

General Uncertainty Estimation with Delta Variances

Simon Schmitt,^{1,2} John Shawe-Taylor,² Hado van Hasselt¹

¹ DeepMind

² University College London, UK
suschmitt@google.com

Abstract

Decision makers may suffer from uncertainty induced by limited data. This may be mitigated by accounting for epistemic uncertainty, which is however challenging to estimate efficiently for large neural networks.

To this extent we investigate Delta Variances, a family of algorithms for epistemic uncertainty quantification, that is computationally efficient and convenient to implement. It can be applied to neural networks and more general functions composed of neural networks. As an example we consider a weather simulator with a neural-network-based step function inside – here Delta Variances empirically obtain competitive results at the cost of a single gradient computation.

The approach is convenient as it requires no changes to the neural network architecture or training procedure. We discuss multiple ways to derive Delta Variances theoretically noting that special cases recover popular techniques and present a unified perspective on multiple related methods. Finally we observe that this general perspective gives rise to a natural extension and empirically show its benefit.

1 Introduction

Decision makers often need to act given limited data. Accounting for the resulting uncertainty (epistemic uncertainty) may be helpful for active learning (MacKay 1992a), exploration (Duff 2002; Auer, Cesa-Bianchi, and Fischer 2002) and safety (Heger 1994).

How to measure epistemic uncertainty efficiently for large neural networks is active research. Computational efficiency is important because even a single evaluation (e.g. a forward pass through a neural network) can be expensive. Popular approaches compute an ensemble of predictions using bootstrapping or MC dropout and incur a multiplicative computational overhead. Other approaches are faster but require changes to the predictors architecture and training procedure (Van Amersfoort et al. 2020).

In this paper we propose the *Delta Variance* family of algorithms which connects and extends Bayesian, frequentist and heuristic notions of variance. Delta Variances require no changes to the architecture or training procedure while incurring the cost of little more than a gradient computation. We present further appealing properties and benefits in Table 1.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The approach can be applied to neural networks or functions that contain neural networks as building blocks to compute a quantity of interest. For instance we could learn a step-by-step dynamics model and then use it to infer some utility function for decision making.

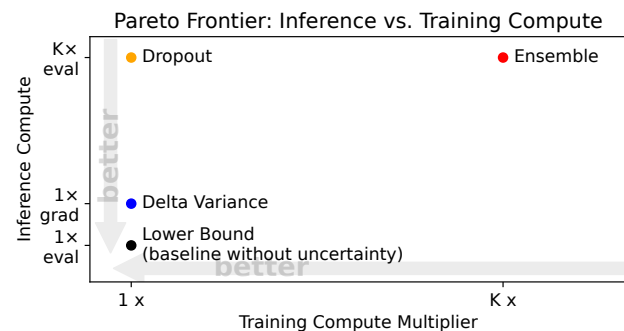


Figure 1: We compare the computational overhead of training and evaluating different variance estimators. Delta Variances are favourable in terms of computational efficiency. They incur negligible training overhead while inference incurs the cost of a regular gradient pass making them more efficient than the alternatives considered. Monte-Carlo Dropout also incurs negligible training overhead, but requires K independent evaluations for inference. Most expensive are Bootstrapped Ensembles requiring $K \times$ repeated computations.

In Section 6 we consider the GraphCast weather forecasting system with a neural network step function (Lam et al. 2023). We then compute the epistemic uncertainty of various derived quantities such as the expected precipitation or wind-turbine-power at a particular location.

Section 5 observes how instances of the Delta Variance family can be derived using different assumptions and theoretical frameworks. We begin with a Bernstein-von Mises plus Delta Method derivation, which relies on strong assumptions. We conclude with an influence function based derivation, which relies only on mild assumptions. Interestingly the resulting instances are not only similar, but also become identical as the number of observed data-points grows.

Formalizing the Delta Variance family allows us to connect Bayesian and frequentist notions of variance, adversarial robustness and anomaly detection in a unified perspective.

