
An Analytic Solution to Covariance Propagation in Neural Networks

Oren Wright
Carnegie Mellon University

Yorie Nakahira
Carnegie Mellon University

José M. F. Moura
Carnegie Mellon University

Abstract

Uncertainty quantification of neural networks is critical to measuring the reliability and robustness of deep learning systems. However, this often involves costly or inaccurate sampling methods and approximations. This paper presents a sample-free moment propagation technique that propagates mean vectors and covariance matrices across a network to accurately characterize the input-output distributions of neural networks. A key enabler of our technique is an analytic solution for the covariance of random variables passed through nonlinear activation functions, such as Heaviside, ReLU, and GELU. The wide applicability and merits of the proposed technique are shown in experiments analyzing the input-output distributions of trained neural networks and training Bayesian neural networks.

1 INTRODUCTION

This paper presents an analytic moment propagation technique to accurately characterize the input-output distributions of deep neural networks. Neural networks have achieved state-of-the-art inference accuracy across many problem domains, but their application to safety-critical domains like autonomous vehicles, industrial robots, and medicine demands not only high accuracy but also high reliability and robustness. Accurate uncertainty quantification in neural networks enables measurements of reliability and robustness, since uncertainty quantification gives us a language in which we can reason about the likelihood and conditions of failure. Uncertainty quantification techniques can be used to analyze the robustness of

trained neural networks to input noise and adversarial attacks, or to train networks that represent predictive uncertainty to provide well-calibrated confidence scores and out-of-distribution detection.

Rather than explicitly representing uncertainty, neural networks are commonly trained to produce point estimates without rigorous uncertainty quantification. Although their optimization process has a probabilistic interpretation—e.g., minimizing a cross-entropy loss function corresponds to maximum likelihood estimation under certain conditions (MacKay, 2002)—networks trained in this way are often poorly calibrated (Guo et al., 2017) and can perform poorly against small input perturbations and adversarial attacks (Szegedy et al., 2014; Carlini and Wagner, 2017).

One approach to quantifying uncertainty in neural networks is by propagating statistical moments across a network, from input to output, layer by layer. In the existing literature, moment propagation has been used 1) to analyze the input-output distributions of any given (or well-trained) networks (Bibi et al., 2018); 2) to train probabilistic networks that explicitly model input uncertainty (Gast and Roth, 2018); and 3) to train fully Bayesian neural networks (BNNs) that model both input uncertainty and parameter uncertainty (Wu et al., 2019).

Uncertainty quantification remains a challenging area of research, because modern neural networks are nonlinear and represent data in high-dimensional spaces. For moment propagation in particular, interactions between the outputs of activation functions are challenging to characterize analytically. The moment propagation techniques proposed in prior work approximate the covariance between activation variables, as no analytic solution for covariance has been published for generalized neural network nonlinearities.

In this paper, we address the above problem by deriving an exact analytic solution for the covariance matrix of nonlinear activation functions with Gaussian inputs. Our covariance theorem, presented in Section 3, is a general solution that can be computed to arbitrary precision, enabling more accurate and widely applica-




Figure 1: Moment propagation for a single network layer with an affine transformation followed by a non-linear activation function. Layer-by-layer mean and covariance statistics can be used as a measure of uncertainty traveling across a network.

ble moment propagation in neural networks. As illustrative examples, we show how our solution computes the covariance for neural networks with Heaviside, rectified linear unit (ReLU), and Gaussian error linear unit (GELU) activation functions. With these examples we also derive other activation statistics missing from the current literature. In Section 4, we demonstrate the application of these theoretical results to both the analysis and synthesis of neural networks, showing in experiments improved accuracy for characterizing input-output distributions of trained neural networks and for training Bayesian neural networks (see Table 3). Code to run these experiments is publicly available¹.

2 RELATED WORK

Below we discuss related work in the quantification of neural network uncertainty.

2.1 Uncertainty Propagation

The earliest work on propagating uncertainty through neural networks appears in research on Bayesian belief networks (Pearl, 1988; Neal, 1992). Frey and Hinton (1999) propose a framework for transformations of Gaussian variables in nonlinear Gaussian belief networks. Gast and Roth (2018) build upon this framework and propose a method to propagate mean and variance statistics for training neural networks with assumed density filtering (Boyen and Koller, 1998; Minka, 2001) and probabilistic output layers. This prior work models the predictive uncertainty of a neural network, but is limited by the fact that only variance, not covariance, is propagated (i.e., only the diagonal of a covariance matrix), due to the lack of exact covariance solutions.

Another application of uncertainty propagation is the measurement of adversarial robustness, which is commonly done by evaluating adversarial accuracy, i.e.,

the network test accuracy under selected adversarial attacks. Evaluating adversarial accuracy is usually done with attacks that have proven empirically effective, such as those proposed in Carlini and Wagner (2017) and Madry et al. (2018). However, because one must select and implement a particular set of attacks, adversarial accuracy is a coarse and sometimes misleading measure of robustness (Ding et al., 2019; Olivier and Raj, 2023). Bibi et al. (2018) propose a technique to propagate moments through ReLU networks via piecewise-linear approximations (PL-DNN) to give a probabilistic model of noise and adversarial robustness, which they use to measure adversarial fooling rates and pixel-level noise sensitivity in computer vision problems. The authors give a closed-form solution for the covariance of Gaussian random variables passed through a ReLU activation, but only for zero-mean Gaussians.

In other work, Astudillo and da Silva Neto (2011) and Abdelaziz et al. (2015) combine approximations with sampling-based methods to propagate uncertainty through networks for automatic speech recognition. Daunizeau (2017) derives accurate approximations for the mean and variance of the sigmoid and softmax functions. Seo et al. (2021) show how the distribution of the penultimate layer of a neural network can be used for generative modeling and knowledge distillation.

2.2 Bayesian Neural Networks

Bayesian inference with neural networks was first proposed by MacKay (1992), modeling both input and parameter uncertainty. Because a BNN computes a full posterior distribution, it is robust to overfitting and can quantify predictive uncertainty. However, because modern neural networks are high-dimensional and non-convex, computing posterior distributions incurs a prohibitive computational cost and typically requires approximative, sampling-based methods (Graves, 2011; Blundell et al., 2015; Hernández-Lobato and Adams, 2015; Kingma et al., 2015).

Due to this cost and difficulty, some approximative methods avoid directly representing the uncertainty of the parameters (Gal and Ghahramani, 2016; Maddox et al., 2019). Loquercio et al. (2020) combine the assumed density filtering approach of Gast and Roth (2018) with the dropout sampling of Gal and Ghahramani (2016) and apply their framework to autonomous systems. Most closely related to our work is Wu et al. (2019), which proposes an alternative, sample-free approach to Bayesian inference, deterministic variational inference (DVI), which analytically propagates the mean and covariance statistics of inputs, activations, and parameters. To compute the activation co-

¹<https://github.com/omwright/cov-prop-nn>

variance matrix, the authors propose closed-form approximations of the covariance of Heaviside and ReLU activations. Mae et al. (2021) build on the approximations of Daunizeau (2017) and Wu et al. (2019) to compute uncertainty for dropout-based Bayesian neural networks. The reliance on activation-specific covariance approximations or variance alone limits the accuracy and applicability of these approaches.

3 MOMENT PROPAGATION

In this section we describe moment propagation in neural networks and propose an analytic solution. We denote scalars, vectors, and matrices by x , \mathbf{x} , and \mathbf{X} , respectively. We do not distinguish between random variables and their instantiations, in part because we may treat a given variable as random or deterministic depending upon the context. For economy of notation, we avoid subscripts and superscripts except when necessary.

3.1 Problem Formulation

Given some uncertainty at the input to a neural network, whether due to noise, adversarial perturbations, or otherwise—or possibly at some intermediate stage of the network—our goal is to compute the statistics or probability distribution of the network output. More explicitly, we treat inputs probabilistically and propagate their mean and covariance layer by layer, as depicted in Figure 1, and use these moments to determine output (or intermediate) uncertainty. Each layer ℓ in a standard multilayer perceptron takes an input $\mathbf{x}^{(\ell)} \in \mathbb{R}^n$ and is composed of an affine transformation

$$\mathbf{y}^{(\ell)} = \mathbf{W}^{(\ell)}\mathbf{x}^{(\ell)} + \mathbf{b}^{(\ell)}$$

followed by a nonlinear, element-wise activation function

$$\mathbf{z}^{(\ell)} = \mathbf{g}^{(\ell)}\left(\mathbf{y}^{(\ell)}\right).$$

The output of the activation function then serves as the input to the affine transformation of the next layer, $\mathbf{x}^{(\ell+1)} = \mathbf{z}^{(\ell)}$.

