

---

# Deep Convolutional Networks as shallow Gaussian Processes

---

**Adrià Garriga-Alonso**  
Department of Engineering  
University of Cambridge  
ag919@cam.ac.uk

**Laurence Aitchison**  
Department of Engineering  
University of Cambridge  
laurence.aitchison@gmail.com

**Carl Edward Rasmussen**  
Department of Engineering  
University of Cambridge  
cer54@cam.ac.uk

## Abstract

We show that the output of a (residual) convolutional neural network (CNN) with an appropriate prior over the weights and biases is a Gaussian process (GP) in the limit of infinitely many convolutional filters, extending similar results for dense networks. For a CNN, the equivalent kernel can be computed exactly and, unlike “deep kernels”, has very few parameters: only the hyperparameters of the original CNN. Further, we show that this kernel has two properties that allow it to be computed efficiently; the cost of evaluating the kernel for a pair of images is similar to a single forward pass through the original CNN with only one filter per layer. The kernel equivalent to a 32-layer ResNet obtains 0.84% classification error on MNIST, a new record for GPs with a comparable number of parameters.<sup>1</sup>

## 1 Introduction

Convolutional Neural Networks (CNNs) have powerful pattern-recognition capabilities that have recently given dramatic improvements in important tasks involving images, such as classification [1]. However, as CNNs are increasingly being applied in real-world, safety-critical domains, their vulnerability to adversarial examples [2, 3], and their poor uncertainty estimates are becoming increasingly problematic. Bayesian inference is a theoretically principled and demonstrably successful [4, 5] framework for learning in the face of uncertainty, which may also help to address the problems of adversarial examples [6]. Unfortunately, Bayesian inference in CNNs is extremely difficult due to the very large number of parameters, requiring highly approximate factorised variational approximations [7, 8], or requiring the storage [9] of large numbers of approximate samples [10, 11].

Other methods such as Gaussian Processes (GPs) are more amenable to Bayesian inference, allowing us to compute the posterior uncertainty exactly [12]. This raises the question of whether it might be possible to combine the pattern-recognition capabilities of CNNs with exact probabilistic computations in GPs. Two such approaches exist in the literature. First, deep convolutional kernels [13] parameterise a GP kernel using the weights and biases of a CNN, which is used to embed the input images into some latent space before computing their similarity. The CNN parameters of the resulting kernel then have to be optimised by gradient descent. However, the large number of kernel parameters in the CNN reintroduces

---

<sup>1</sup>Code to replicate this paper is available at <https://github.com/rhaps0dy/convnets-as-gps>

the risk of overconfidence and overfitting, because they are treated in a non-Bayesian way. To avoid this risk, we need to infer a posterior over the CNN kernel parameters, which is as difficult as directly inferring a posterior over the parameters of the original CNN. Second, it is possible to define a convolutional GP [14] or a deep convolutional GP [15] by defining a GP that takes an image patch as input, and using that GP as a component in a larger CNN-like system. However, inference in such systems is very computationally expensive, at least without the use of potentially severe variational approximations [14].

An alternative approach is suggested by the underlying connection between Bayesian neural networks (NNs) and GPs. In particular, Neal [16] showed that the function defined by a single-layer fully-connected NN with infinitely many hidden units, and random independent zero-mean weights and biases is equivalent to a GP, implying that we can do exact Bayesian inference in such a NN by working with the equivalent GP. Recently, this result was extended to arbitrarily deep fully-connected NNs with infinitely many hidden units at each layer [17, 18]. However, these fully-connected networks are rarely used in practice, as they are unable to exploit important properties of images such as translational invariance, raising the question of whether state-of-the-art architectures such as CNNs [19] and ResNets [20] have equivalent GP representations. Here, we answer in the affirmative, giving the GP kernel corresponding to arbitrarily deep CNNs including networks with upconvolutions (“deconvolutions”), and to residual neural networks [20]. In this case, if each hidden layer has an infinite number of convolutional *filters*, the network prior is equivalent to a GP.

Furthermore, we show that two properties of the GP kernel induced by a CNN allow it to be computed very efficiently. First, in previous work it was necessary to compute the covariance matrix for the output of a single convolutional filter applied at all possible locations within a single image [14], which was prohibitively computationally expensive. In contrast, under our prior, the downstream weights are independent with zero-mean, which decorrelates the contribution from each location, and implies that it is necessary only to track the patch variances, and not their covariances. Second, while it is still necessary to compute the variance of the output of a convolutional filters applied at all locations within the image, the specific structure of the kernel induced by the CNN means that the variance at every location can be computed simultaneously and efficiently as a convolution.

Finally, we empirically demonstrate the performance increase coming from adding translation-invariant structure to the GP prior. Without computing any gradients, and without augmenting the training set (e.g. using translations), we obtain 0.84% error rate on the MNIST classification benchmark, setting a new record for nonparametric GP-based methods.

## 2 GP behaviour in a CNN

For clarity of exposition, we will treat the case of a 2D convolutional NN. The result applies straightforwardly to  $n$ D convolutions, dilated convolutions and upconvolutions (“deconvolutions”), since they can be represented as linear transformations with tied coefficients (see figure 1).

