_{\nu} P(\mathbf{y}_\nu|\mathbf{v}_\nu, \mathbf{H}) \\ &= \prod_{\nu} \mathcal{N}(\mathbf{y}_\nu; \mathbf{H}\mathbf{v}_\nu, \sigma^2\mathbf{I}). \end{aligned} \quad (3)$$

Thus, we can integrate over the output weights, \mathbf{v}_ν , to obtain a distribution over \mathbf{Y} conditioned on \mathbf{H} ,

$$\begin{aligned} P(\mathbf{Y}|\mathbf{H}) &= \prod_{\nu} P(\mathbf{y}_\nu|\mathbf{H}) \\ &= \prod_{\nu} \mathcal{N}(\mathbf{y}_\nu; \mathbf{0}, \frac{1}{H}\mathbf{H}\mathbf{H}^T + \sigma^2\mathbf{I}). \end{aligned} \quad (4)$$

This is the classical Gaussian process representation of Bayesian linear regression (Rasmussen & Williams, 2006). Remembering that the hidden activities, \mathbf{H} , is a deterministic function of the weights, \mathbf{W} , and inputs, \mathbf{X} , we can write this distribution as,

$$\begin{aligned} P(\mathbf{y}_\nu|\mathbf{H}) &= P(\mathbf{y}_\nu|\mathbf{W}, \mathbf{X}) \\ &= \mathcal{N}(\mathbf{y}_\nu; \mathbf{0}, \frac{1}{H}\mathbf{X}\mathbf{W}\mathbf{W}^T\mathbf{X}^T + \sigma^2\mathbf{I}). \end{aligned} \quad (5)$$

Thus, the first-layer weights, \mathbf{W} , act as kernel hyperparameters in a Gaussian process: they control the covariance of the outputs, \mathbf{y}_ν . To evaluate the model evidence we need to

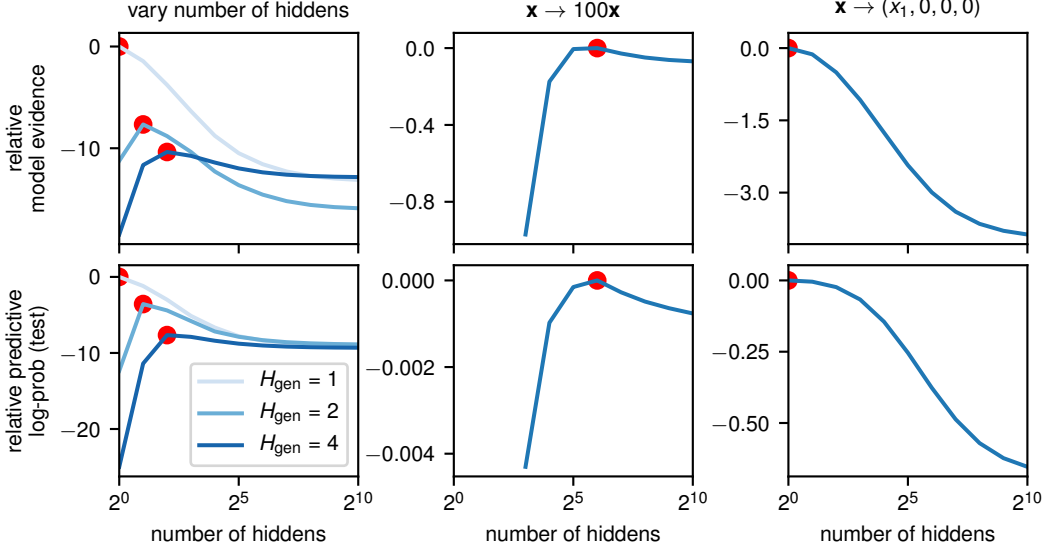


Figure 1. A toy fully-connected, two-layer Bayesian linear network showing situations in which smaller networks perform better than larger networks. The red dots indicate the optimal number of hidden units in that simulation. Left: training data generated from the prior with H_{gen} hidden units. Middle: training data generated from the prior with $H_{\text{gen}} = 4$ but where we scale-up the inputs by a factor of 100. Right: training data generated from the prior with $H_{\text{gen}} = 4$, but where we zero-out all but the first input dimension. Top: Bayesian model evidence. Bottom: predictive log-probability, or equivalently test-error.