As we will show, we can compute layer-wise mean and covariance across this structure under certain assumptions, giving us a general, sample-free framework for propagating uncertainty through neural networks. This can be used to analyze the robustness of trained networks (Bibi et al., 2018), or to train probabilistic neural networks that can represent predictive uncertainty (Hernández-Lobato and Adams, 2015; Gast and Roth, 2018; Wu et al., 2019).

3.2 Activation Covariance

To make the problem tractable, we assume that the input to each nonlinearity is Gaussian, $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{y}}, \boldsymbol{\Sigma}_{\mathbf{y}})$. Hernández-Lobato and Adams (2015); Bibi et al. (2018); Gast and Roth (2018); Wu et al. (2019) make similar Gaussian assumptions in their moment propagation algorithms. Hernández-Lobato and Adams (2015); Gast and Roth (2018) use assumed density filtering (Boyen and Koller, 1998; Minka, 2001), which requires parametric distributions like the Gaussian. Wu et al. (2019) observe that, empirically, for the high dimensions and conditions typical of practical neural networks, the Gaussian assumption can be loosely justified by the central limit theorem, regardless of the shape of the distribution of \mathbf{x} . The Gaussian assumption may also be justified from the principle of maximum entropy (Jaynes, 2003).

Even if the distribution of the activation input \mathbf{y} is Gaussian, computing the moments of the activation output \mathbf{z} is non-trivial. For the covariance matrix $\boldsymbol{\Sigma}_{\mathbf{z}}$ in particular, no exact solution exists in the literature, and existing methods use activation-specific approximations (see, e.g., Bibi et al. (2018); Wu et al. (2019)). We present a general method to compute the covariance between functions of Gaussian random variables below in Theorem 1, a complete proof of which can be found in Appendix A.

Definition 1. A vector function $\mathbf{g} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is element-wise independent and identical if

$$\mathbf{g}(\mathbf{y}) = [g(y_1) \quad g(y_2) \quad \dots \quad g(y_n)]^\top$$

with $g : \mathbb{R} \rightarrow \mathbb{R}$.

Theorem 1. Suppose $\mathbf{y} \in \mathbb{R}^n$ is a multivariate Gaussian random vector, $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, and $\mathbf{z} = \mathbf{g}(\mathbf{y})$ is an element-wise independent and identical function of \mathbf{y} . If the mean and variance of any input element y_i is (μ_i, σ_i^2) , and any two input elements y_i, y_j are related by Pearson correlation coefficient $\rho_{ij} = \mathbb{E}[(y_i - \mu_i)(y_j - \mu_j)]/(\sigma_i\sigma_j)$, then the covariance between output elements z_i and z_j is

$$\text{Cov}(z_i, z_j) = \sum_{k=1}^{\infty} \frac{\rho_{ij}^k}{k!} \left(\sigma_i^k \frac{\partial^k \mathbb{E}[z_i]}{\partial \mu_i^k} \right) \left(\sigma_j^k \frac{\partial^k \mathbb{E}[z_j]}{\partial \mu_j^k} \right).$$

Corollary 1.1. Given $z = g(y)$, a function of a Gaussian random variable $y \sim \mathcal{N}(\mu, \sigma^2)$, the variance of z is

$$\text{Var}(z) = \sum_{k=1}^{\infty} \frac{1}{k!} \left(\sigma^k \frac{\partial^k \mathbb{E}[z]}{\partial \mu^k} \right)^2.$$

Theorem 1 is a general solution that can be computed to arbitrary precision. Computing the covariance matrix of $\mathbf{z} \in \mathbb{R}^n$ with a k -th order expansion of

Theorem 1 has an asymptotic runtime complexity of $O(kn^2)$. Theorem 1 is not limited to any particular activation function, assuming only the Gaussianity of the input and that the activation mean and its derivatives are well-defined. As we show below in Section 3.3, we can readily apply Theorem 1 to several common activation functions and successfully propagate across network layers.

3.3 Propagation Across a Network Layer

The input and output of a neural network layer are related by

$$\mathbf{y} = \mathbf{W}\mathbf{x} + \mathbf{b} \quad (1)$$

$$\mathbf{z} = \mathbf{g}(\mathbf{y}) \quad (2)$$

where \mathbf{g} takes the element-wise independent and identical form given in Definition 1. If the input \mathbf{x} has mean $\boldsymbol{\mu}_x$ and covariance $\boldsymbol{\Sigma}_x$, the mean and covariance of \mathbf{y} are

$$\boldsymbol{\mu}_y = \mathbf{W}\boldsymbol{\mu}_x + \mathbf{b} \quad (3)$$

$$\boldsymbol{\Sigma}_y = \mathbf{W}\boldsymbol{\Sigma}_x\mathbf{W}^\top. \quad (4)$$

We next compute the mean and covariance of \mathbf{z} in terms of the mean and covariance \mathbf{y} , which, together with Equations (3) and (4) relate the output moments to the input moments.

For simplicity, consider from \mathbf{z} a single element, $z = g(y)$, $y \sim \mathcal{N}(\mu, \sigma^2)$. With the standard normal Gaussian probability density function and standard normal cumulative density function

$$\phi(x) = \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{x^2}{2}\right\}$$

$$\Phi(x) = \frac{1}{2} \left[1 + \operatorname{erf}\left(\frac{x}{\sqrt{2}}\right) \right]$$

the mean of the output for Heaviside, ReLU, and GELU activation functions can be expressed compactly as

$$\text{Heaviside} \quad \mathbb{E}[z] = \Phi\left(\frac{\mu}{\sigma}\right) \quad (5)$$

$$\text{ReLU} \quad \mathbb{E}[z] = \mu\Phi\left(\frac{\mu}{\sigma}\right) + \sigma\phi\left(\frac{\mu}{\sigma}\right) \quad (6)$$

$$\text{GELU} \quad \mathbb{E}[z] = \mu\Phi\left(\frac{\mu}{\sqrt{1+\sigma^2}}\right) + \frac{\sigma^2}{\sqrt{1+\sigma^2}}\phi\left(\frac{\mu}{\sqrt{1+\sigma^2}}\right) \quad (7)$$

and the first derivative with respect to μ is

$$\text{Heaviside} \quad \frac{\partial \mathbb{E}[z]}{\partial \mu} = \frac{1}{\sigma}\phi\left(\frac{\mu}{\sigma}\right) \quad (8)$$

$$\text{ReLU} \quad \frac{\partial \mathbb{E}[z]}{\partial \mu} = \Phi\left(\frac{\mu}{\sigma}\right) \quad (9)$$

$$\text{GELU} \quad \frac{\partial \mathbb{E}[z]}{\partial \mu} = \Phi\left(\frac{\mu}{\sqrt{1+\sigma^2}}\right) + \frac{1}{(1+\sigma^2)^{3/2}}\phi\left(\frac{\mu}{\sqrt{1+\sigma^2}}\right). \quad (10)$$

Equations (5), (6), and (7) allow us to compute $\boldsymbol{\mu}_z$ in terms of the moments of \mathbf{y} , and with their derivatives we can apply Theorem 1 to compute $\boldsymbol{\Sigma}_z$. Equations (5) and (6) can be derived with direct integration (Gradshteyn and Ryzhik, 1994), and in the context of neural networks solutions first appear in Frey and Hinton (1999). We derive Equation (7), and give formulae to compute the k -th derivative of Equations (5), (6), and (7), respectively, in Appendix B.

We again note that Theorem 1 assumes the activation mean and its derivatives are well-defined, and applies only to activation functions that take the element-wise independent and identical form given in Definition 1. Some activation functions, like the sigmoid, don't have a closed-form mean, but accurate approximations exist (Daunizeau, 2017). We discuss the sigmoid function in detail in Appendix B.4. Some common output layers like the softmax function don't meet the element-wise independent and identical requirement, but also do not present a major obstacle since moments can still be computed for output logits. However, extensions are needed to admit other functions frequently used at intermediate layers, like pooling and batch normalization. Shriver (2022) shows that such functions can be refactored into sequences of linear and ReLU operations, to which the moment propagation techniques above can be applied. (Convolution, a linear operation, can be covered straightforwardly by Equations (3) and (4).) For the case of pooling, Gast and Roth (2018) approximate the mean and variance with closed-form approximations. As an alternative, Springenberg et al. (2015) show that pooling layers can be replaced with strided convolutions without loss of inference accuracy.