### 2.1 A 2D convolutional network prior

The network takes in an arbitrary input image  $\mathbf{X}$  of height  $H^{(0)}$  and width  $D^{(0)}$ , which is a  $C^{(0)} \times (H^{(0)}D^{(0)})$  real matrix. Each row, which we denote  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{C^{(0)}}$ , corresponds to a channel of the image (e.g.  $C^{(0)} = 3$  for RGB), flattened to form a vector. The first activations  $\mathbf{A}^{(1)}(\mathbf{X})$  are a linear transformation of the inputs. For  $j = 1, \dots, C^{(1)}$ :

$$\mathbf{a}_j^{(1)}(\mathbf{X}) := b_j^{(0)} + \sum_{i=1}^{C^{(0)}} \mathbf{W}_{j,i}^{(0)} \mathbf{x}_i. \quad (1)$$

We consider a network with  $L$  hidden layers. The other activations of the network, from  $\mathbf{A}^{(2)}(\mathbf{X})$  up to  $\mathbf{A}^{(L+1)}(\mathbf{X})$ , are defined recursively:

$$\mathbf{a}_j^{(\ell+1)}(\mathbf{X}) := b_j^{(\ell)} + \sum_{i=1}^{C^{(\ell)}} \mathbf{W}_{j,i}^{(\ell)} \phi(\mathbf{a}_i^{(\ell)}(\mathbf{X})). \quad (2)$$

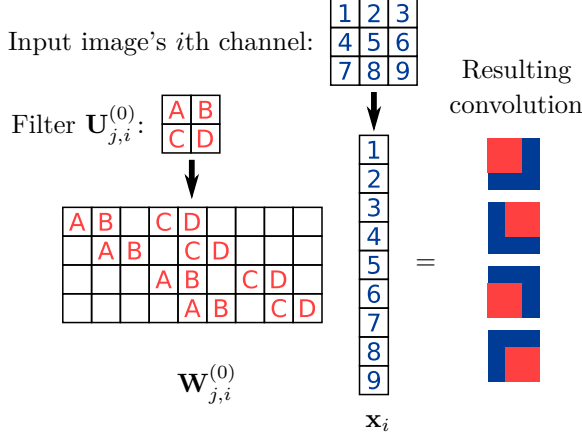


Figure 1: The 2D convolution  $\mathbf{U}_{j,i}^{(0)} * \mathbf{x}_i$  as the dot product  $\mathbf{W}_{j,i}^{(0)} \mathbf{x}_i$ . The blank elements of  $\mathbf{W}_{j,i}^{(0)}$  are zeros. The  $k$ th row of  $\mathbf{W}_{j,i}^{(0)}$  corresponds to applying the filter to the  $k$ th convolutional patch of the channel  $\mathbf{x}_i$ .

The activations  $\mathbf{A}^{(\ell)}(\mathbf{X})$  are  $C^{(\ell)} \times (H^{(\ell)} D^{(\ell)})$  matrices. Each row  $\mathbf{a}_j^{(\ell)}$  represents the flattened  $j$ th channel of the image that results from applying a convolutional filter to  $\phi(\mathbf{A}^{(\ell-1)}(\mathbf{X}))$ .

The structure of the pseudo-weight matrices  $\mathbf{W}_{j,i}^{(\ell)}$  and biases  $b_j^{(\ell)}$ , for  $i = 1, \dots, C^{(\ell)}$  and  $j = 1, \dots, C^{(\ell+1)}$ , depends on the architecture. For a convolutional layer, each row of  $\mathbf{W}_{j,i}^{(\ell)}$  represents a *position* of the filter, such that the dot product of all the rows with the image vector  $\mathbf{x}_i$  represents applying the convolutional filter  $\mathbf{U}_{j,i}^{(\ell)}$  to the  $i$ th channel. Thus, the elements of each row of  $\mathbf{W}_{j,i}^{(\ell)}$  are: 0 where the filter does not apply and the corresponding element of  $\mathbf{U}_{j,i}^{(\ell)}$  where it does, as illustrated in figure 1. For the  $g$ th position of the filter, the elements of the previous layer’s activation where the filter applies form the  $g$ th *convolutional patch*.

The outputs of the network are the last activations,  $\mathbf{A}^{(L+1)}(\mathbf{X})$ . In the classification or regression setting, the outputs are not spatially extended, so we have  $H^{(L+1)} = D^{(L+1)} = 1$ , which is equivalent to a fully-connected output layer. In this case, the pseudo-weights  $\mathbf{W}_{j,i}^{(L)}$  only have one row, and the activations  $\mathbf{a}_j^{(L+1)}$  are single-element vectors.

Finally, we define the prior distribution over functions by making the filters  $\mathbf{U}_{j,i}^{(\ell)}$  and biases  $b_j^{(\ell)}$  be independent Gaussian random variables (RVs). For each layer  $\ell$ , channels  $j, i$  and locations within the filter  $x, y$ :

$$u_{j,i,x,y}^{(\ell)} \sim \mathcal{N}\left(0, \sigma_w^2 / C^{(\ell)}\right), \quad b_j^{(\ell)} \sim \mathcal{N}\left(0, \sigma_b^2\right). \quad (3)$$

Note that, to keep the activation variance constant, the weight variance is divided by the number of input channels. The weight variance can also be divided by the number of elements of the filter, which makes it equivalent to the NN weight initialisation scheme introduced by He et al. [20].

## 2.2 Argument for GP behaviour

We follow the proofs by Lee et al. [17] and Matthews et al. [18] to show that the output of the CNN described in the previous section,  $\mathbf{A}^{(L+1)}$ , defines a GP indexed by the inputs,  $\mathbf{X}$ . Their proof [17] proceeds by applying the multivariate Central Limit Theorem (CLT) to each layer in sequence, i.e. taking the limit as  $N^{(1)} \rightarrow \infty$ , then  $N^{(2)} \rightarrow \infty$  etc, where  $N^{(\ell)}$  is the number of hidden units in layer  $\ell$ . By analogy, we sequentially apply the multivariate CLT by taking the limit as the number of channels goes to infinity, i.e.  $C^{(1)} \rightarrow \infty$ , then

$C^{(2)} \rightarrow \infty$  etc. While this is the simplest approach to taking the limits, other potentially more realistic approaches also exist [18].

The fundamental quantity we consider is a vector formed by concatenating the feature maps (or equivalently channels),  $\mathbf{a}_j^{(\ell)}(\mathbf{X})$  and  $\mathbf{a}_j^{(\ell)}(\mathbf{X}')$  from data points  $\mathbf{X}$  and  $\mathbf{X}'$ ,

$$\mathbf{a}_j^{(\ell)}(\mathbf{X}, \mathbf{X}') = \begin{pmatrix} \mathbf{a}_j^{(\ell)}(\mathbf{X}) \\ \mathbf{a}_j^{(\ell)}(\mathbf{X}') \end{pmatrix}. \quad (4)$$

This quantity (and the following arguments) can all be extended to the case of arbitrarily many input points.