integrate over \mathbf{W} ,

$$\begin{aligned} P(\mathbf{Y}|\mathbf{X}) &= \int d\mathbf{W} P(\mathbf{W}) \prod_{\nu} P(y_{\nu}|\mathbf{W}, \mathbf{X}) \\ &= \mathbb{E}_{P(\mathbf{W})} \left[\prod_{\nu} P(y_{\nu}|\mathbf{W}, \mathbf{X}) \right] \end{aligned} \quad (6)$$

and we estimate this integral by drawing 64 000 samples from the prior, $P(\mathbf{W})$. Importantly, while \mathbf{W} provides flexibility in the kernel in finite networks, this flexibility gradually disappears as we consider wider hidden layers networks. In particular,

$$\begin{aligned} \lim_{H \rightarrow \infty} \frac{1}{H} \mathbf{W}\mathbf{W}^T &= \lim_{H \rightarrow \infty} \frac{1}{H} \sum_{\mu=1}^H \mathbf{w}_{\mu} \mathbf{w}_{\mu}^T \\ &= \mathbb{E} [\mathbf{w}_{\mu} \mathbf{w}_{\mu}^T] = \frac{1}{X} \mathbf{I}. \end{aligned} \quad (7)$$

Therefore, in this limit, the distribution over \mathbf{Y} converges to,

$$\lim_{H \rightarrow \infty} P(\mathbf{Y}|\mathbf{X}) = \prod_{\nu} \mathcal{N}(y_{\nu}; \mathbf{0}, \frac{1}{X} \mathbf{X}\mathbf{X}^T + \sigma^2 \mathbf{I}). \quad (8)$$

This is exactly the distribution we would expect from Bayesian linear regression in a one-layer network. Thus, by taking the infinite limit, we have eliminated the additional flexibility afforded by the two-layer network, and we can see that the superior performance of smaller networks in Fig. 1 emerges because they give additional flexibility in

the covariance of the outputs, which gradually disappears as network size increases. Finally, note that sampling from the prior works well here both because of the concentration result above, and because we use relatively small amount of data, 20 points.

2. Infinite networks with finite bottlenecks

In the previous section, we considered the simplest networks in which these phenomena emerge: a two-layer, linear network. In this section, we setup a full infinite network with bottlenecks and show that activity flowing through this network can be understood entirely in terms of kernel and covariance matrices.

Consider a single layer within a fully-connected network, where the potentially infinite activity at the previous layer, $\mathbf{H}_{\ell-1}$, corresponding to a batch containing all inputs, is multiplied by a weight matrix, \mathbf{W}_{ℓ} , to give a finite number of activations, \mathbf{A}_{ℓ} . This activation matrix, \mathbf{A}_{ℓ} is multiplied by another matrix, \mathbf{M}_{ℓ} , to give a potentially infinite updated activation matrix, \mathbf{A}'_{ℓ} , which is then passed through a non-linearity, ϕ , to give the potentially infinite activity at this layer, \mathbf{H}_{ℓ} . Note that following Matthews et al. (2018), we use “activation” pre-nonlinearity and “activity” post-nonlinearity.

$$\mathbf{A}_{\ell} = \mathbf{H}_{\ell-1} \mathbf{W}_{\ell} \quad \mathbf{A}'_{\ell} = \mathbf{A}_{\ell} \mathbf{M}_{\ell} \quad \mathbf{H}_{\ell} = \phi(\mathbf{A}'_{\ell}) \quad (9)$$




Figure 2. The relationships between the feature-space and kernel representations of the neural network. For a typical finite neural network, $M_\ell = \mathbf{I}$, so $M_\ell = N_\ell$. For a finite-infinite network (which allows us to compute \mathbf{L}_ℓ from \mathbf{K}_ℓ), we send $M_\ell \rightarrow \infty$, and draw the elements of \mathbf{M}_ℓ IID from a Gaussian distribution with zero mean and variance $1/M_\ell$.

where the input data is $\mathbf{H}_0 = \mathbf{X}$, and

$$\begin{aligned} \mathbf{H}_\ell &\in \mathbb{R}^{P \times M_\ell} & \mathbf{A}_\ell &\in \mathbb{R}^{P \times N_\ell} & \mathbf{A}'_\ell &\in \mathbb{R}^{P \times M_\ell} \\ \mathbf{W}_\ell &\in \mathbb{R}^{M_{\ell-1} \times N_\ell} & \mathbf{M}_\ell &\in \mathbb{R}^{N_\ell \times M_\ell} \end{aligned} \quad (10)$$

For an infinite network with bottlenecks, we take the limit as M_ℓ goes to infinity, leaving N_ℓ finite. As such, the activity before, \mathbf{A}'_ℓ , and after, \mathbf{H}_ℓ , the nonlinearity is infinite, with a finite linear bottleneck formed by \mathbf{A}_ℓ .