3.4 Propagation Across Neural Networks

Since $\mathbf{x}^{(\ell+1)} = \mathbf{z}^{(\ell)}$, the moment propagation equations in Section 3.3 yield the following algorithm, versions of which appear in Frey and Hinton (1999), Gast and Roth (2018), and Wu et al. (2019):

Algorithm 1 lets us analytically estimate the output mean and covariance of a neural network in terms of

Algorithm 1 Neural network moment propagation

-
- 1: Get input mean $\boldsymbol{\mu}_x$ and covariance $\boldsymbol{\Sigma}_x$
 - 2: **for** each network layer **do**
 - 3: $\boldsymbol{\mu}_y, \boldsymbol{\Sigma}_y \leftarrow \mathbf{W}\boldsymbol{\mu}_x + \mathbf{b}, \mathbf{W}\boldsymbol{\Sigma}_x\mathbf{W}^\top$
 - 4: Compute $\boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z$ from $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_y, \boldsymbol{\Sigma}_y)$
 - 5: $\boldsymbol{\mu}_x, \boldsymbol{\Sigma}_x \leftarrow \boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z$
 - 6: **end for**
-

the input. Output mean and covariance can be used to analyze the robustness of a neural network to noise, adversarial attacks, and other perturbations, as demonstrated by Bibi et al. (2018).

3.5 Training Neural Networks

With moment propagation, we can train neural networks that achieve high inference accuracy, as with standard neural network training, and represent output uncertainty due to input uncertainty (Gast and Roth, 2018; Loquercio et al., 2020). This uncertainty is sometimes defined as *aleatoric* uncertainty—uncertainty inherent to the underlying process or system. By altering standard output layers and loss functions—e.g., for regression problems, we can replace mean-squared-error loss with a log-likelihood loss to capture both mean and variance of the output—we can use moment propagation in a supervised training regime to characterize the aleatoric uncertainty of the model’s predictions. In their probabilistic training, Gast and Roth (2018) and Loquercio et al. (2020) only model variance for each neuron, i.e., the diagonal of the full activation covariance matrix $\boldsymbol{\Sigma}_z$. Using Algorithm 1 will capture the interactions between a layer’s activation random variables, at the cost of added computational complexity. For a fuller discussion of training probabilistic networks of this form, we refer readers to Gast and Roth (2018).

We can also use moment propagation in the training of fully Bayesian neural networks (Wu et al., 2019). As discussed in Section 2.2, BNNs represent both aleatoric uncertainty and *epistemic* uncertainty—uncertainty in the model, which can be resolved with additional training data—by attempting to characterize the entire predictive distribution instead of learning a single estimate of network parameters. For a dataset of D input-output pairs $\mathcal{D} = \{(\mathbf{x}^{(i)}, \mathbf{z}^{(i)})\}_{i=1}^D$, and a model parameterized by $\boldsymbol{\theta}$, the predictive distribution is (MacKay, 2002):

$$p(\mathbf{z}|\mathbf{x}, \mathcal{D}) = \int_{\boldsymbol{\theta}} p(\mathbf{z}|\mathbf{x}, \boldsymbol{\theta})p(\boldsymbol{\theta}|\mathcal{D})d\boldsymbol{\theta}. \quad (11)$$

Because computing the posterior distribution is intractable for modern neural networks, approximate methods like Monte Carlo variational inference

(MCVI) are typically used (Graves, 2011; Kingma et al., 2015; Wu et al., 2019). MCVI methods minimize the Kullback-Leibler (KL) divergence D_{KL} between a parameterized distribution $q_\phi(\boldsymbol{\theta})$ and the true posterior $p(\boldsymbol{\theta}|\mathcal{D})$, where ϕ represents the parameters of the variational distribution. Because the true posterior is unknown, MCVI methods instead maximize a surrogate objective, the evidence lower bound (ELBO):

$$L(\phi) = \mathbb{E}_{q_\phi(\boldsymbol{\theta})} [\log p(\mathcal{D}|\boldsymbol{\theta})] - D_{KL}(q_\phi(\boldsymbol{\theta})\|p(\boldsymbol{\theta})). \quad (12)$$

By choosing a tractable prior $p(\boldsymbol{\theta})$ and variational distribution $q_\phi(\boldsymbol{\theta})$ such as factorized (i.e., diagonal covariance matrix) Gaussians, one can iteratively optimize the ELBO with gradient ascent, similar to the process of standard neural network training. However, the first term in Equation (12), sometimes called the reconstruction term, can’t be computed analytically for modern neural networks, and Monte Carlo sampling is used instead.

Wu et al. (2019) present an alternative, sample-free variational inference method, deterministic variational inference (DVI), which they use to estimate Equation (12) analytically. DVI uses a version of Algorithm 1 to approximate MCVI in the limit of infinite samples, $s \rightarrow \infty$. To propagate through a single network layer with an $n \times n$ transform, MCVI has an asymptotic runtime complexity of $O(sn^2)$, and DVI $O(n^3)$. For an affine transformation $\mathbf{y} = \mathbf{W}\mathbf{x} + \mathbf{b}$, Wu et al. (2019) show that, under assumptions of independence between $\mathbf{x}, \mathbf{W}, \mathbf{b}$,

$$\begin{aligned} \text{Cov}(y_i, y_k) &= \text{Cov}(b_i, b_k) + \\ &\sum_j \sum_\ell \mathbb{E}[x_j x_\ell] \text{Cov}(W_{ij}, W_{k\ell}) + \\ &\mathbb{E}[W_{ij}] \mathbb{E}[W_{k\ell}] \text{Cov}(x_j, x_\ell). \end{aligned} \quad (13)$$

Because \mathbf{W} and \mathbf{b} are often chosen to be factorized Gaussians for tractability, the only difficult terms to model in Equation (13) are those involving $\text{Cov}(x_j, x_\ell)$. ($\text{Cov}(W_{ij}, W_{k\ell}) = 0$ for $j \neq \ell$.) To compute $\text{Cov}(x_j, x_\ell)$ in Equation (13), Wu et al. (2019) rely on activation-specific approximations of Heaviside and ReLU. We can instead apply Theorem 1 at no additional asymptotic runtime cost, and extend DVI to any supported activation function.

4 EXPERIMENTS

In experiments, we numerically assess the accuracy of Theorem 1 in Section 4.1, apply it to the analysis of trained neural networks in Section 4.2, and apply it to training Bayesian neural networks in Section 4.3. Error analysis in Section 4.1 was performed on an Apple M1 Max, and subsequent experiments were performed on an Nvidia Titan RTX GPU.




Figure 2: The covariance between a bivariate Gaussian passed through a ReLU, calculated numerically, plotted over input means μ_1 and μ_2 , with fixed input parameters $\sigma_1 = 1, \sigma_2 = 1, \rho = 0.5$.




Figure 3: The (a) maximum and (b) mean absolute error of Theorem 1 applied to a ReLU by Taylor order. Error is determined by comparing against the numerical calculation depicted in Figure 2, over $\mu_1, \mu_2 \in [-5, 5]$.

4.1 Error Analysis

Figure 2 shows the covariance of a bivariate Gaussian passed through a ReLU, calculated numerically. We can compare this numerical calculation to the analytic solution of Theorem 1. Figure 3 shows the error by Taylor order, and Figure 4 compares the error with the approximation proposed by Wu et al. (2019). The Wu et al. (2019) approximation has a maximum absolute error of 1.160×10^{-1} . A first-order approximation of Theorem 1 has a maximum absolute error of 2.034×10^{-2} , and a fourth-order approximation 3.740×10^{-4} . We can similarly analyze the error of Theorem 1 applied to other activation functions for which covariance approximations don’t appear in previous work, like the GELU or the approximated sigmoid (Daunizeau, 2017), both shown in Figure 5. The GELU and sigmoid are discussed more fully in Appendix B.




Figure 4: Absolute error for (a) a fourth-order expansion of Theorem 1 applied to a ReLU and (b) the ReLU approximation proposed by Wu et al. (2019), plotted over input means μ_1 and μ_2 . Error is determined by comparing against the numerical calculation depicted in Figure 2.