**Induction base case.** For any pair of data points,  $\mathbf{X}$  and  $\mathbf{X}'$  the feature-maps corresponding to the  $j$ th channel,  $\mathbf{a}_j^{(1)}(\mathbf{X}, \mathbf{X}')$  have a multivariate Gaussian joint distribution. This is because each element is a linear combination of shared Gaussian random variables: the biases,  $\mathbf{b}_j^{(0)}$  and the filters,  $\mathbf{U}_{j,:}^{(0)}$ . Following equation (1),

$$\mathbf{a}_j^{(1)}(\mathbf{X}, \mathbf{X}') = b_j^{(0)} \mathbf{1} + \sum_{i=1}^{C^{(0)}} \begin{pmatrix} \mathbf{W}_{j,i}^{(0)} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_{j,i}^{(0)} \end{pmatrix} \begin{pmatrix} \mathbf{x}_i \\ \mathbf{x}'_i \end{pmatrix}, \quad (5)$$

where  $\mathbf{1}$  is a vector of all-ones. While the feature maps themselves,  $\mathbf{a}_j^{(1)}(\mathbf{X}, \mathbf{X}')$ , display strong correlations, different feature maps are independent and identically distributed (iid) conditioned on  $\mathbf{X}$  and  $\mathbf{X}'$  (i.e.  $\mathbf{a}_j^{(1)}(\mathbf{X}, \mathbf{X}')$  and  $\mathbf{a}_{j'}^{(1)}(\mathbf{X}, \mathbf{X}')$  are iid for  $j \neq j'$ ), because the parameters for different feature-maps (i.e. the biases,  $b_j^{(0)}$  and the filters,  $\mathbf{U}_{j,:}^{(0)}$ ) are themselves iid.

**Induction step.** Consider the feature maps at the  $\ell$ th layer,  $\mathbf{a}_i^{(\ell)}(\mathbf{X}, \mathbf{X}')$ , to be iid multivariate Gaussian RVs (i.e. for  $i \neq i'$ ,  $\mathbf{a}_i^{(\ell)}(\mathbf{X}, \mathbf{X}')$  and  $\mathbf{a}_{i'}^{(\ell)}(\mathbf{X}, \mathbf{X}')$  are iid). Our goal is to show that, taking the number of channels at layer  $\ell$  to infinity (i.e.  $C^{(\ell)} \rightarrow \infty$ ), the same properties hold at the next layer (i.e. all feature maps,  $\mathbf{a}_j^{(\ell+1)}(\mathbf{X}, \mathbf{X}')$ , are iid multivariate Gaussian RVs). Writing eq. (2) for two training examples,  $\mathbf{X}$  and  $\mathbf{X}'$ , we obtain,

$$\mathbf{a}_j^{(\ell+1)}(\mathbf{X}, \mathbf{X}') = b_j^{(\ell)} \mathbf{1} + \sum_{i=1}^{C^{(\ell)}} \begin{pmatrix} \mathbf{W}_{j,i}^{(\ell)} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_{j,i}^{(\ell)} \end{pmatrix} \phi(\mathbf{a}_i^{(\ell)}(\mathbf{X}, \mathbf{X}')) \quad (6)$$

We begin by showing that  $\mathbf{a}_j^{(\ell+1)}(\mathbf{X}, \mathbf{X}')$  is a multivariate Gaussian RV. The first term is multivariate Gaussian, as it is a linear combination of  $b_j^{(\ell)}$ , which is itself iid Gaussian. We can apply the multivariate CLT to show that the second term is also Gaussian, because, in the limit as  $C^{(\ell)} \rightarrow \infty$ , it is the sum of infinitely many iid terms:  $\mathbf{a}_i^{(\ell)}(\mathbf{X}, \mathbf{X}')$  are iid by assumption, and  $\mathbf{W}_{j,i}^{(\ell)}$  are iid by definition. Note that the same argument applies to all feature maps jointly, so all elements of  $\mathbf{A}^{(\ell)}(\mathbf{X}, \mathbf{X}')$  (defined by analogy with eq. 4) are jointly multivariate Gaussian.

To complete the proof, we need to show that the output feature maps are iid, i.e.  $\mathbf{a}_j^{(\ell+1)}(\mathbf{X}, \mathbf{X}')$  and  $\mathbf{a}_{j'}^{(\ell+1)}(\mathbf{X}, \mathbf{X}')$  are iid for  $j \neq j'$ . They are identically distributed, as  $b_j^{(\ell)}$  and  $\mathbf{W}_{j,i}^{(\ell)}$  are iid and  $\phi(\mathbf{a}_i^{(\ell)}(\mathbf{X}, \mathbf{X}'))$  is shared. To show that they are independent, remember that  $\mathbf{a}_j^{(\ell+1)}(\mathbf{X}, \mathbf{X}')$  and  $\mathbf{a}_{j'}^{(\ell+1)}(\mathbf{X}, \mathbf{X}')$  are jointly Gaussian, so it is sufficient to show that they are uncorrelated, and we can show that they are uncorrelated because the weights,  $\mathbf{W}_{j,i}^{(\ell)}$  are independent with zero-mean, eliminating any correlations that might arise through the shared RV,  $\phi(\mathbf{a}_i^{(\ell)}(\mathbf{X}, \mathbf{X}'))$ .

**Technical notes** Following Matthews et al. [18] we could take a more realistic limit to infinity: the number of channels  $C^{(\ell)}$  for all the layers grow together, and are all finite at

every point before convergence. More formally: there are strictly increasing width functions  $f_\ell : \mathbb{N} \rightarrow \mathbb{N}$  such that  $C^{(1)} = f_1(n), \dots, C^{(L)} = f_L(n)$ , and we take the limit  $n \rightarrow \infty$ . This is a more problematic case, because at no point are the feature maps  $\mathbf{a}_j^{(\ell)}(\mathbf{X}, \mathbf{X}')$  *exactly* independent and Gaussian, so the multivariate CLT does not directly apply.