For a fully-connected network, the columns of \mathbf{W}_ℓ and \mathbf{M}_ℓ , denoted \mathbf{w}_λ^ℓ and \mathbf{m}_λ^ℓ are generated IID from a Gaussian distribution,

$$\mathbb{P}(\mathbf{W}_\ell) = \prod_{\lambda=1}^{N_\ell} \mathbb{P}(\mathbf{w}_\lambda^\ell) = \prod_{\lambda=1}^{N_\ell} \mathcal{N}\left(\mathbf{w}_\lambda^\ell; \mathbf{0}, \frac{1}{M_{\ell-1}} \mathbf{I}\right) \quad (11)$$

$$\mathbb{P}(\mathbf{M}_\ell) = \prod_{\lambda=1}^{M_\ell} \mathbb{P}(\mathbf{m}_\lambda^\ell) = \prod_{\lambda=1}^{M_\ell} \mathcal{N}\left(\mathbf{m}_\lambda^\ell; \mathbf{0}, \frac{1}{N_\ell} \mathbf{I}\right). \quad (12)$$

where the normalization constants, $1/M_{\ell-1}$ and $1/N_{\ell-1}$, ensure that activations remain normalized as they flow through the network.

Following the infinite network literature, we would like to characterise activity flowing through the network in terms of the activation kernel, \mathbf{K}_ℓ and activity kernel, \mathbf{L}_ℓ ,

$$\begin{aligned} \mathbf{K}_\ell &\equiv \frac{1}{N_\ell} \mathbf{A}_\ell \mathbf{A}_\ell^T \\ \mathbf{L}_\ell &\equiv \frac{1}{M_\ell} \mathbf{H}_\ell \mathbf{H}_\ell^T & \mathbf{L}_0 &= \frac{1}{M_0} \mathbf{X} \mathbf{X}^T \end{aligned} \quad (13)$$

We begin by characterising the relationship between \mathbf{A}'_ℓ and \mathbf{A}_ℓ . As each channel (column) of \mathbf{A}'_ℓ is a linear function of the corresponding channel of the weights, $\mathbf{a}'_\lambda{}^\ell = \mathbf{A}_\ell \mathbf{m}_\lambda^\ell$, these activations are Gaussian and IID conditioned on \mathbf{A}_ℓ ,

$$\begin{aligned} \mathbb{P}(\mathbf{A}'_\ell | \mathbf{A}_\ell) &= \prod_{\lambda=1}^{M_\ell} \mathbb{P}(\mathbf{a}'_\lambda{}^\ell | \mathbf{A}_\ell) \\ &= \prod_{\lambda=1}^{M_\ell} \mathcal{N}(\mathbf{a}'_\lambda{}^\ell; \mathbf{0}, \mathbf{K}_\ell) = \mathbb{P}(\mathbf{A}'_\ell | \mathbf{K}_\ell) \end{aligned} \quad (14)$$

and taking the limit of $M_\ell \rightarrow \infty$,

$$\begin{aligned} \lim_{M_\ell \rightarrow \infty} \frac{1}{M_\ell} \mathbf{M}_\ell \mathbf{M}_\ell^T &= \frac{1}{N_\ell} \mathbf{I} \\ \lim_{M_\ell \rightarrow \infty} \frac{1}{M_\ell} \mathbf{A}'_\ell (\mathbf{A}'_\ell)^T &= \lim_{M_\ell \rightarrow \infty} \frac{1}{M_\ell} \mathbf{A}_\ell \mathbf{M}_\ell \mathbf{M}_\ell^T \mathbf{A}_\ell^T \\ &= \frac{1}{N_\ell} \mathbf{A}_\ell \mathbf{A}_\ell^T = \mathbf{K}_\ell, \end{aligned} \quad (15)$$

Thus, the kernel for \mathbf{A}_ℓ is equivalent to the kernel for \mathbf{A}'_ℓ in infinite networks with finite bottlenecks (Fig. 2).

Next, consider computing \mathbf{K}_ℓ from $\mathbf{L}_{\ell-1}$. As each channel (column) of the activations is a linear function of the corresponding channel of the weights, $\mathbf{a}'_\lambda{}^\ell = \mathbf{H}_{\ell-1} \mathbf{w}_\lambda^\ell$, the activations are Gaussian and IID conditioned on the activity at the previous layer,

$$\begin{aligned} \mathbb{P}(\mathbf{A}_\ell | \mathbf{H}_{\ell-1}) &= \prod_{\mu=1}^{N_\ell} \mathbb{P}(\mathbf{a}'_\mu{}^\ell | \mathbf{H}_{\ell-1}) \\ &= \prod_{\mu=1}^{N_\ell} \mathcal{N}(\mathbf{a}'_\mu{}^\ell; \mathbf{0}, \mathbf{J}_\ell) = \mathbb{P}(\mathbf{A}_\ell | \mathbf{J}_\ell), \end{aligned} \quad (16)$$

with covariance \mathbf{J}_ℓ . For a fully connected network, the covariance, \mathbf{J}_ℓ , is equal to the previous layer's activity-kernel, \mathbf{L}_ℓ ,

$$\mathbf{J}_\ell = \mathbf{L}_{\ell-1} = \frac{1}{M_{\ell-1}} \mathbf{H}_{\ell-1} \mathbf{H}_{\ell-1}^T, \quad (17)$$

but the relationship is more complex in convolutional architectures (Garriga-Alonso et al., 2019; Novak et al., 2019) (Appendix A.2). As \mathbf{A}_ℓ is always finite and random, \mathbf{K}_ℓ is also a random variable, and inspecting the above expressions, its distribution can be written as a Wishart, centered on $\mathbf{L}_{\ell-1}$.