Figure 5: Absolute error of a fourth-order expansion of Theorem 1 applied to (a) a GELU and (b) an approximated sigmoid (Daunizeau, 2017), plotted over input means μ_1 and μ_2 and using the fixed input parameters of Figure 2.

4.2 Characterizing Trained Neural Networks

We study how well moment propagation can characterize the input-output distribution of a trained network by comparing to Monte Carlo estimation.

Following the experiments in Bibi et al. (2018), we first test our method on synthetic data and networks. We construct both fully connected and convolutional networks of varying depth, with ReLU activations throughout and i.i.d. network parameters generated with Kaiming initialization (He et al., 2015). Input data are \mathbb{R}^{100} for the fully connected networks and $\mathbb{R}^{20 \times 20}$ for the convolutional networks, and input noise is represented with a multivariate Gaussian distribution. In the fully connected networks, each hidden layer is 100 units. In the convolutional networks, each convolutional layer is 10 channels with 3×3 filter size and “same” padding to preserve feature map dimensions. The mean is drawn from a standard normal Gaussian, and the covariance is based on randomly

Table 1: Tightness measures for synthetic data and networks of varying type and depth, comparing the proposed moment propagation method to PL-DNN (Bibi et al., 2018). We report the mean and standard deviation of the ratio of mean estimates $\mathcal{Q}\hat{\mu} = \hat{\mu}_{MC}/\hat{\mu}_A$ and ratio of variance estimates $\mathcal{Q}\hat{\sigma}^2 = \hat{\sigma}_{MC}^2/\hat{\sigma}_A^2$ between the Monte Carlo and analytic methods, calculated over 200 instances. Ideally the ratio approaches 1 ± 0 .

Network	Proposed		PL-DNN	
	$\mathcal{Q}\hat{\mu}$	$\mathcal{Q}\hat{\sigma}^2$	$\mathcal{Q}\hat{\mu}$	$\mathcal{Q}\hat{\sigma}^2$
FC-4	1.000 ± 0.014	1.010 ± 0.011	0.999 ± 0.643	1.155 ± 0.311
FC-8	1.000 ± 0.001	1.016 ± 0.017	1.012 ± 0.169	1.067 ± 0.419
CNN-4	1.001 ± 0.021	1.007 ± 0.006	0.997 ± 0.168	1.222 ± 0.240
CNN-8	1.000 ± 0.000	1.009 ± 0.007	1.001 ± 0.026	1.184 ± 0.347

Table 2: Tightness measures for a convolutional network trained on MNIST data, comparing the proposed moment propagation method to PL-DNN (Bibi et al., 2018). For each of the 10 class logits, we report the mean and standard deviation of the ratio of mean estimates $\mathcal{Q}\hat{\mu} = \hat{\mu}_{MC}/\hat{\mu}_A$ and ratio of variance estimates $\mathcal{Q}\hat{\sigma}^2 = \hat{\sigma}_{MC}^2/\hat{\sigma}_A^2$ between the Monte Carlo and analytic methods, calculated over 200 instances.

Logit	Proposed		PL-DNN	
	$\mathcal{Q}\hat{\mu}$	$\mathcal{Q}\hat{\sigma}^2$	$\mathcal{Q}\hat{\mu}$	$\mathcal{Q}\hat{\sigma}^2$
0	1.000 ± 0.000	1.001 ± 0.007	1.016 ± 0.102	1.422 ± 0.229
1	1.001 ± 0.001	1.004 ± 0.005	0.971 ± 0.185	1.553 ± 0.254
2	1.001 ± 0.001	1.009 ± 0.008	0.998 ± 0.267	1.531 ± 0.222
3	0.999 ± 0.007	1.009 ± 0.008	1.018 ± 0.410	1.512 ± 0.252
4	1.000 ± 0.001	1.005 ± 0.006	0.998 ± 0.141	1.383 ± 0.244
5	1.001 ± 0.015	1.006 ± 0.006	0.980 ± 0.696	1.733 ± 0.267
6	1.001 ± 0.014	1.004 ± 0.007	0.998 ± 0.096	1.382 ± 0.217
7	1.000 ± 0.003	1.009 ± 0.009	0.975 ± 0.178	1.500 ± 0.161
8	1.000 ± 0.002	1.005 ± 0.006	1.068 ± 0.684	1.453 ± 0.205
9	1.000 ± 0.002	1.004 ± 0.005	0.993 ± 0.043	1.557 ± 0.297

drawn eigenvalues such that the maximum variance is 1.0. All networks output a single scalar value.

In each experiment, we generate a random mean and covariance for the input distribution, from which we draw 7.5×10^4 Monte Carlo samples to pass through the network and form an estimate of the output mean and variance. We then analytically propagate the input mean and covariance using both our proposed method and the PL-DNN technique proposed in Bibi et al. (2018). Each experiment is repeated 200 times.

Table 1 shows tightness measures for fully connected and convolutional networks with 4 and 8 layers (denoted as FC-4, FC-8, CNN-4, and CNN-8). The tightness of the analytic methods is measured by taking the ratio of the Monte Carlo estimates of mean and variance to the respective analytic estimates. (We use $\mathcal{Q}\hat{\mu}$ to represent the ratio of estimates as one might use $\Delta\hat{\mu}$ to represent the difference of estimates.) We see that the proposed method achieves high estimation accuracy for both mean and variance, with a better

average tightness and lower standard deviation than PL-DNN. We conjecture that this is because, while both methods approximate unknown distributions as Gaussian, PL-DNN approximates the known structure of the network whereas our method preserves it.

The weights of a trained neural network are not random, of course, and training may induce correlations between network parameters. As in Bibi et al. (2018), we also compare methods using a network trained on MNIST data (LeCun and Cortes, 2010) with a convolutional architecture based on LeNet (LeCun et al., 1999). Because our proposed technique doesn’t admit pooling layers without extension, we replace pooling layers with strided convolutions (Springenberg et al., 2015). The network has two convolutional layers with ReLU activations, each followed by a strided convolution in place of a pooling layer, then two fully connected layers. The network input is $\mathbb{R}^{28 \times 28}$ corresponding to an image, and the output is a set of 10 logits corresponding to the 10 possible classes. For

Table 3: Test log-likelihoods for UCI regression datasets. We compare the proposed method to DVI (Wu et al., 2019) and MCVI, keeping the training routine and model architecture otherwise identical. For each dataset we report the mean and standard deviation for 20 randomized data splits. A higher test log-likelihood is better, and best results are in boldface.

Dataset			Method		
Name	Samples	Features	Proposed	DVI	MCVI
boston	506	13	-2.31 ± 0.19	-2.33 ± 0.19	-2.40 ± 0.23
concrete	1,030	8	-2.98 ± 0.11	-2.99 ± 0.11	-3.04 ± 0.10
energy	768	8	-1.27 ± 0.21	-1.30 ± 0.22	-1.40 ± 0.18
kin8nm	8,192	8	1.11 ± 0.03	1.11 ± 0.03	1.20 ± 0.03
naval	11,934	16	5.81 ± 0.16	5.73 ± 0.22	5.92 ± 0.24
power	9,568	4	-2.82 ± 0.04	-2.82 ± 0.04	-2.82 ± 0.04
protein	45,730	9	-2.97 ± 0.01	-2.97 ± 0.02	-3.05 ± 0.02
wine	1,588	11	-0.91 ± 0.07	-0.92 ± 0.07	-0.94 ± 0.07
yacht	308	6	-0.25 ± 0.20	-0.23 ± 0.22	-0.53 ± 0.19

each experiment, we represent the input as zero-mean Gaussian noise with a randomly generated covariance matrix added to an image \mathbf{M} randomly selected from the MNIST test set. As in the synthetic experiments, we generate 7.5×10^4 Monte Carlo samples to estimate output mean and variance, analytically propagate input mean and covariance using both the proposed method and PL-DNN, and repeat each experiment 200 times.

Table 2 shows tightness measures for each output logit, computed in the same manner as in the experiments on synthetic data. Similarly, the proposed method achieves a consistently high estimation accuracy for both mean and variance, with a better average tightness and lower standard deviation than PL-DNN.