The proof of Gaussian behaviour for this case involves bounding the divergence from a Gaussian at every layer, and propagating the bound inductively. Matthews et al. [18] wrote the rigorous proof for dense networks, and we expect their bounds to also apply (within a constant factor) to convolutional networks.

Finally, note that as long as we apply the multivariate CLT, we do not need the weights to be Gaussian, allowing us to capture alternative initialization schemes (e.g. drawing parameters from a zero-mean Uniform distribution, such as in ‘‘Xavier’’ weight initialisation [21]).

### 3 The ConvNet and ResNet kernels

Here we derive a computationally efficient kernel corresponding to the CNN described in the previous section. It is surprising that we can compute the kernel efficiently because the feature maps,  $\mathbf{a}_j^{(\ell)}(\mathbf{X})$ , display rich covariance structure due to the shared convolutional filter. Computing and representing these covariances would be prohibitively computationally expensive. However, in many cases we only need the variance of the output, e.g. in the case of classification or regression with a final dense layer. It turns out that this propagates backwards through the convolutional network, implying that for every layer, we only need the *marginal covariance* of the activations: the covariance between the corresponding elements of  $\mathbf{a}_j^{(\ell)}(\mathbf{X})$  and  $\mathbf{a}_j^{(\ell)}(\mathbf{X}')$ .

#### 3.1 GP mean and covariance

A GP is completely specified by its mean and covariance (kernel) functions. These give the parameters of the joint Gaussian distribution of the RVs indexed by any two inputs,  $\mathbf{X}$  and  $\mathbf{X}'$ . For the purposes of computing the mean and covariance, it is easiest to consider the network as being written entirely in index notation,

$$A_{j,g}^{(\ell+1)}(\mathbf{X}) = b_j^{(\ell)} + \sum_{i=1}^{C^{(\ell)}} \sum_{h=1}^{H^{(\ell)}D^{(\ell)}} W_{j,i,g,h}^{(\ell)} \phi(A_{i,h}^{(\ell)}(\mathbf{X})).$$

where  $\ell$  and  $\ell + 1$  denote the input and output layers respectively,  $i$  and  $j$  denote the input and output channels, and  $h$  and  $g$  denote the location within the input and output channel or feature-maps.

The mean function is thus easy to compute

$$\mathbb{E} \left[ A_{j,g}^{(\ell+1)}(\mathbf{X}) \right] = \mathbb{E} \left[ b_j^{(\ell)} \right] + \sum_{i=1}^{C^{(\ell)}} \sum_{h=1}^{H^{(\ell)}D^{(\ell)}} \mathbb{E} \left[ W_{j,i,g,h}^{(\ell)} \phi(A_{i,h}^{(\ell)}(\mathbf{X})) \right] = 0.$$

as  $b_j^{(\ell)}$  and  $W_{j,i,g,h}^{(\ell)}$  have zero mean, and  $W_{j,i,g,h}^{(\ell)}$  are independent of the activations at the previous layer,  $\phi(A_{i,h}^{(\ell)}(\mathbf{X}))$ .

Now we show that it is possible to efficiently compute the covariance function. This is surprising because for many networks, we need to compute the covariance of activations between all locations in the feature map (i.e.  $\mathbb{C} \left[ A_{j,g}^{(\ell+1)}(\mathbf{X}), A_{j,g'}^{(\ell+1)}(\mathbf{X}') \right]$ ) and this object is extremely high-dimensional,  $N^2(H^{(\ell)}D^{(\ell)})^2$ . However, it turns out that we only need to consider the covariance across different data points,  $\mathbf{X}$  and  $\mathbf{X}'$  at the same location within the feature map,  $g$ , which is a more manageable quantity of size  $N^2(H^{(\ell)}D^{(\ell)})$ .

This is true at the output layer ( $L + 1$ ): in order to achieve an output suitable for classification or regression, we use only a single output location  $H^{(L+1)} = D^{(L+1)} = 1$ , so it is only possible to compute the covariance at that single location. We now show that, if we only

need the covariance at corresponding locations in the outputs, we only need the covariance at corresponding locations in the inputs, and this requirement propagates backwards through the network.

Formally, as the activations are composed of a sum of terms, their covariance is the sum of the covariances of all those underlying terms,

$$\begin{aligned} \mathbb{C} \left[ A_{j,g}^{(\ell+1)}(\mathbf{X}), A_{j,g}^{(\ell+1)}(\mathbf{X}') \right] &= \mathbb{V} \left[ b_j^{(\ell)} \right] + \\ &\sum_{i=1}^{C^{(\ell)}} \sum_{i'=1}^{C^{(\ell)}} \sum_{h=1}^{H^{(\ell)}D^{(\ell)}} \sum_{h'=1}^{H^{(\ell)}D^{(\ell)}} \mathbb{C} \left[ W_{j,i,g,h}^{(\ell)} \phi(A_{i,h}^{(\ell)}(\mathbf{X})), W_{j,i',g,h'}^{(\ell)} \phi(A_{i',h'}^{(\ell)}(\mathbf{X}')) \right]. \end{aligned} \quad (7)$$

As the terms in the covariance have mean zero, and as the weights and activations from the previous layer are independent,

$$\begin{aligned} \mathbb{C} \left[ A_{j,g}^{(\ell+1)}(\mathbf{X}), A_{j,g}^{(\ell+1)}(\mathbf{X}') \right] &= \sigma_b^2 + \\ &\sum_{i=1}^{C^{(\ell)}} \sum_{i'=1}^{C^{(\ell)}} \sum_{h=1}^{H^{(\ell)}D^{(\ell)}} \sum_{h'=1}^{H^{(\ell)}D^{(\ell)}} \mathbb{E} \left[ W_{j,i,g,h}^{(\ell+1)} W_{j,i',g,h'}^{(\ell+1)} \right] \mathbb{E} \left[ \phi(A_{i,h}^{(\ell)}(\mathbf{X})) \phi(A_{i',h'}^{(\ell)}(\mathbf{X}')) \right]. \end{aligned} \quad (8)$$