Finally consider computing \mathbf{L}_ℓ from \mathbf{K}_ℓ . Note that as both \mathbf{A}'_ℓ and \mathbf{H}_ℓ are infinite, we can directly use standard results from infinite neural networks, i.e. those from Cho & Saul (2009), as in Lee et al. (2018); Matthews et al. (2018); Garriga-Alonso et al. (2019); Novak et al. (2019).

Linear infinite networks with finite bottlenecks can be obtained by setting $\mathbf{H}_\ell = \phi(\mathbf{A}'_\ell) = \mathbf{A}'_\ell$, implying that

$\mathbf{L}_\ell = \mathbf{K}_\ell$. Critically, this is equivalent to a deep linear network obtained by in addition setting $\mathbf{M}_\ell = \mathbf{I}$ so that $\mathbf{A}'_\ell = \mathbf{A}_\ell$ and $M_\ell = N_\ell$, as these choices imply that the $\mathbf{A}_\ell = \mathbf{A}'_\ell = \mathbf{H}_\ell$ so that again, $\mathbf{L}_\ell = \mathbf{K}_\ell$.

2.1. DNNs are deep GPs

Given this setup, we can see that even a finite nonlinear network (i.e. with $\mathbf{M}_\ell = \mathbf{I}$) is a deep Gaussian process. In particular, in a deep Gaussian process, the activations at layer ℓ , denoted \mathbf{A}_ℓ , consist of N_ℓ IID channels that are Gaussian-process distributed (Eq. 16), with a kernel/covariance determined by the activations at the previous layer. For a fully connected network,

$$\mathbf{J}_\ell = \mathbf{L}_{\ell-1} = \frac{1}{N_{\ell-1}} \phi(\mathbf{A}_{\ell-1}) \phi^T(\mathbf{A}_{\ell-1}). \quad (18)$$

The relationship between finite neural networks and deep GPs is worth noting, because the same intuition, of the lower-layers shaping the top-layer kernel, arises in both scenarios (e.g. Bui et al., 2016), and because there is potential for applying GP inference methods for neural networks, and vice versa.

3. The prior view on kernel flexibility

We can analyse how flexibility in the kernel emerges by looking at the variability (i.e. the variance and covariance) of \mathbf{J}_ℓ , \mathbf{K}_ℓ and \mathbf{L}_ℓ . If the prior gives a stochastic kernel with higher variance, then it will be easier to shape that kernel by conditioning on data. In the appendix, we derive recursive updates for deep, linear, convolutional networks, but here, for simplicity we give the fully-connected updates,

$$\mathbb{C}[J_{ij}^\ell, J_{kl}^\ell | \mathbf{L}^0] = \mathbb{C}[L_{ij}^{\ell-1}, L_{kl}^{\ell-1} | \mathbf{K}^0] \quad (19a)$$

$$\begin{aligned} \mathbb{C}[K_{ij}^\ell, K_{kl}^\ell | \mathbf{L}^0] &\approx \mathbb{C}[J_{ij}^\ell, J_{kl}^\ell | \mathbf{L}^0] \\ &+ \frac{1}{N_\ell} (\langle J_{ik}^\ell \rangle \langle J_{jl}^\ell \rangle + \langle J_{il}^\ell \rangle \langle J_{jk}^\ell \rangle) \end{aligned} \quad (19b)$$

$$\mathbb{C}[L_{ij}^\ell, L_{kl}^\ell | \mathbf{L}^0] = \mathbb{C}[K_{ij}^\ell, K_{kl}^\ell | \mathbf{L}^0] \quad (19c)$$

where,

$$\langle J_{ij}^\ell \rangle = \mathbb{E}[J_{ij}^\ell | \mathbf{L}^0] = L_{ij}^0 \quad (19d)$$

where i, j, k and l index datapoints.