	Delta Variances	Ensemble	MC-Dropout
Efficiency			
Inference cost	$1 \times$ gradient	$K \times$ evaluations	$K \times$ evaluations
Training overhead	$1 \times$	$K \times$	$1 \times$
Memory overhead	$2 \times$	$K \times$	$1 \times$
Ease of Use			
No architecture requirements	✓	✓	
No change to training procedure	✓		
Deterministic result	✓	✓	

Table 1: Delta Variances have appealing benefits: The prototypical variant with diagonal Σ is computationally efficient requiring only a gradient pass for inference while other methods evaluate the neural network multiple times. Furthermore they easily build on the existing training procedure and can even be added post hoc after training. Delta Variances do not require architecture changes such as introducing dropout layers or training procedure changes as needed for ensembling or even hyper parameter search. Finally Delta Variances have a simple closed form expression that yields reproducible deterministic results.

This perspective can be used to answer questions such as: What happens if we use a Bayesian variance, but our neural network does not meet all theoretical assumptions – what is a theoretically sound interpretation of the number that we compute? To further highlight the generality of this unified perspective we propose a novel Delta Variance in Section 7.1 for which we observe empirically improvements in Section 6.

When applied to the state-of-the-art GraphCast weather forecasting system we observe favourable results. In comparison to popular related approaches such as ensemble methods our method exhibits similar quality while requiring fewer computational resources.

What is Epistemic Variance? Given limited data the parameters θ of a parametric model f_θ can only be identified with limited certainty. The resulting parameter uncertainty translates into uncertainty in the outputs of $f_\theta(x)$ and any other function $u_\theta(z)$ that depends on θ . We define Epistemic Variance as the output variance induced by the posterior distribution over parameters given the model f and its training data $p(\theta|f, \mathcal{D})$:

$$\mathbb{V}_{\theta \sim p(\theta|f, \mathcal{D})} [f_\theta(x)]$$

Section 3 extends this definition to any function u_θ that depends on parameters θ that were estimated using f_θ and \mathcal{D} . As an illustration consider Section 6 where f_θ is a learned weather dynamics model and u_θ are weather dependent utility functions - e.g. a wind turbine power yield forecast. The epistemic variance of $u_\theta(z)$ is then $\mathbb{V}_{\theta \sim p(\theta|f, \mathcal{D})} [u_\theta(z)]$.

What is a Delta Variance Estimator? Variance estimators of the Delta Variance family are all of the following parametric form:

$$\Delta_{f(x)}^\top \Sigma \Delta_{f(x)}$$

being a vector-matrix-vector product using the gradient vector of f : $\Delta_{f(x)} := \nabla_\theta f_\theta(x)$ while leaving some flexibility in the choice of matrix Σ . In Section 5 and Table 2 we discuss different choices of Σ and their properties. For $\Sigma_f = \frac{1}{N} F_f^{-1}$ being the inverse Fisher information matrix divided by the number of data-points N , it can be shown under suitable conditions that

$$\mathbb{V}_{\theta \sim p(\theta|f, \mathcal{D})} [f_\theta(x)] \approx \Delta_{f(x)}^\top \Sigma_f \Delta_{f(x)}$$

Its name is inspired by the closely related Delta Method (Lambert 1765; Gauss 1823; Doob 1935) – see Gorroochurn (2020) for a historic account – which provides one of many ways to derive Delta Variances.

What is a Quantity of Interest? Sometimes we use a neural network f_θ to learn predictions, that are then used to compute a downstream quantity of interest. For instance we could learn a step-by-step weather dynamics model $f_\theta(x)$ and then use it to infer some utility function for decision making $u_\theta(z)$. Given limited training data neither the neural network’s prediction nor the downstream quantity of interest will be exact. The definition of Epistemic Variance and the Delta Variance family both extend conveniently to quantities of interest u_θ . In fact we only need to replace $\Delta_{f(x)}$ by $\Delta_{u(z)} := \nabla_\theta u_\theta(z)$:

$$\mathbb{V}_{\theta \sim p(\theta|f, \mathcal{D})} [u_\theta(z)] \approx \Delta_{u(z)}^\top \Sigma_f \Delta_{u(z)}$$

Conveniently we can still use the same Σ_f as before. It is independent of u_θ and can be re-used for various quantities of interest.