The weights are independent for different channels:  $W_{j,i,g,h}^{(\ell+1)}$  and  $W_{j,i',g,h'}^{(\ell+1)}$  are iid for  $i \neq i'$ , so their covariance is 0. Each row  $g$  of the weight matrices  $\mathbf{W}_{j,i}^{(\ell+1)}$  only contains independent variables or zeros (figure 1), so  $\mathbb{C} \left[ W_{j,i,g,h}^{(\ell+1)}, W_{j,i',g,h'}^{(\ell+1)} \right] = 0$  for  $h \neq h'$ . Thus, we can eliminate the sums over  $i'$  and  $h'$ :

$$\mathbb{C} \left[ A_{j,g}^{(\ell+1)}(\mathbf{X}), A_{j,g}^{(\ell+1)}(\mathbf{X}') \right] = \sigma_b^2 + \sum_{i=1}^{C^{(\ell)}} \sum_{h=1}^{H^{(\ell)}D^{(\ell)}} \mathbb{E} \left[ W_{j,i,g,h}^{(\ell+1)} W_{j,i,g,h}^{(\ell+1)} \right] \mathbb{E} \left[ \phi(A_{i,h}^{(\ell)}(\mathbf{X})) \phi(A_{i,h}^{(\ell)}(\mathbf{X}')) \right]. \quad (9)$$

The  $g$ th row of  $\mathbf{W}_{i,j}^{(\ell+1)}$  is zero for indices  $h$  that don't belong to its convolutional patch, so we can restrict the sum over  $h$  to that region. We also define  $v_g^{(1)}(\mathbf{X}, \mathbf{X}')$ , to emphasise that the covariances are independent of the output channel,  $j$ . The variance of the first layer is

$$v_g^{(1)}(\mathbf{X}, \mathbf{X}') = \mathbb{C} \left[ A_{j,g}^{(1)}(\mathbf{X}), A_{j,g}^{(1)}(\mathbf{X}') \right] = \sigma_b^2 + \frac{\sigma_w^2}{C^{(0)}} \sum_{i=1}^{C^{(0)}} \sum_{h \in g\text{th patch}} X_{i,h} X'_{i,h}. \quad (10)$$

And we do the same for the other layers,

$$v_g^{(\ell+1)}(\mathbf{X}, \mathbf{X}') = \mathbb{C} \left[ A_{j,g}^{(\ell+1)}(\mathbf{X}), A_{j,g}^{(\ell+1)}(\mathbf{X}') \right] = \sigma_b^2 + \sigma_w^2 \sum_{h \in g\text{th patch}} s_h^{(\ell)}(\mathbf{X}, \mathbf{X}'), \quad (11)$$

where

$$s_h^{(\ell)}(\mathbf{X}, \mathbf{X}') = \mathbb{E} \left[ \phi(A_{i,h}^{(\ell)}(\mathbf{X})) \phi(A_{i,h}^{(\ell)}(\mathbf{X}')) \right] \quad (12)$$

is the covariance of the activations, which is again independent of the channel.

### 3.2 Covariance of the activities

The elementwise covariance in the right-hand side of equation (11) can be computed in closed form for many choices of  $\phi$  if the activations are Gaussian. For each element of the activations, one needs to keep track of the 3 distinct entries of the bivariate covariance matrix between the inputs,  $v_g^{(\ell)}(\mathbf{X}, \mathbf{X})$ ,  $v_g^{(\ell)}(\mathbf{X}, \mathbf{X}')$  and  $v_g^{(\ell)}(\mathbf{X}', \mathbf{X}')$ .

For example, for the ReLU nonlinearity ( $\phi(x) = \max(0, x)$ ), one can adapt Cho and Saul [22] in the same way as Matthews et al. [18, section 3] to obtain

$$s_g^{(\ell)}(\mathbf{X}, \mathbf{X}') = \frac{\sqrt{v_g^{(\ell)}(\mathbf{X}, \mathbf{X})v_g^{(\ell)}(\mathbf{X}', \mathbf{X}')}}{\pi} \left( \sin \theta_g^{(\ell)} + (\pi - \theta_g^{(\ell)}) \cos \theta_g^{(\ell)} \right) \quad (13)$$

---

**Algorithm 1** The ConvNet kernel  $k(\mathbf{X}, \mathbf{X}')$ 

---

- 1: *Input*: two images,  $\mathbf{X}, \mathbf{X}' \in \mathbb{R}^{C^{(0)} \times (H^{(0)} W^{(0)})}$ .
  - 2: Compute  $v_g^{(1)}(\mathbf{X}, \mathbf{X})$ ,  $v_g^{(1)}(\mathbf{X}, \mathbf{X}')$ , and  $v_g^{(1)}(\mathbf{X}', \mathbf{X}')$ ; using equation (10).
  - 3: **for**  $\ell = 1, 2, \dots, L$  **do**
  - 4:   Compute  $s_g^{(\ell)}(\mathbf{X}, \mathbf{X}')$ ,  $s_g^{(\ell)}(\mathbf{X}, \mathbf{X})$  and  $s_g^{(\ell)}(\mathbf{X}', \mathbf{X}')$  using equation (13), (14), or some other nonlinearity.
  - 5:   Compute  $v_g^{(\ell+1)}(\mathbf{X}, \mathbf{X})$ ,  $v_g^{(\ell+1)}(\mathbf{X}, \mathbf{X}')$ , and  $v_g^{(\ell+1)}(\mathbf{X}', \mathbf{X}')$ ; using equation (11).
  - 6: **end for**
  - 7: Output the scalar  $v_1^{(L+1)}(\mathbf{X}, \mathbf{X}')$ .
- 