4.3 Training Bayesian Neural Networks

In our training experiments, we train small BNNs on UCI regression datasets (Kelly et al., 2023). We compare different implementations of variational inference: baseline Monte Carlo variational inference (MCVI), deterministic variational inference (DVI) as proposed by Wu et al. (2019), and our version of DVI which uses Theorem 1 to calculate ReLU covariance more accurately. We manually tune the learning rate, but otherwise use the same training routine, hyperparameters, and model architecture across the different datasets and implementations. The model is heteroskedastic (i.e., it outputs both a mean and variance) and consists of a single hidden layer of 50 units with fixed priors and ReLU activations (except for the `protein` dataset, for which we use 100 units). Following Wu et al. (2019), we choose a factorized Gaussian for our variational distribution (i.e., the covariance matrix of

trainable parameters is diagonal) to reduce computational complexity. We use the AdamW optimizer (Loshchilov and Hutter, 2019), the batch size is fixed at 10 and training is run for 50 epochs (except for the larger `protein` dataset, which is run for 25). For MCVI, the number of Monte Carlo samples is fixed at 10. Each dataset is randomly split into 90% training data and 10% test data, and each training experiment is repeated over 20 randomized data splits.

In Table 3 we report the mean and standard deviation of the best test log-likelihood—the reconstruction term in Equation (12)—for each dataset and method. On average, the more accurate covariance calculation from Theorem 1 modestly improves BNN training, but this effect is not uniform, and in some cases performance drops. We conjecture that in such cases the ReLU covariance approximation used by Wu et al. (2019) is sufficiently accurate, or that the factorized Gaussian assumption curbs the impact of accurate covariance calculations and other aspects of the learning problem predominate. We also note that the DVI and MCVI results do not exactly match those reported in Wu et al. (2019), as we run the training experiments with our own BNN implementation; we attribute discrepancies to differences in hyperparameter selection.

5 CONCLUSION

We introduced a general, analytic solution to compute activation covariance and illustrated how this can propagate moments across layers in neural networks. Our experiments confirm the validity of this approach. Our moment propagation method can be applied to characterizing the input-output distributions of trained neural networks and to training probabilis-

tic neural networks that model predictive uncertainty, both of which offer a path to neural networks that are quantifiably reliable and robust.

Prominent research questions remain for successful uncertainty quantification in neural networks. Investigating the application of Theorem 1 to Laplace approximation (Daxberger et al., 2021) may be fruitful. Because computing a covariance matrix is of course $\mathcal{O}(n^2)$ in the number of dimensions n , covariance-based methods still face significant scaling challenges. Also germane to this work are studying more thoroughly the effect of activation covariance, and formulating a theory for the applicability of the Gaussian approximation.

Acknowledgements

We thank Anqi Wu and Adel Bibi for their generous correspondence; Eric Heim, Shannon Gallagher, and David Shriver for their comments; and J. Nelson Wright for his insight into Fourier analysis. This material is based upon work funded in part by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center; in part by JST, PRESTO Grant Number JPMJPR2136, Japan; in part by the Department of the Navy, Office of Naval Research, grant number N00014-23-1-2252; in part by Mobility21 National University Transportation Center, which is sponsored by the Department of Transportation; and in part by NSF grant CCF 2327905.

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution. Please see Copyright notice for non-US Government use and distribution. DM24-0157.

References

- Abdelaziz, A. H., Watanabe, S., Hershey, J. R., Vincent, E., and Kolossa, D. (2015). Uncertainty propagation through deep neural networks. In *Proc. Interspeech 2015*, pages 3561–3565.
- Astudillo, R. F. and da Silva Neto, J. P. (2011). Propagation of uncertainty through multilayer perceptrons for robust automatic speech recognition. In *Proc. Interspeech 2011*, pages 461–464.
- Bibi, A., Alfadly, M., and Ghanem, B. (2018). Analytic expressions for probabilistic moments of pl-dnn with Gaussian input. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9099–9107.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. (2015). Weight uncertainty in neural networks. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ICML’15, page 1613–1622. JMLR.org.
- Boyen, X. and Koller, D. (1998). Tractable inference for complex stochastic processes. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, UAI’98, page 33–42, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Carlini, N. and Wagner, D. (2017). Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57, Los Alamitos, CA, USA. IEEE Computer Society.
- Daunizeau, J. (2017). Semi-analytical approximations to statistical moments of sigmoid and softmax mappings of normal variables. *arXiv preprint arXiv:1703.00091*.
- Daxberger, E., Kristiadi, A., Immer, A., Eschenhagen, R., Bauer, M., and Hennig, P. (2021). Laplace Redux - Effortless Bayesian Deep Learning. In *Advances in Neural Information Processing Systems*, volume 34, pages 20089–20103. Curran Associates, Inc.
- Ding, G. W., Lui, K. Y.-C., Jin, X., Wang, L., and Huang, R. (2019). On the sensitivity of adversarial robustness to input data distributions. In *International Conference on Learning Representations*.
- Frey, B. J. and Hinton, G. E. (1999). Variational learning in nonlinear Gaussian belief networks. *Neural Computation*, 11(1):193–213.
- Gal, Y. and Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML’16, page 1050–1059. JMLR.org.
- Gast, J. and Roth, S. (2018). Lightweight probabilistic deep networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3369–3378.
- Gradshteyn, I. S. and Ryzhik, I. M. (1994). *Table of Integrals, Series, and Products*. Academic Press, 5th edition edition.
- Graves, A. (2011). Practical variational inference for neural networks. In Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., and Weinberger, K., editors, *Advances in Neural Information Processing Systems*, volume 24, page 2348–2356. Curran Associates, Inc.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. (2017). On calibration of modern neural networks.

- In *Proceedings of the 34th International Conference on Machine Learning*, pages 1321–1330. PMLR.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1026–1034.
- Hendrycks, D. and Gimpel, K. (2016). Gaussian error linear units (GELUs). *arXiv preprint arXiv:1606.08415*.
- Hernández-Lobato, J. M. and Adams, R. P. (2015). Probabilistic backpropagation for scalable learning of Bayesian neural networks. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML’15*, page 1861–1869. JMLR.org.
- Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge.
- Kelly, M., Longjohn, R., and Nottingham, K. (2023). The UCI Machine Learning Repository.
- Kingma, D. P., Salimans, T., and Welling, M. (2015). Variational dropout and the local reparameterization trick. In Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 28, page 2575–2583. Curran Associates, Inc.
- LeCun, Y. and Cortes, C. (2010). MNIST handwritten digit database.
- LeCun, Y., Haffner, P., Bottou, L., and Bengio, Y. (1999). Object recognition with gradient-based learning. In Forsyth, D., Mundy, J., di Gesu, V., and Cipolla, R., editors, *Shape, Contour and Grouping in Computer Vision*, Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), pages 319–345. Springer Verlag.
- Loquercio, A., Segu, M., and Scaramuzza, D. (2020). A general framework for uncertainty estimation in deep learning. *IEEE Robotics and Automation Letters*, 5(2):3153–3160.
- Loshchilov, I. and Hutter, F. (2019). Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- MacKay, D. J. C. (1992). A practical Bayesian framework for backpropagation networks. *Neural Computation*, 4(3):448–472.
- MacKay, D. J. C. (2002). *Information Theory, Inference & Learning Algorithms*. Cambridge University Press, USA.
- Maddox, W. J., Garipov, T., Izmailov, P., Vetrov, D., and Wilson, A. G. (2019). A simple baseline for Bayesian uncertainty in deep learning. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 13153–13164, Red Hook, NY, USA. Curran Associates Inc.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. (2018). Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*.
- Mae, Y., Kumagai, W., and Kanamori, T. (2021). Uncertainty propagation for dropout-based Bayesian neural networks. *Neural Networks*, 144:394–406.
- Minka, T. P. (2001). *A Family of Algorithms for Approximate Bayesian Inference*. PhD thesis, USA. AAI0803033.
- Neal, R. M. (1992). Connectionist learning of belief networks. *Artificial Intelligence*, 56(1):71–113.
- Olivier, R. and Raj, B. (2023). How many perturbations break this model? Evaluating robustness beyond adversarial accuracy. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J., editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 26583–26598. PMLR.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Francisco (CA).
- Seo, M., Lee, Y., and Kwak, S. (2021). On the distribution of penultimate activations of classification networks. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, pages 1141–1151. PMLR. ISSN: 2640-3498.
- Shriver, D. (2022). *Increasing the Applicability of Verification Tools for Neural Networks*. PhD thesis, USA.
- Springenberg, J. T., Dosovitskiy, A., Brox, T., and Riedmiller, M. A. (2015). Striving for simplicity: The all convolutional net. In Bengio, Y. and LeCun, Y., editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Workshop Track Proceedings*.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. (2014). Intriguing properties of neural networks. In *International Conference on Learning Representations*.
- Wu, A., Nowozin, S., Meeds, E., Turner, R. E., Hernández-Lobato, J. M., and Gaunt, A. L. (2019). Deterministic variational inference for robust Bayesian neural networks. In *International Conference on Learning Representations*.