2 Notation

We consider a function $f_\theta(x)$ with parameters θ trained on a dataset \mathcal{D} of size N . We strive to estimate the uncertainty introduced by training on limited data. To admit Bayesian interpretations we assume that $f_\theta(x)$ is a density model that is trained with log-likelihood (or a function that is trained with a log-likelihood-equivalent loss, such as L_2 regression or cross-entropy). Unless otherwise specified we assume that θ has been trained until convergence – i.e. equals the (local) maximum likelihood estimate $\hat{\theta}$. Under appropriate conditions $\hat{\theta}$ converges to the true distribution parameters θ_{True} . Let H_f be the Hessian of the log-likelihood of all \mathcal{D} evaluated at $\hat{\theta}$. When adopting a Bayesian view with prior belief $p(\theta)$ the posterior over parameters is defined as $p(\theta|\mathcal{D}) \propto p(\mathcal{D}|\theta)p(\theta)$. $\mathbb{E}_{z \sim p(z)}$ refers to the expectation with respect to random variable z with distribution $p(z)$ – which we shorten to \mathbb{E}_z or \mathbb{E} when the distribution is clear. Similarly let $\mathbb{V}[X] := \mathbb{E}[X^2] - \mathbb{E}[X]^2$.

Let $F_f := \mathbb{E}_{x \sim f_{\theta_{\text{True}}}} \nabla_{\theta} \log f_{\theta}(x) \nabla_{\theta} \log f_{\theta}(x)^{\top} |_{\theta=\theta_{\text{True}}}$ be the Fisher information matrix and let $\hat{F}_f := \frac{1}{N} \sum_{x_i \in \mathcal{D}} \nabla_{\theta} \log f_{\theta}(x_i) \nabla_{\theta} \log f_{\theta}(x_i)^{\top} |_{\theta=\bar{\theta}}$ be the empirical Fisher information. Note that H_f^{-1} and \hat{F}_f^{-1}/N are both $O(N^{-1})$. When strong conditions are met (see van der Vaart 1998, for details) the Bernstein-von Mises theorem ensures that the Bayesian posterior converges to the Gaussian distribution $\mathcal{N}(\bar{\theta}, \frac{1}{N} \hat{F}_f^{-1})$ in total variation norm independently of the choice of prior as the number of data-points N increases. In Definition 1 we consider *Quantities of Interest* that we denote u_{θ} . In practice $u_{\theta}(z)$ may be a utility function that depends on some context provided by z . For a simpler exposition but without loss of generality we assume that $u_{\theta}(z)$ is scalar valued. We require f_{θ} and u_{θ} to have bounded second derivatives wrt. θ in order to perform first order Taylor expansions. To simplify notations we assume that θ is evaluated at the learned parameters $\bar{\theta}$ unless specified otherwise: in particular we write $\nabla_{\theta} f_{\theta}(x)$ in place of $\nabla_{\theta} f_{\theta}(x) |_{\theta=\bar{\theta}}$ and $\Delta_{f(x)} := \nabla_{\theta} f_{\theta}(x) |_{\theta=\bar{\theta}}$.

3 Epistemic Variance of Quantities of Interest

Sometimes we use a neural network f_{θ} to learn predictions, that are then used to compute a downstream quantity of interest u_{θ} – see motivational examples below. Given limited training data neither the neural network’s prediction nor the downstream quantity of interest will be exact. This motivates our research question:

If we estimate the parameters θ of u_{θ} by learning $f_{\theta}(x)$, how can we quantify the epistemic uncertainty of $u_{\theta}(z)$?

For a simpler exposition and without loss of generality we assume that $u_{\theta}(z)$ predicts scalar quantities. The prototypical example is a utility function that depends on some context provided by z and internally uses f_{θ} to compute a utility value. The derivations carry over naturally to the multi-variate case. Note that f_{θ} and u_{θ} may have different input spaces. Our research focuses on the general case where $f_{\theta} \neq u_{\theta}$ which has received little attention. This naturally includes the case where $f_{\theta} = u_{\theta}$.