where  $\theta_g^{(\ell)} = \cos^{-1} \left( v_g^{(\ell)}(\mathbf{X}, \mathbf{X}') / \sqrt{v_g^{(\ell)}(\mathbf{X}, \mathbf{X}) v_g^{(\ell)}(\mathbf{X}', \mathbf{X}')} \right)$ . For  $\phi$  as the error function (very similar to the hyperbolic tangent), Williams [23] gives the expression

$$s_g^{(\ell)}(\mathbf{X}, \mathbf{X}') = \frac{2}{\pi} \sin^{-1} \left( \frac{2v_g^{(\ell)}(\mathbf{X}, \mathbf{X}')}{\sqrt{(1 + 2v_g^{(\ell)}(\mathbf{X}, \mathbf{X}))(1 + 2v_g^{(\ell)}(\mathbf{X}', \mathbf{X}'))}} \right). \quad (14)$$

### 3.3 Efficiency of the ConvNet kernel

We now have all the pieces for computing the kernel, as written in Algorithm 1.

Putting together equations (11), and (13) or (14), gives us the surprising result that the marginal covariances of the activations at layer  $\ell + 1$  only depend on the marginal covariances of the activations at layer  $\ell$ . This is very important, because it makes the computational cost of the kernel be within a constant factor of the cost of a forward pass for the equivalent CNN with 1 filter per layer.

Thus, the algorithm is more efficient than one would naively think. A priori, one needs to compute the covariance between all the elements of  $\mathbf{a}_j^{(\ell)}(\mathbf{X})$  and  $\mathbf{a}_j^{(\ell)}(\mathbf{X}')$  combined, yielding a  $2H^{(\ell)}D^{(\ell)} \times 2H^{(\ell)}D^{(\ell)}$  covariance matrix for every pair of points. Instead, we only need to keep track of a  $H^{(\ell)}D^{(\ell)}$ -dimensional vector per layer and pair of points, and an amortised  $H^{(\ell)}D^{(\ell)}$ -dimensional vector per *point* and layer.

Furthermore, the particular form for the kernel (eq. 1 and eq. 2) implies that the required variances and covariances at all required locations can be computed efficiently as a convolution.

### 3.4 Kernel for a residual CNN

The induction step in the argument for GP behaviour from section 2.2 depends only on the previous activations being iid Gaussian. Since all the activations are iid Gaussian, we can add skip connections between the activations of different layers while preserving GP behaviour, e.g.  $\mathbf{A}^{(\ell+1)}$  and  $\mathbf{A}^{(\ell-s)}$  where  $s$  is the number of layers that the skip connection spans. If we change the NN recursion (equation 2) to

$$\mathbf{a}_j^{(\ell+1)}(\mathbf{X}) := \mathbf{a}_j^{(\ell-s)}(\mathbf{X}) + \mathbf{b}_j^{(\ell)} + \sum_{i=1}^{C^{(\ell)}} \mathbf{W}_{j,i}^{(\ell)} \phi \left( \mathbf{a}_i^{(\ell)}(\mathbf{X}) \right), \quad (15)$$

then the kernel recursion (equation 11) becomes

$$v_g^{(\ell+1)}(\mathbf{X}, \mathbf{X}') = v_g^{(\ell-s)}(\mathbf{X}, \mathbf{X}') + \sigma_b^2 + \sigma_w^2 \sum_{h \in g\text{th patch}} s_g^{(\ell)}(\mathbf{X}, \mathbf{X}'). \quad (16)$$

This way of adding skip connections is equivalent to the “pre-activation” shortcuts described by He et al. [24]. Remarkably, the natural way of adding residual connections to NNs is the one that performed best in their empirical evaluations.

Method	#samples	Validation error	Test error
NNGP [17]	$\approx 250$	–	1.21%
Convolutional GP [14]	SGD	–	1.17%
Deep Conv. GP [15]	SGD	–	1.34%
ConvNet GP	27	0.71%	1.03%
Residual CNN GP	27	0.72%	0.96%
ResNet GP	–	0.68%	<b>0.84%</b>
GP + parametric deep kernel [25]	SGD	–	0.60%
Vanilla ResNet [26]	–	–	<b>0.41%</b>

Table 1: MNIST classification results. #samples gives the number of kernels that were randomly sampled for the hyperparameter search. “ConvNet GP” and “Residual CNN GP” are random CNN architectures with a fixed filter size, whereas “ResNet GP” is a slight modification of the architecture by He et al. [24]. Entries labeled “SGD” used Stochastic Gradient Descent for tuning hyperparameters, by maximising the likelihood of the training set. The last two methods use parametric neural networks. The ResNet GP performs better than all other nonparametric approaches, but is still not as good as methods with many parameters.

## 4 Experiments

We evaluate our kernel on the MNIST handwritten digit classification task. Classification likelihoods are not conjugate for GPs, so we must make an approximation, and we follow Lee et al. [17], in re-framing classification as multi-output regression.

The training set is split into  $N = 50000$  training and 10000 validation examples. The regression targets  $\mathbf{Y} \in \{-1, 1\}^{N \times 10}$  are a one-hot encoding of the example’s class:  $y_{n,c} = 1$  if the  $n$ th example belongs to class  $c$ , and  $-1$  otherwise.

Training is exact conjugate likelihood GP regression with noiseless targets  $\mathbf{Y}$  [12]. First we compute the  $N \times N$  kernel matrix  $\mathbf{K}_{xx}$ , which contains the kernel between every pair of examples. Then we compute  $\mathbf{K}_{xx}^{-1}\mathbf{Y}$  using a linear system solver.

The test set has  $N_T = 10000$  examples. We compute the  $N_T \times N$  matrix  $\mathbf{K}_{x^*x}$ , the kernel between each test example and all the training examples. The predictions are given by the row-wise maximum of  $\mathbf{K}_{x^*x}\mathbf{K}_{xx}^{-1}\mathbf{Y}$ .