This expression predicts that the variance of the kernel is proportional to the depth (including the last layer; $L + 1$) and inversely proportional to the width, N ,

$$\mathbb{C}[K_{ij}^{L+1}, K_{kl}^{L+1} | \mathbf{L}^0] \approx \frac{L+1}{N} (L_{ik}^0 L_{jl}^0 + L_{il}^0 L_{jk}^0). \quad (20)$$

This expression is so simple because, for a fully-connected linear network, the expected covariance at each layer is the same. For nonlinear and convolutional or locally-connected networks the covariance is still proportional to $1/N$, but the

depth-dependence becomes more complex, as the covariance changes as it propagates through layers.

To check the validity of these expressions, we sampled 10,000 neural networks from the prior, and evaluated the variance of the kernel for a single input (Fig. 3). These inputs were either spatially unstructured (i.e. white noise), or spatially structured, in which case the inputs were the same across the whole image. For fully connected networks, we confirmed that the variance of the kernel is proportional to the depth including the last layer, $L + 1$, and inversely proportional to width, N (Fig. 3A). For locally connected networks, we found that structured and unstructured inputs gave the same kernel variance, which is expected as any spatial structure is destroyed after the first layer (Fig. 3B). Further, for convolutional networks with structured input, the variance of the kernel was proportional to network depth (Fig. 3C bottom), but whenever that spatial structure was absent, either because it was absent in the inputs or because it was eliminated by an LCN (Fig. 3BC bottom) the variance of the kernel was almost constant with depth (see Appendix A.2.1).

The large decrease in kernel flexibility for locally connected networks might be one reason behind the result in Novak et al. (2019) that locally connected networks have performance that is very similar to an infinite-width network, in which all flexibility has been eliminated. In essence, for a locally connected network, we sample the weights for each spatial region independently, so we in effect average over more IID random variables, reducing the variance of the kernel at the next layer, and hence reducing the possibility for data to shape that representation. In contrast, for a convolutional network we share weights across locations, increasing the variance in the kernel, and hence increasing the possibility for data to shape the representation. Finally, as the spatial input size, S , increases, for convolutional networks with spatially structured inputs, the variance of the kernel is constant, whereas for locally connected or spatially unstructured inputs, the variance falls (Fig. 3D).

4. The posterior view on kernel flexibility

An alternative approach to understanding flexibility in finite neural networks is to consider the posterior viewpoint: how learning shapes top-level representations. To obtain analytical insights, we considered maximum a-posteriori and sampling based inference in a deep, fully-connected, linear network. In both cases, we found that learned neural networks shift the representation from being close to the input kernel, defined by,

$$\mathbf{K}_0 = \mathbf{L}_0 = \frac{1}{M_0} \mathbf{X}\mathbf{X}^T, \quad (21)$$