A COVARIANCE PROPAGATION, PROOF

We recall Theorem 1 and provide a complete derivation below.

Theorem 1. *Suppose $\mathbf{y} \in \mathbb{R}^n$ is a multivariate Gaussian random vector, $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, and $\mathbf{z} = \mathbf{g}(\mathbf{y})$ is an element-wise independent and identical function of \mathbf{y} . If the mean and variance of any input element y_i is (μ_i, σ_i^2) , and any two input elements y_i, y_j are related by Pearson correlation coefficient $\rho_{ij} = \mathbb{E}[(y_i - \mu_i)(y_j - \mu_j)] / (\sigma_i \sigma_j)$, then the covariance between output elements z_i and z_j is*

$$\text{Cov}(z_i, z_j) = \sum_{k=1}^{\infty} \frac{\rho_{ij}^k}{k!} \left(\sigma_i^k \frac{\partial^k \mathbb{E}[z_i]}{\partial \mu_i^k} \right) \left(\sigma_j^k \frac{\partial^k \mathbb{E}[z_j]}{\partial \mu_j^k} \right).$$

Proof. With Gaussian random variables, the integrals of interests are (omitting subscripts unless necessary):

$$\mathbb{E}[z] = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} g(y) \exp \left\{ -\frac{(y - \mu)^2}{2\sigma^2} \right\} dy \quad (14)$$

$$\mathbb{E}[z^2] = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} g(y)^2 \exp \left\{ -\frac{(y - \mu)^2}{2\sigma^2} \right\} dy \quad (15)$$

$$\begin{aligned} \mathbb{E}[z_i z_j] &= \frac{1}{2\pi\sigma_i\sigma_j\sqrt{1-\rho^2}} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(y_i)g(y_j) \\ &\exp \left\{ -\frac{1}{2(1-\rho^2)} \left[\frac{(y_i - \mu_i)^2}{\sigma_i^2} - \frac{2\rho(y_i - \mu_i)(y_j - \mu_j)}{\sigma_i\sigma_j} + \frac{(y_j - \mu_j)^2}{\sigma_j^2} \right] \right\} dy_i dy_j. \end{aligned} \quad (16)$$

We recognize the integral in Equation (14) as a convolution in μ , and make use of the Fourier transform convolution property. Let $G(\xi) = \mathfrak{F}\{g(\mu)\}$, the Fourier transform of $g(\mu)$. The one-dimensional Fourier transforms of interest are:

$$\begin{aligned} g(\mu) &\xrightarrow{\mathfrak{F}} G(\xi) \\ g(\mu)^2 &\xrightarrow{\mathfrak{F}} (G * G)(\xi) \\ \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left\{ -\frac{\mu^2}{\sigma^2} \right\} &\xrightarrow{\mathfrak{F}} \exp \{-2\pi^2\xi^2\sigma^2\}. \end{aligned}$$

Let $H(\xi) = (G * G)(\xi)$. We can express Equations (14) and (15) as inverse Fourier transforms:

$$\mathbb{E}[z] = \mathfrak{F}^{-1} \{ G(\xi) \exp\{-2\pi^2\xi^2\sigma^2\} \} \quad (17)$$

$$\mathbb{E}[z^2] = \mathfrak{F}^{-1} \{ H(\xi) \exp\{-2\pi^2\xi^2\sigma^2\} \}. \quad (18)$$

Similarly, we can treat Equation (16) as a two-dimensional convolution and exploit this to determine $\mathbb{E}[z_i z_j]$. With the two-dimensional Fourier transform of a bivariate Gaussian

$$\frac{1}{2\pi\sigma_i\sigma_j\sqrt{1-\rho^2}} \exp \left\{ -\frac{1}{2(1-\rho^2)} \left[\frac{\mu_i^2}{\sigma_i^2} - \frac{2\rho\mu_i\mu_j}{\sigma_i\sigma_j} + \frac{\mu_j^2}{\sigma_j^2} \right] \right\} \xrightarrow{\mathfrak{F}} \exp \{-2\pi^2[\xi_i^2\sigma_i^2 + 2\rho\xi_i\xi_j\sigma_i\sigma_j + \xi_j^2\sigma_j^2]\}$$

we can express Equation (16) as an inverse Fourier transform:

$$\mathbb{E}[z_i z_j] = \mathfrak{F}^{-1} \left\{ G(\xi_i)G(\xi_j) \exp\{-2\pi^2[\xi_i^2\sigma_i^2 + 2\rho\xi_i\xi_j\sigma_i\sigma_j + \xi_j^2\sigma_j^2]\} \right\}. \quad (19)$$

Using Equations (17) and (19), we can then express the covariance $\text{Cov}(z_i, z_j) = \mathbb{E}[z_i z_j] - \mathbb{E}[z_i]\mathbb{E}[z_j]$ as

$$\text{Cov}(z_i, z_j) = \mathfrak{F}^{-1} \left\{ G(\xi_i) G(\xi_j) \exp \{-2\pi^2\xi_i^2\sigma_i^2\} [\exp \{-4\pi^2\rho\xi_i\xi_j\sigma_i\sigma_j\} - 1] \exp \{-2\pi^2\xi_j^2\sigma_j^2\} \right\}.$$

To solve this inverse Fourier transform, we can express $\exp \{-4\pi^2\rho\xi_i\xi_j\sigma_i\sigma_j\} - 1$ as a Taylor series in ρ :

$$\exp \{-4\pi^2\rho\xi_i\xi_j\sigma_i\sigma_j\} - 1 = \sum_{k=1}^{\infty} \frac{\rho^k}{k!} (-4\pi^2\xi_i\xi_j\sigma_i\sigma_j)^k.$$

We then move the summation out of the inverse Fourier operation:

$$\text{Cov}(z_i, z_j) = \sum_{k=1}^{\infty} \frac{\rho^k}{k!} (-4\pi^2 \sigma_i \sigma_j)^k \mathfrak{F}^{-1} \left\{ \xi_i^k \xi_j^k G(\xi_i) G(\xi_j) \exp \{-2\pi^2 \xi_i^2 \sigma_i^2\} \exp \{-2\pi^2 \xi_j^2 \sigma_j^2\} \right\}.$$

Then, using the separability property of two-dimensional Fourier transforms:

$$\text{Cov}(z_i, z_j) = \sum_{k=1}^{\infty} \frac{\rho^k}{k!} (-4\pi^2 \sigma_i \sigma_j)^k \mathfrak{F}^{-1} \left\{ \xi_i^k G(\xi_i) \exp \{-2\pi^2 \xi_i^2 \sigma_i^2\} \right\} \mathfrak{F}^{-1} \left\{ \xi_j^k G(\xi_j) \exp \{-2\pi^2 \xi_j^2 \sigma_j^2\} \right\}.$$

Using the Fourier differentiation property and Equation (17), we see that

$$\xi^k G(\xi) \exp \{-2\pi^2 \xi^2 \sigma^2\} \xrightarrow{\mathfrak{F}^{-1}} \frac{1}{(2\pi\sqrt{-1})^k} \frac{\partial^k \mathbb{E}[z]}{\partial \mu^k}.$$

Hence, the covariance can be expressed as:

$$\text{Cov}(z_i, z_j) = \sum_{k=1}^{\infty} \frac{\rho_{ij}^k}{k!} \left(\sigma_i^k \frac{\partial^k \mathbb{E}[z_i]}{\partial \mu_i^k} \right) \left(\sigma_j^k \frac{\partial^k \mathbb{E}[z_j]}{\partial \mu_j^k} \right).$$

□

B MOMENTS OF COMMON ACTIVATION FUNCTIONS

B.1 Heaviside

The Heaviside step function $u(y)$ is defined as

$$u(y) = \begin{cases} 0 & \text{if } y < 0 \\ 1 & \text{if } y \geq 0. \end{cases}$$