Definition 1. We call the real-valued function $u_{\theta}(z)$ quantity of interest if it depends on the same parameters θ as a related parametric model $f_{\theta}(x)$.

3.1 Motivational Examples

We consider three motivational examples for training on f_{θ} but evaluating a different quantity of interest u_{θ} . We will see that training f_{θ} is straightforward while training a predictor for $u_{\theta}(z)$ is inefficient, impractical, or even impossible.

1. As a simple motivation let us consider estimating the 10-year survival chance using a neural network predictor $f_{\theta}(x)$ of 1-year outcomes given patient features x :

$$u_{\theta}(x) = f_{\theta}(x)^{10}$$

This example illustrates that it may be **impossible to train $u_{\theta}(x)$ directly** unless we collect data for 9 more years, hence we train $f_{\theta}(x)$ and evaluate $u_{\theta}(z)$.

2. Distinct input spaces: $u_{\theta}(z)$ might aggregate predictions of $f_{\theta}(x)$ for sets $z = \{x_1, \dots, x_k\}$: E.g. the survival chance of everyone in set of patients z via $u_{\theta}(z) := \prod_{x_i \in z} f_{\theta}(x_i)$, or the average value of some basket of items z , or the chance of any advertisement from a presented set being clicked. Here training f_{θ} may be more convenient than training u_{θ} .
3. Multiple derived quantities: In Section 6 we compute multiple quantities of interest using the GraphCast weather forecasting system (Lam et al. 2023). Training a separate u_{θ} for each of them would be cumbersome and expensive.

3.2 Epistemic Variance

Here we define Epistemic Variance first from a Bayesian and then from a frequentist perspective. This allows us to formalize and quantify how parameter uncertainty from training f_{θ} translates to uncertainty of any quantity of interest $u_{\theta}(z)$ that also depends on θ .

Bayesian Definition Epistemic uncertainty can be formalized with a Bayesian posterior distribution over parameters given training data: $p(\theta|\mathcal{D})$. The Epistemic Variance of a function evaluation $u_{\theta}(z)$ is then defined to be the variance induced by the posterior over θ :

Definition 2. Given any function u_{θ} and a posterior over parameters $p(\theta|f, \mathcal{D})$ resulting from training f_{θ} on data \mathcal{D} the Epistemic Variance of $u_{\theta}(z)$ is defined as

$$\mathbb{V}_{\theta \sim p(\theta|f, \mathcal{D})} [u_{\theta}(z)]$$

where $\mathbb{V}[X] := \mathbb{E}[X^2] - \mathbb{E}[X]^2$.

Frequentist Definition Leave-one-out cross-validation (Quenouille 1949) is a frequentist counterpart to Epistemic Variance. It computes the variance of $u_{\theta}(z)$ induced by removing a random element from the training data and re-estimating the parameters θ .

Definition 3. Let $\theta_{\setminus i}$ be the leave-one-out parameters resulting from training f_{θ} on data $\mathcal{D} \setminus \{x_i\}$, then the Leave-one-out Variance is defined as

$$\mathbb{V}_{\theta \sim LOO} [u_{\theta}(z)] := \mathbb{V}_{i \sim U(1, \dots, N)} [u_{\theta_{\setminus i}}(z)]$$

where $U(1, \dots, N)$ is the uniform distribution over indices.

4 Delta Variance Approximators

Delta Variance estimators are a family of efficient and convenient approximators of epistemic uncertainty. They can be used to compute the Epistemic Variance of a quantity of interest $u_{\theta}(z)$ where the parameters θ are obtained by learning f_{θ} with limited data. Given any quantity of interest $u_{\theta}(z)$ they approximate both the Bayesian Epistemic Variance as well as the frequentist leave-one-out analogue:

$$\underbrace{\mathbb{V}_{\theta \sim p(\theta|f, \mathcal{D})} [u_{\theta}(z)]}_{\text{Epistemic Variance}} \approx \underbrace{\Delta_{u(z)}^{\top} \Sigma \Delta_{u(z)}}_{\text{Delta Variance}} \approx \underbrace{\mathbb{V}_{\theta \sim LOO} [u_{\theta}(z)]}_{\text{LOO Variance}}$$

Here the *Delta* $\Delta_{u(z)} := \nabla_{\theta} u_{\theta}(z)$ is the gradient vector of u_{θ} evaluated at the input z . Σ is a suitable matrix for which the canonical choice is an approximation of the scaled inverse Fisher Information matrix of f_{θ} .