For the “ConvNet GP” and “Residual CNN GP”, we optimise the kernel hyperparameters by random search. We draw  $M$  random hyperparameter samples, compute the resulting kernel’s performance in the validation set, and pick the highest performing run. The kernel hyperparameters are:  $\sigma_b^2$ ,  $\sigma_w^2$ ; the number of layers; the convolution stride, filter sizes and edge behaviour; the nonlinearity (we consider the error function and ReLU); and the frequency of residual skip connections (for Residual CNN GPs). We do not retrain the model on the validation set after choosing hyperparameters.

The “ResNet GP” is the kernel equivalent to a 32-layer version of the basic residual architecture by He et al. [20]. The differences are: a  $3 \times 3$  convolutional initial layer, no max-pooling, and a dense layer instead of average pooling at the end. We chose to remove the pooling because computing its output variance requires the off-diagonal elements of the filter covariance, which would invalidate the efficiency gains described in section 3.3. See section 6 for a more detailed explanation.

**Computational efficiency.** Asymptotically, computing the kernel matrix takes  $O(N^2LD)$  time, where  $L$  is the number of layers in the network and  $D$  is the dimensionality of the input, and inverting the kernel matrix takes  $O(N^3)$ . As such, we expect that for very large datasets, inverting the kernel matrix will dominate the computation time. However, on MNIST,  $N^3$  is only around a factor of 10 larger than  $N^2LD$ . In practice, we found that it was more expensive to *compute* the kernel matrix than to invert it. For the ResNet kernel, the most expensive, computing  $\mathbf{K}_{xx}$ , and  $\mathbf{K}_{xx^*}$  for validation and test took 3h 40min on two Tesla P100 GPUs. In contrast, inverting  $\mathbf{K}_{xx}$  and computing validation and test performance took  $12.58 \pm 3$  min on two 14-core Xeon E5-2680 CPUs, totalling 56 threads.



## 5 Related Work

Van der Wilk et al. [14] also adapted GPs to image classification. They defined a prior on functions  $f$  that takes an image and outputs a scalar. First, draw a function  $g \sim \mathcal{GP}(0, k_p(\mathbf{X}, \mathbf{X}'))$ . Then,  $f$  is the sum of the output of  $g$  applied to each of the convolutional patches. The result is that  $f$  also follows a GP, with kernel  $k(\mathbf{X}, \mathbf{X}') := \sum_{i \in \text{patches}} \sum_{j \in \text{patches}} k_p(\mathbf{X}_i, \mathbf{X}'_j)$ .

Their approach is also inspired by convolutional NNs, but their kernel  $k_p$  is applied to all pairs of patches of  $\mathbf{X}$  and  $\mathbf{X}'$ . This makes their convolutional kernel expensive to evaluate, requiring inter-domain inducing point approximations to remain tractable. The kernels in this work, directly motivated by the infinite-filter limit of a CNN, only apply something like  $k_p$  to the *corresponding* pairs of patches within  $\mathbf{X}$  and  $\mathbf{X}'$  (equation 10). As such, the CNN kernels are cheaper to compute and exhibit superior performance (Table 1), despite the use of an approximate likelihood function.

Kumar et al. [15] define a prior over functions by stacking several GPs with van der Wilk’s convolutional kernel, forming a “Deep GP” [27]. In contrast, the kernel in this paper confines all hierarchy to the definition of the kernel, and the resulting GPs is shallow.

Wilson et al. [13] introduced and Bradshaw et al. [25] improved deep kernel learning. The inputs to a classic GP kernel  $k$  (e.g. RBF) are preprocessed by applying a feature extractor  $g$  (a deep NN) prior to computing the kernel:  $k_{\text{deep}}(\mathbf{X}, \mathbf{X}') := k(g(\mathbf{X}; \theta), g(\mathbf{X}'; \theta))$ . The NN parameters are optimised by gradient ascent using the likelihood as the objective, as in standard GP kernel learning [12, Chapter 5]. Since deep kernel learning incorporates a state-of-the-art NN with over  $10^6$  parameters, we expect it to perform similarly to a NN applied directly to the task of image classification. At present both CNNs and deep kernel learning display superior performance to the GP kernels in this work. However, the kernels defined here have far fewer parameters (around 10, compared to their  $10^6$ ).

Finally, Borovkyh [28] also suggests that a CNN exhibits GP behaviour. However, they take the infinite limit with respect to the *filter size*, not the number of filters. Thus, their infinite network is inapplicable to real data which is always of finite dimension.

## 6 Future work

Our work suggests three directions for future research: improving the classification performance by incorporating an improved likelihood and more features from state-of-the-art neural networks, improving the speed of training using approximation schemes and exploiting the ability of the presented GPs to perform close-to-ideal Bayesian inference.

At present our GP CNNs do not achieve state-of-the-art performance, as compared with a CNN applied directly to the task of image classification. We hope to close this gap by making two improvements. First, our networks do not currently incorporate average- or max-pooling, which are used on most state-of-the-art CNNs. CNNs with pooling are also equivalent to GPs, since they also keep the channels independent in equation (6). This is because pooling only operates separately on each activation channel, so the activations would still be iid after pooling. While the resulting kernels would be less efficient as some off-diagonal elements of the covariance would be necessary, the additional computational burden might be manageable for pooling over small regions (e.g.  $2 \times 2$ ). Second, we currently approximate classification as a regression problem, using an actual classification likelihood should lead to better models.

While we can compute and invert the kernel matrix exactly for intermediate size datasets (e.g. MNIST with 60,000 examples), this may become impossible with larger datasets such as ImageNet. As such, the practical application of these techniques in large-scale settings will require the use of GP posterior approximation schemes, for instance those based on inducing points.