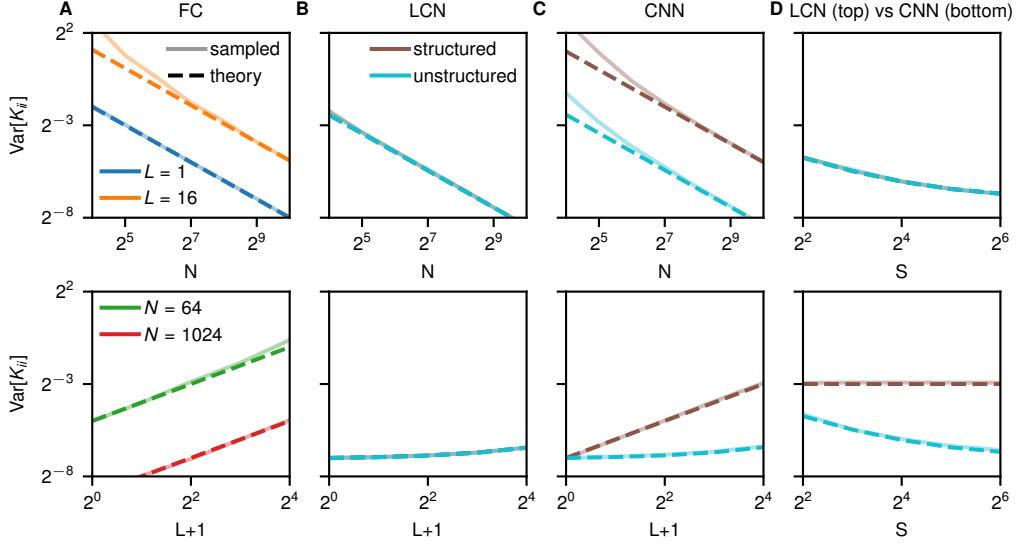


Figure 3. The variance of the stochastic kernel induced by randomly sampling weights in finite, linear, fully connected and convolutional networks, with spatially structured and unstructured inputs. We use normalized inputs and circular convolutions to ensure that the kernel’s expected value remains equal to 1 at all locations as it propagates through the network. The dashed lines in all plots display the theoretical approximation (Eq. 19) which is valid when the width is much greater than the number of layers. The solid lines display the empirical variance of the kernel from 10,000 simulations. **A** The variance of the kernel for fully connected networks, plotted against network width, N , for shallow (blue; $L + 1 = 1$), and deep (orange; $L + 1 = 16$) networks (top) and plotted against network depth, $L + 1$, for narrow (green; $N = 64$) and wide (red; $N = 1024$) networks. **B** The variance of the kernel for locally connected networks with spatially structured and unstructured inputs, plotted against the number of channels, N , and against network depth, $L + 1$. Note that the structured line lies underneath the unstructured line. The inputs are 1-dimensional with $S = 32$ spatial locations, and 100 input channels. **C** As in **B**, but for convolutional networks. **D** The variance of the kernel as a function of the input spatial size, S , for deep ($L + 1 = 16$) LCNs (top) and CNNs (bottom) with spatially structured and unstructured inputs.

to being close the output kernel, defined by,

$$\mathbf{K}_{L+1} = \frac{1}{N_{L+1}} \mathbf{Y} \mathbf{Y}^T. \quad (22)$$

In particular, under MAP inference, the shape of the kernel smoothly transitions from the input to the output kernel (Appendix B.2),

$$\mathbf{K}_\ell = \left(\frac{N_{\ell <}}{N_{\leq \ell}} \right)^{\frac{\ell(L+1-\ell)}{L+1}} (\mathbf{K}_{L+1} \mathbf{K}_0^{-1})^{\ell/(L+1)} \mathbf{K}_0, \quad (23)$$

where $N_{\ell <}$ is the geometric average of the width in layers $\ell + 1$ to $L + 1$, and $N_{\leq \ell}$ is the geometric average of the width in layers 1 to ℓ . Thus, the kernels (and the underlying weights) at each layer can be made arbitrarily large or small by changing the width, despite the prior distribution being chosen specifically to ensure that the scale of the kernels was invariant to network width. This is an issue inherent to the use of MAP inference, which often finds modes that give a poor characterisation of the Bayesian posterior. In contrast, if we sample the weights using Langevin sampling (Appendix C), and set all the intermediate widths, from N_1

to N_L to N , then we get a similar expression,

$$\mathbf{K}_\ell = (\mathbf{K}_L \mathbf{K}_0^{-1})^{\ell/L} \mathbf{K}_0, \quad (24)$$

where the kernels slowly transition from \mathbf{K}_0 to \mathbf{K}_L . The key difference is that the similarity between the top-layer representation, \mathbf{K}_L , and the output kernel, \mathbf{K}_{L+1} , depends on the ratio between the network width, N , and the number of output units, $Y = N_{L+1}$. In particular, if $Y = N$, then we get a relationship very similar to that for MAP inference,

$$\mathbf{K}_\ell = (\mathbf{K}_{L+1} \mathbf{K}_0^{-1})^{\ell/(L+1)} \mathbf{K}_0, \quad (25)$$

However, as the network width grows very large, the prior begins to dominate, and the posterior becomes dominated by the prior,

$$\lim_{N/Y \rightarrow \infty} \mathbf{K}_\ell = \mathbf{K}_0, \quad (26)$$

as $\mathbf{K}_L = \mathbf{K}_0$. Finally, if the network width is small in comparison to the number of units,

$$\lim_{N/Y \rightarrow 0} \mathbf{K}_\ell = (\mathbf{K}_{L+1} \mathbf{K}_0^{-1})^{\ell/L} \mathbf{K}_0, \quad (27)$$