Canonical Choice of Σ The family of Delta Variances in principle supports any positive definite matrix Σ . We will see in Section 5.1 that it intuitively represents the posterior covariance of the parameters θ after learning f_θ on the training data \mathcal{D} . The canonical choice is $\Sigma := \frac{1}{N} \hat{F}_f^{-1}$ being the inverse empirical Fisher Information matrix scaled by the number of data-points N . Plugged into the Delta Variance formula we obtain the following estimate for the Epistemic Variance of $u_\theta(z)$:

$$\nabla_{\theta \sim p(\theta|f,\mathcal{D})} [u_\theta(z)] \approx \frac{1}{N} \nabla_\theta u_\theta(z)^\top \hat{F}_f^{-1} \nabla_\theta u_\theta(z)$$

It is worth emphasizing that the Fisher information is computed using $f_\theta(x)$ (the model that was used for training θ) while the gradient delta vectors come from $u_\theta(z)$ the quantity of interest that is evaluated. Hence \hat{F}_f^{-1} can be precomputed and reused for various choices of u_θ .

Intuition Section 5 explores multiple ways to theoretically justify the Delta Variance family. The Bayesian intuition is that Σ captures the posterior covariance of the parameters θ while $\Delta_{u(z)} = \nabla_\theta u_\theta(z)$ translates this parameter uncertainty from variations in θ to variations in $u_\theta(z)$. In Figure 2 we consider an illustrative example, where a survival rate of $f_\theta(x) = \theta$ has been estimated and is used to make predictions 10 years ahead via $u_\theta(z) = \theta^{10}$.

Theoretical Motivation The family of Delta Variance estimators is motivated because under strong conditions (see Section 5.1) and for number of data-points N it can be shown to recover the Epistemic Variance up to a diminishing error:

$$\underbrace{\nabla_{\theta \sim p(\theta|f,\mathcal{D})} [u_\theta(z)]}_{\text{Epistemic Variance}} = \underbrace{\Delta_{u(z)}^\top \Sigma \Delta_{u(z)}}_{\text{Delta Variance}} + O(N^{-1.5})$$

An additional motivation is that it can be derived using mild assumptions from a leave-one-out or an adversarial robustness perspective (see Sections 5.2 and 5.3).

Computational Convenience Delta Variances are convenient because $\Delta_{u(z)} := \nabla_\theta u_\theta(z)$ can be computed using any auto-differentiation framework and because Σ does not depend on u_θ (e.g. can be re-used for many different quantities of interest u_θ). It is efficient because it is a vector-matrix-vector product, where the matrix can be approximated efficiently (e.g. diagonally, low-rank, or using KFAC (Martens 2014)).

4.1 How to choose Σ

Principled choices of Σ Theory suggests three principled choices for the covariance matrix, which all scale as $\Sigma \propto \frac{1}{N}$. Each choice can be derived in at least two ways using statistics or using influence functions (see Section 5 for details and Table 2 for an overview).

1. The inverse Fisher Information F_f^{-1} divided by N .
2. The inverse Hessian of the training loss H_f^{-1} .
3. The sandwich $H_f^{-1} F_f H_f^{-1}$ times N .