The possibility of doing close-to-ideal Bayesian inference in large-scale, state-of-the-art networks has many applications, three of which are particularly salient. First, these networks are ideal for use in scenarios with small numbers of training examples, as they efficiently exploit all available data, and have far fewer hyperparameters (i.e. training method, learning

rate, momentum etc.) that need to be tuned. Second, there are suggestions that idealised Bayesian inference might help to reduce issues with adversarial examples [6]. It is important to investigate the degree to which infinite-filter networks might realise these benefits. Third, these kernels provide a tool for understanding the function of various NN features. In particular, does any given feature (e.g. dropout, pooling) simply make the network easier to optimise, or does it improve the underlying generalisation capability of the model?

## 7 Conclusions

We have shown that deep Bayesian CNNs with infinitely many filters are equivalent to a GP with a recursive kernel. We also derived the kernel for the GP equivalent to a CNN, and showed that, in handwritten digit classification, it outperforms all previous GP approaches that do not incorporate a parametric NN into the kernel. Given that most state of the art neural networks incorporate structure (convolutional or otherwise) into their architecture, the equivalence between CNNs and GPs is potentially of considerable practical relevance.

## Acknowledgements

We wish to thank Will Tebbutt, Richard Turner, and others, for helpful comments and feedback on this work.

## References

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [2] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, “Intriguing properties of neural networks,” *arXiv preprint arXiv:1312.6199*, 2013.
- [3] A. Kurakin, I. Goodfellow, and S. Bengio, “Adversarial examples in the physical world,” *arXiv preprint arXiv:1607.02533*, 2016.
- [4] J. Snoek, H. Larochelle, and R. P. Adams, “Practical bayesian optimization of machine learning algorithms,” in *Advances in neural information processing systems*, 2012, pp. 2951–2959.
- [5] M. Deisenroth and C. E. Rasmussen, “PILCO: A model-based and data-efficient approach to policy search,” in *Proceedings of the 28th International Conference on machine learning (ICML-11)*, 2011, pp. 465–472.
- [6] Y. Gal and L. Smith, “Idealised bayesian neural networks cannot have adversarial examples: Theoretical and empirical study,” *arXiv preprint arXiv:1806.00667*, 2018.
- [7] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, “Weight uncertainty in neural network,” in *International Conference on Machine Learning*, 2015, pp. 1613–1622.
- [8] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” *arXiv preprint arXiv:1506.02142*, 2015.
- [9] B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” in *Advances in Neural Information Processing Systems*, 2017, pp. 6402–6413.
- [10] M. Welling and Y. W. Teh, “Bayesian learning via stochastic gradient langevin dynamics,” in *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, 2011, pp. 681–688.
- [11] S. Mandt, M. D. Hoffman, and D. M. Blei, “Stochastic gradient descent as approximate bayesian inference,” *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 4873–4907, 2017.
- [12] C. E. Rasmussen and C. K. Williams, *Gaussian processes for machine learning*. MIT press Cambridge, 2006, vol. 1.

- [13] A. G. Wilson, Z. Hu, R. Salakhutdinov, and E. P. Xing, “Deep kernel learning,” in *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, A. Gretton and C. C. Robert, Eds., vol. 51. Cadiz, Spain: PMLR, 09–11 May 2016, pp. 370–378. [Online]. Available: <http://proceedings.mlr.press/v51/wilson16.html>
- [14] M. van der Wilk, C. E. Rasmussen, and J. Hensman, “Convolutional gaussian processes,” in *Advances in Neural Information Processing Systems*, 2017, pp. 2845–2854.
- [15] V. Kumar, V. Singh, P. Srijith, and A. Damianou, “Deep gaussian processes with convolutional kernels,” *arXiv preprint arXiv:1806.01655*, 2018.
- [16] R. M. Neal, *Bayesian Learning for Neural Networks*. Berlin, Heidelberg: Springer-Verlag, 1996.
- [17] J. Lee, Y. Bahri, R. Novak, S. S. Schoenholz, J. Pennington, and J. Sohl-Dickstein, “Deep neural networks as gaussian processes,” *arXiv preprint arXiv:1711.00165*, 2017.
- [18] A. G. d. G. Matthews, J. Hron, M. Rowland, R. E. Turner, and Z. Ghahramani, “Gaussian process behaviour in wide deep neural networks,” in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=H1-nGgWC->
- [19] Y. LeCun, B. E. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. E. Hubbard, and L. D. Jackel, “Handwritten digit recognition with a back-propagation network,” in *Advances in neural information processing systems*, 1990, pp. 396–404.
- [20] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778. [Online]. Available: <https://arxiv.org/pdf/1512.03385.pdf>
- [21] X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 2010, pp. 249–256.
- [22] Y. Cho and L. K. Saul, “Kernel methods for deep learning,” in *Advances in neural information processing systems*, 2009, pp. 342–350. [Online]. Available: <http://papers.nips.cc/paper/3628-kernel-methods-for-deep-learning.pdf>
- [23] C. K. Williams, “Computing with infinite networks,” in *Advances in neural information processing systems*, 1997, pp. 295–301.
- [24] K. He, X. Zhang, S. Ren, and J. Sun, “Identity mappings in deep residual networks,” in *European conference on computer vision*. Springer, 2016, pp. 630–645.
- [25] J. Bradshaw, A. G. d. G. Matthews, and Z. Ghahramani, “Adversarial examples, uncertainty, and transfer testing robustness in gaussian process hybrid deep networks,” *arXiv preprint arXiv:1707.02476*, 2017.
- [26] T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud, “Neural ordinary differential equations,” *arXiv preprint arXiv:1806.07366*, 2018.
- [27] A. Damianou and N. Lawrence, “Deep gaussian processes,” in *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, C. M. Carvalho and P. Ravikumar, Eds., vol. 31. Scottsdale, Arizona, USA: PMLR, 29 Apr–01 May 2013, pp. 207–215. [Online]. Available: <http://proceedings.mlr.press/v31/damianou13a.html>
- [28] A. Borovykh, “A gaussian process perspective on convolutional neural networks,” *ResearchGate*, 05 2018.