When the input is a Gaussian random variable $y \sim \mathcal{N}(\mu, \sigma^2)$, the mean and variance of $z = u(y)$, first reported in the context of neural networks by Frey and Hinton (1999), are

$$\mathbb{E}[z] = \Phi\left(\frac{\mu}{\sigma}\right) \tag{20}$$

$$\text{Var}(z) = \Phi\left(\frac{\mu}{\sigma}\right) \left[1 - \Phi\left(\frac{\mu}{\sigma}\right)\right] \tag{21}$$

To compute the covariance of the output of a Heaviside step function using Theorem 1, the first few terms of interest are:

$$\begin{aligned} \sigma \frac{\partial \mathbb{E}[z]}{\partial \mu} &= \phi\left(\frac{\mu}{\sigma}\right) \\ \sigma^2 \frac{\partial^2 \mathbb{E}[z]}{\partial \mu^2} &= -\left(\frac{\mu}{\sigma}\right) \phi\left(\frac{\mu}{\sigma}\right) \\ \sigma^3 \frac{\partial^3 \mathbb{E}[z]}{\partial \mu^3} &= \left[\left(\frac{\mu}{\sigma}\right)^2 - 1\right] \phi\left(\frac{\mu}{\sigma}\right) \\ \sigma^4 \frac{\partial^4 \mathbb{E}[z]}{\partial \mu^4} &= -\left(\frac{\mu}{\sigma}\right) \left[\left(\frac{\mu}{\sigma}\right)^2 - 3\right] \phi\left(\frac{\mu}{\sigma}\right) \\ \sigma^5 \frac{\partial^5 \mathbb{E}[z]}{\partial \mu^5} &= \left[\left(\frac{\mu}{\sigma}\right)^4 - 6\left(\frac{\mu}{\sigma}\right)^2 + 3\right] \phi\left(\frac{\mu}{\sigma}\right). \end{aligned}$$

With the probabilist's Hermite polynomial

$$\text{He}_k(x) = \frac{(-1)^k \partial^k \phi(x)}{\phi(x) \partial x^k}$$

terms can be expressed as

$$\sigma^k \frac{\partial^k \mathbb{E}[z]}{\partial \mu^k} = (-1)^{k-1} \text{He}_{k-1}\left(\frac{\mu}{\sigma}\right) \phi\left(\frac{\mu}{\sigma}\right). \tag{22}$$

B.2 Rectified Linear Unit

The rectified linear unit (ReLU) is defined as

$$\text{ReLU}(y) = \begin{cases} 0 & \text{if } y < 0 \\ y & \text{if } y \geq 0. \end{cases}$$

When the input is a Gaussian random variable $y \sim \mathcal{N}(\mu, \sigma^2)$, the mean and variance of $z = \text{ReLU}(y)$, first reported in the context of neural networks by Frey and Hinton (1999), are

$$\mathbb{E}[z] = \mu\Phi\left(\frac{\mu}{\sigma}\right) + \sigma\phi\left(\frac{\mu}{\sigma}\right) \quad (23)$$

$$\text{Var}(z) = (\mu^2 + \sigma^2)\Phi\left(\frac{\mu}{\sigma}\right) + \mu\sigma\phi\left(\frac{\mu}{\sigma}\right) - \mathbb{E}[z]^2 \quad (24)$$

To compute the covariance of the output of a ReLU using Theorem 1, the first few terms of interest are:

$$\begin{aligned} \sigma \frac{\partial \mathbb{E}[z]}{\partial \mu} &= \sigma\Phi\left(\frac{\mu}{\sigma}\right) \\ \sigma^2 \frac{\partial^2 \mathbb{E}[z]}{\partial \mu^2} &= \sigma\phi\left(\frac{\mu}{\sigma}\right) \\ \sigma^3 \frac{\partial^3 \mathbb{E}[z]}{\partial \mu^3} &= -\sigma\left(\frac{\mu}{\sigma}\right)\phi\left(\frac{\mu}{\sigma}\right) \\ \sigma^4 \frac{\partial^4 \mathbb{E}[z]}{\partial \mu^4} &= \sigma\left[\left(\frac{\mu}{\sigma}\right)^2 - 1\right]\phi\left(\frac{\mu}{\sigma}\right) \\ \sigma^5 \frac{\partial^5 \mathbb{E}[z]}{\partial \mu^5} &= -\sigma\left(\frac{\mu}{\sigma}\right)\left[\left(\frac{\mu}{\sigma}\right)^2 - 3\right]\phi\left(\frac{\mu}{\sigma}\right). \end{aligned}$$

With the probabilist's Hermite polynomial, terms for $k > 1$ can be expressed as

$$\sigma^k \frac{\partial^k \mathbb{E}[z]}{\partial \mu^k} = \sigma(-1)^k \text{He}_{k-2}\left(\frac{\mu}{\sigma}\right)\phi\left(\frac{\mu}{\sigma}\right). \quad (25)$$

B.3 Gaussian Error Linear Unit

The Gaussian Error Linear Unit (GELU), first proposed by Hendrycks and Gimpel (2016), is defined as

$$\text{GELU}(y) = y\Phi(y)$$

When the input is a Gaussian random variable $y \sim \mathcal{N}(\mu, \sigma^2)$, the mean of $z = \text{GELU}(y)$ is

$$\mathbb{E}[z] = \mu\Phi\left(\frac{\mu}{\sqrt{1+\sigma^2}}\right) + \frac{\sigma^2}{\sqrt{1+\sigma^2}}\phi\left(\frac{\mu}{\sqrt{1+\sigma^2}}\right). \quad (26)$$

Unlike the Heaviside and ReLU moments, Equation (26) does not appear in prior work. We provide a derivation below.

Proof. The integral of interest is

$$\mathbb{E}[z] = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} y\Phi(y) \exp\left\{-\frac{(y-\mu)^2}{2\sigma^2}\right\} dy. \quad (27)$$

With the definition of $\Phi(y)$ we can separate this into two integrals. The first integral

$$\frac{1}{2} \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} y \exp\left\{-\frac{(y-\mu)^2}{2\sigma^2}\right\} dy$$

evaluates to $\frac{\mu}{2}$. The second integral

$$\frac{1}{2} \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} y \operatorname{erf}\left(\frac{y}{\sqrt{2}}\right) \exp\left\{-\frac{(y-\mu)^2}{2\sigma^2}\right\} dy$$

can be evaluated via Fourier analysis to obtain an exact solution. We recognize that Equation (27) is a convolution in μ and exploit the Fourier convolution property. The Fourier transforms of interest are:

$$\begin{aligned} \frac{1}{2} \mu \operatorname{erf}\left(\frac{\mu}{\sqrt{2}}\right) &\xrightarrow{\mathfrak{F}} \frac{-(4\pi^2\xi^2 + 1) \exp\{-2\pi^2\xi^2\}}{4\pi^2\xi^2} \\ \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{\mu^2}{2\sigma^2}\right\} &\xrightarrow{\mathfrak{F}} \exp\{-2\pi^2\xi^2\sigma^2\}. \end{aligned}$$

In the Fourier domain, our convolution becomes the product of the transforms:

$$-\frac{4\pi^2\xi^2 + 1}{4\pi^2\xi^2} \exp\{-2\pi^2\xi^2\} \exp\{-2\pi^2\xi^2\sigma^2\} = -\exp\{-2\pi^2\xi^2(1 + \sigma^2)\} - \frac{1}{4\pi^2\xi^2} \exp\{-2\pi^2\xi^2(1 + \sigma^2)\}.$$

This representation gives us two terms to which we can apply the inverse Fourier transform:

$$\begin{aligned} -\exp\{-2\pi^2\xi^2(1 + \sigma^2)\} &\xrightarrow{\mathfrak{F}^{-1}} -\frac{1}{\sqrt{1 + \sigma^2}} \phi\left(\frac{\mu}{\sqrt{1 + \sigma^2}}\right) \\ -\frac{1}{4\pi^2\xi^2} \exp\{-2\pi^2\xi^2(1 + \sigma^2)\} &\xrightarrow{\mathfrak{F}^{-1}} \frac{1}{2} \mu \operatorname{erf}\left(\frac{\mu}{\sqrt{2(1 + \sigma^2)}}\right) + \sqrt{1 + \sigma^2} \phi\left(\frac{\mu}{\sqrt{1 + \sigma^2}}\right) \end{aligned}$$

Combining all results yields:

$$\mathbb{E}[z] = \mu \Phi\left(\frac{\mu}{\sqrt{1 + \sigma^2}}\right) + \frac{\sigma^2}{\sqrt{1 + \sigma^2}} \phi\left(\frac{\mu}{\sqrt{1 + \sigma^2}}\right).$$