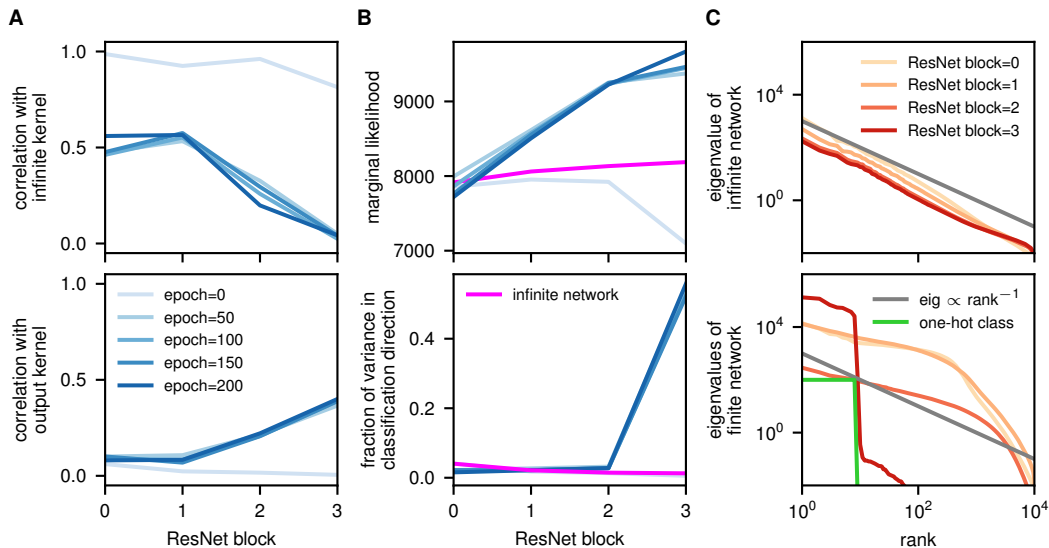


Figure 4. Comparison of kernels for finite and infinite neural networks at different layers. All kernels are computed on test data. **A** (top) Correlation (coefficient) between the kernel defined by the infinite network, and kernel defined by a finite network after different numbers of training epochs. **A** (bottom) Correlation (coefficient) between the kernel defined by the infinite network, and the output kernel defined by taking the inner product of one-hot vectors representing the class label. **B** (top) The Gaussian process marginal likelihood for the 10 functions given by the one-hot class labels, evaluated using the kernel output by different ResNet blocks. **B** (bottom) The fraction of variance in the direction of the one-hot output class labels. **C** (top) The eigenvalues of the kernel defined by the infinite network as we progress through layers, and compared to a -1 power law (grey). **C** (bottom) The eigenvalues of the kernel defined by the finite network after 200 training epochs, as we progress through ResNet blocks.

as the top-layer kernel converges to the output, $\mathbf{K}_L = \mathbf{K}_{L+1}$.

The above results suggest that finite neural networks perform well by giving flexibility to interpolate between the input kernel and output kernel. To see how this happens in real neural networks, we considered a 34-layer ResNet without batchnorm corresponding to the infinite network in Garriga-Alonso et al. (2019) trained on CIFAR-10. We began by computing the correlation between elements of the finite and infinite kernel (Fig. 4A top) as we go through ResNet blocks (x -axis), and as we go through training (blue lines). As expected, the randomly initialized, untrained network retains a high correlation with the infinite kernel at all layers, though the correlation is somewhat smaller for higher layers, as there have been more time for random sampling to build up discrepancies. However, for trained networks, this correspondence between the finite and infinite networks is far weaker: even at the first layer, the correlation is only around 0.5, and as we go through layers, the correlation decreases to almost zero. To understand whether kernels were being actively shaped, we computed the correlation between the kernel for the finite network and the output kernel, defined by taking the inner product of vectors representing the one-hot class labels (Fig. 4A bottom). We found that while the correlation for the untrained net-

work decreased across layers, training gives strong positive correlations with the output kernel, and these correlations increase as we move through network layers. Combined, these results indicate that the top-layer representation is much closer to the output kernel, as suggested by the deep linear results, than it is to the corresponding infinite network. While correlation is a useful simple measure of similarity, there are other measures of similarity that take into account the special structure of kernel matrices. In particular, we considered the marginal likelihood for the one-hot outputs corresponding to the class label, under a GP, with a kernel given by a scaled sum of the kernel at that ResNet block, and the identity (see Appendix D; Fig. 4B top). For the infinite network, the marginal likelihood increased somewhat as we moved through network layers, and the untrained finite network performed similarly, except that there was a decrease in performance at the last layer. In contrast, the marginal likelihood for the finite, trained networks was initially very close to the infinite networks, but grew rapidly as we move through ResNet blocks.

To gain an insight into how training shaped the neural network kernels, we computed the variance in the subspace defined by the one-hot outputs (i.e. the classification directions; Fig. 4B bottom). We might have expected to see a steady increase in the variance in this subspace as we move

through layers, but in fact the level was very small, only rising appreciably at the final block, and only for trained networks. To try to understand these results, we computed the eigenvalue spectrum of the kernels. For the infinite network (Fig. 4C top), we found that the eigenvalue spectrum at all levels decayed as a -1 power law. This is expected at the lowest level due to the well known $1/f$ power spectrum of images (Van der Schaaf & van Hateren, 1996), but is not necessarily the case at higher-levels. Given the power-spectrum of the output kernel is just a small set of equal-sized eigenvalues corresponding to the class labels (Fig. 4C bottom, green line), we might expect the eigenspectrum of finite networks to gradually get steeper as we move through network layers. In fact, we find the opposite: for intermediate layers, the eigenvalue spectrum becomes flatter, which can be interpreted as the network attempting to retain as much information as possible about all aspects of the image. It is only at the last layer where the relevant information is selected, giving an eigenvalue spectrum with around 10 – 1 large and roughly equally-sized eigenvalues, followed by much smaller eigenvalues, which mirrors the spectrum of the output kernel. This again confirms that the top-layer representation in trained networks is much closer to the output kernel than it is to corresponding infinite network.

5. Related work

Agrawal et al. (2020) independently introduced infinite networks with finite bottlenecks, but then made a very different contribution in that context. In particular, they highlighted that if we take the limit as some layers of a neural network go to infinity, convergence to the infinite networks with bottlenecks considered here is not immediate, but requires the neural network components to exhibit sufficient uniformity with respect to their inputs. In contrast, we show that finite bottlenecks can introduce flexibility and thereby improve performance even in two-layer linear networks, give analytic results in the case of linear networks, and show that these considerations are likely to be important in realistic large-scale networks, by showing that the kernel for a trained ResNet differs dramatically from that for the corresponding infinite network.

Technically, our work bears similarity to classical work on the dynamics of gradient descent in unregularised deep linear networks (Saxe et al., 2013). Importantly, the lack of regularisation in this work implies that infinitely many optimal solutions are available (e.g. all the lower-layer weights being fixed to the identity). In contrast, we focused on Bayesian inference, but also considered the optimal solution for regularised networks, which are much more constrained.

6. Conclusions

We have shown that finite Bayesian neural networks have more flexibility than infinite networks, and that this may explain the superior performance of finite networks. Thus, we introduced infinite networks with bottlenecks, and argue that they may be as incorporate flexibility and are able to perform representation learning, they may be a better model of real neural networks. We then assessed the flexibility of deep linear networks from two perspectives. First, we looked at the prior viewpoint: the variability in the top-layer kernel induced by the prior over a finite neural network. Second, we looked at the posterior viewpoint: the ability of the learning process to shape the top-layer kernel. Under both MAP inference and sampling in finite networks, learning gradually shaped top-layer representations so as to match the output-kernel. But, as Bayesian neural networks increase in width, the kernels become gradually less flexible, eliminating the possibility for learning to shape the kernel. In contrast, for MAP inference, the degree of kernel shaping is not affected by network width, and this additional flexibility might be an avenue for overfitting.

Acknowledgements

I would like to thank Adrià Garriga-Alonso, Sebastian Ober and Vidhi Lalchand for useful discussions.

References

- Agrawal, D., Papamarkou, T., and Hinkle, J. Wide neural networks with bottlenecks are deep gaussian processes. *arXiv preprint arXiv:2001.00921*, 2020.
- Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R., and Wang, R. On exact computation with an infinitely wide neural net. *arXiv preprint arXiv:1904.11955*, 2019.
- Bui, T., Hernández-Lobato, D., Hernandez-Lobato, J., Li, Y., and Turner, R. Deep gaussian processes for regression using approximate expectation propagation. In *International Conference on Machine Learning*, pp. 1472–1481, 2016.
- Chen, T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. Neural ordinary differential equations. In *Advances in neural information processing systems*, pp. 6571–6583, 2018.
- Cho, Y. and Saul, L. K. Kernel methods for deep learning. *NeurIPS*, 2009.
- Garriga-Alonso, A., Rasmussen, C. E., and Aitchison, L. Deep convolutional networks as shallow Gaussian processes. *ICLR*, 2019.

- He, K., Zhang, X., Ren, S., and Sun, J. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016.
- Huh, M., Agrawal, P., and Efros, A. A. What makes imagenet good for transfer learning? *arXiv preprint arXiv:1608.08614*, 2016.
- Lee, J., Bahri, Y., Novak, R., Schoenholz, S. S., Pennington, J., and Sohl-Dickstein, J. Deep neural networks as Gaussian processes. *ICLR*, 2018.
- Li, Z., Wang, R., Yu, D., Du, S. S., Hu, W., Salakhutdinov, R., and Arora, S. Enhanced convolutional neural tangent kernels. *arXiv preprint arXiv:1911.00809*, 2019.
- Matthews, A., Rowland, M., Hron, J., Turner, R., and Ghahramani, Z. Gaussian process behaviour in wide deep neural networks. *ICLR*, 2018.
- Novak, R., Xiao, L., Bahri, Y., Lee, J., Yang, G., Hron, J., Abolafia, D. A., Pennington, J., and Sohl-Dickstein, J. Bayesian deep convolutional networks with many channels are Gaussian processes. *ICLR*, 2019.
- Ramsey, F. P. Truth and probability. In *Readings in Formal Epistemology*, pp. 21–45. Springer, 1926.
- Rasmussen, C. E. and Williams, C. K. *Gaussian processes for machine learning*. MIT press, 2006.
- Saxe, A. M., McClelland, J. L., and Ganguli, S. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.
- Van der Schaaf, v. A. and van Hateren, J. v. Modelling the power spectra of natural images: statistics and information. *Vision research*, 36(17):2759–2770, 1996.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.