□

Let $\alpha = \sigma/\sqrt{1 + \sigma^2}$. We can rewrite Equation (26) as

$$\mathbb{E}[z] = \mu \Phi\left(\frac{\alpha\mu}{\sigma}\right) + \alpha\sigma \phi\left(\frac{\alpha\mu}{\sigma}\right). \quad (28)$$

The variance of a GELU doesn't appear to have a closed-form solution, but can be expressed analytically with Corollary 1.1. To apply Theorem 1 or Corollary 1.1, the first few terms of interest are:

$$\begin{aligned} \sigma \frac{\partial \mathbb{E}[z]}{\partial \mu} &= \sigma \Phi\left(\frac{\alpha\mu}{\sigma}\right) + \alpha(1 - \alpha^2) \mu \phi\left(\frac{\alpha\mu}{\sigma}\right) \\ \sigma^2 \frac{\partial^2 \mathbb{E}[z]}{\partial \mu^2} &= \alpha \left[-\left(\frac{\alpha\mu}{\sigma}\right)^2 (1 - \alpha^2) + (2 - \alpha^2) \right] \sigma \phi\left(\frac{\alpha\mu}{\sigma}\right) \\ \sigma^3 \frac{\partial^3 \mathbb{E}[z]}{\partial \mu^3} &= -\alpha^2 \left[-\left(\frac{\alpha\mu}{\sigma}\right)^2 (1 - \alpha^2) + (4 - 3\alpha^2) \right] \sigma \left(\frac{\alpha\mu}{\sigma}\right) \phi\left(\frac{\alpha\mu}{\sigma}\right) \\ \sigma^4 \frac{\partial^4 \mathbb{E}[z]}{\partial \mu^4} &= \sigma \alpha^3 \left[-\left(\frac{\alpha\mu}{\sigma}\right)^4 (1 - \alpha^2) + \left(\frac{\alpha\mu}{\sigma}\right)^2 (7 - 6\alpha^2) - (4 - 3\alpha^2) \right] \phi\left(\frac{\alpha\mu}{\sigma}\right) \\ \sigma^5 \frac{\partial^5 \mathbb{E}[z]}{\partial \mu^5} &= -\sigma \alpha^4 \left[-\left(\frac{\alpha\mu}{\sigma}\right)^4 (1 - \alpha^2) + \left(\frac{\alpha\mu}{\sigma}\right)^2 (11 - 10\alpha^2) - 3(6 - 5\alpha^2) \right] \left(\frac{\alpha\mu}{\sigma}\right) \phi\left(\frac{\alpha\mu}{\sigma}\right). \end{aligned}$$

With the probabilist's Hermite polynomial, we can express terms for $k > 1$ as

$$\sigma^k \frac{\partial^k \mathbb{E}[z]}{\partial \mu^k} = \alpha^{k-1} \sigma (-1)^k \left[\operatorname{He}_{k-2}\left(\frac{\alpha\mu}{\sigma}\right) - (1 - \alpha^2) \operatorname{He}_k\left(\frac{\alpha\mu}{\sigma}\right) \right] \phi\left(\frac{\alpha\mu}{\sigma}\right). \quad (29)$$

We note that as $\sigma \rightarrow \infty$, $\alpha \rightarrow 1$, Equation (28) reduces to Equation (23), and Equation (29) reduces to Equation (25).

B.4 Sigmoid

Sigmoid functions frustrate many standard analytical tools, and moments of Gaussian random variables passed through sigmoid functions do not appear to have closed-form solutions. The logistic sigmoid is defined as

$$s(y) = \frac{1}{1 + e^{-y}}.$$

For a Gaussian input $y \sim \mathcal{N}(\mu, \sigma^2)$, Daunizeau (2017) proposes the following approximation for the mean of $z = s(y)$:

$$\mathbb{E}[z] \approx s\left(\frac{\mu}{\sqrt{1 + \alpha\sigma^2}}\right) \quad (30)$$

where $\alpha = 0.368$ is an empirically derived parameter. We note in passing that the best-fit α is sensitive to the range of considered input means and variances. For example, we find that $\alpha = \pi/8$, derived from $\tanh(y) \approx \operatorname{erf}\left(\frac{\sqrt{\pi}}{2}y\right)$, is sometimes a better fit.

Let $\beta = \sqrt{1 + \alpha\sigma^2}$. To compute the approximate covariance of the output of a sigmoid function using Theorem 1 and Equation (30), the first few terms of interest are:

$$\begin{aligned} \sigma \frac{\partial \hat{\mathbb{E}}[z]}{\partial \mu} &= \sigma \frac{e^{\mu/\beta}}{\beta (1 + e^{\mu/\beta})^2} \\ \sigma^2 \frac{\partial^2 \hat{\mathbb{E}}[z]}{\partial \mu^2} &= \sigma^2 \frac{e^{\mu/\beta} (-1 + e^{\mu/\beta})}{\beta^2 (1 + e^{\mu/\beta})^3} \\ \sigma^3 \frac{\partial^3 \hat{\mathbb{E}}[z]}{\partial \mu^3} &= \sigma^3 \frac{e^{\mu/\beta} (1 - 4e^{\mu/\beta} + e^{2\mu/\beta})}{\beta^3 (1 + e^{\mu/\beta})^4} \\ \sigma^4 \frac{\partial^4 \hat{\mathbb{E}}[z]}{\partial \mu^4} &= \sigma^4 \frac{e^{\mu/\beta} (-1 + 11e^{\mu/\beta} - 11e^{2\mu/\beta} + e^{3\mu/\beta})}{\beta^4 (1 + e^{\mu/\beta})^5} \\ \sigma^5 \frac{\partial^5 \hat{\mathbb{E}}[z]}{\partial \mu^5} &= \sigma^5 \frac{e^{\mu/\beta} (1 - 26e^{\mu/\beta} + 66e^{2\mu/\beta} - 26e^{3\mu/\beta} + e^{4\mu/\beta})}{\beta^5 (1 + e^{\mu/\beta})^6}. \end{aligned}$$

C ADDITIONAL EXPERIMENT DETAILS

For the MNIST experiments discussed in Section 4.2, the architecture of the trained CNN is summarized in Table 4. As mentioned in Section 4.2, we use strided, activation-less convolutional layers in place of pooling layers (Springenberg et al., 2015) to apply moment propagation.

A visualization of the ratio of covariance estimates, complementing the ratio of variance estimates $\Omega \hat{\sigma}^2 = \hat{\sigma}_{MC}^2 / \hat{\sigma}_A^2$ in Table 2, is shown in Figure 6. Overall, for both the proposed method and PL-DNN, the means are similar to the mean of the corresponding ratio of variance estimates, but the standard deviations are greater.

Table 4: Model summary for MNIST tightness experiments with stride information.

Layer	Type	Stride	Output Size	Activation
1	Convolutional	-	$28 \times 28 \times 32$	ReLU
2	Convolutional	2	$14 \times 14 \times 32$	None
3	Convolutional	-	$14 \times 14 \times 64$	ReLU
4	Convolutional	2	$7 \times 7 \times 64$	None
5	Fully Connected	-	1024	ReLU
6	Fully Connected	-	10	Softmax

For the UCI regression experiments discussed in Section 4.3, the considered hyperparameters are shown in Table 5. As noted in Section 4.3 that the hidden-layer width and number of training epochs is 50 except for the larger `protein` dataset, for which we use 100 hidden units and and train over 25 epochs.




Figure 6: Tightness measures of covariance for a convolutional network trained on MNIST data, comparing the proposed moment propagation method to PL-DNN (Bibi et al., 2018). Complementary to the ratio of variance estimates $\mathfrak{Q}\hat{\sigma}^2 = \hat{\sigma}_{MC}^2 / \hat{\sigma}_A^2$ in Table 2, for the 10 class logits we visualize the (a) mean and (b) standard deviation of the ratio of covariance estimates using moment propagation, and the (c) mean and (d) standard deviation of the ratio of covariance estimates using PL-DNN, calculated over 200 instances. Lower is better.

Table 5: BNN hyperparameters considered for UCI regression tasks.

Hyperparameter	Value
Activation	ReLU
Batch size	10
Hidden layers	1
Hidden-layer width	50*
Lambda	1e-3
Learning rate	{1e-3, 3e-3, 5e-3, 7e-3, 1e-2}
Optimizer	AdamW
Training epochs	50*