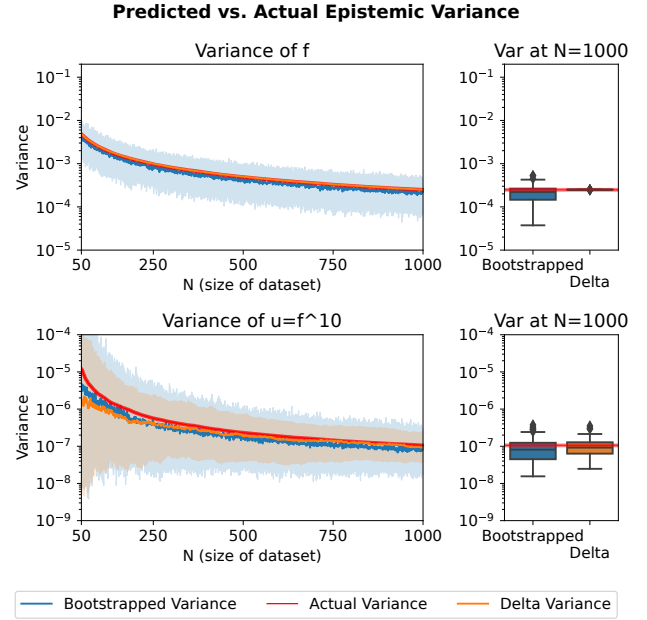


Figure 2: Illustrative survival prediction example. Actual epistemic variance (red) vs. predicted variance using the Delta Variance (orange) or a 10-fold Bootstrap (blue) as the dataset size N grows. Shaded confidence areas contain 95% of the variance predictions. Bold lines are the median. Observe that the orange median line of the Delta Variance and the actual variance in red overlap largely. Top: variance of learned function $f_\theta(x) = \theta$ Bottom: variance of quantity of interest $u_\theta(x) := \theta^{10}$ evaluations. All methods yield reasonable results for $N > 10$ with ensemble methods exhibiting higher variance. Generally the variance for u_θ is harder to estimate than for f_θ .

For well-specified models the three covariance matrices become eventually equivalent as Hessian divided by N and empirical Fisher converge to the true Fisher as data increases. In practice they need to be efficiently approximated from finite data (e.g. diagonally or using KFAC) and safely inverted. The first and third Bayesian approach use F_f which can be approximated by the empirical Fisher information \hat{F}_f and is easily invertible as it is non-negative by construction. Frequentist analogues use the empirical \hat{F}_f directly. In contrast H_f is only non-negative at a maximum which may not be reached precisely with stochastic optimization. Hence inverting H_f requires more careful regularization (Martens 2014). For simplicity we select Σ to be a diagonal approximation of the empirical Fisher in our experiments, which alleviates the question of regularization.

Fine-tuning or Learning Σ The analytic form of the Delta Variance permits to back-propagate into the values of Σ . This enables approaches that learn better values for Σ from scratch or improve the values via fine-tuning. We explore a simple example in Section 7.1 that improves empirically over the regular Fisher information by re-scaling some of its entries.

Bayesian Interpretations (Section 5.1)	Frequentist Interpretations	Choice of Σ (modulo factors of N)
Bernstein-von Mises Posterior ⁴	OOD Detection ¹ (Sec. 5.4)	F^{-1}
Misspecified Bernstein-von Mises Posterior ⁴	Leave-one-out Variance ² (Sec. 5.2)	$H^{-1}FH^{-1}$
Laplace Posterior ³	Adversarial Robustness ² (Sec. 5.3)	H^{-1}

Table 2: For each choice of Σ there exists a Bayesian and a frequentist interpretation. Each interpretation requires different assumptions on $f_\theta(x)$ and its loss. Due to their milder assumptions the frequentist interpretations can serve as fall-back interpretations if the stricter conditions on the Bayesian interpretations are not met. For example observe that assuming a Bernstein-von Mises Posterior is computationally equivalent to performing OOD Detection. Interestingly the former makes strong assumptions about f_θ which typically do not apply to neural networks, while the later only requires differentiability of f_θ and u_θ and that θ converges locally. The Hessian is computed with respect to the training loss of f_θ . The Bayesian interpretations start with the true Fisher matrix and approximate it from data e.g. with \hat{F} . The frequentist approximations work with \hat{F} directly. In practice both H and \hat{F} are computed at the locally optimal parameters $\hat{\theta}$. We consider Σ modulo factors of N as they do not change the interpretation.

5 Analysis

In this section we will investigate multiple ways to derive and motivate Delta Variances. Broadly speaking they can be separated into three classes:

1. In Section 5.1 we begin with the easiest derivations, which approximate the Bayesian posterior and make strong assumptions that may not always apply to neural networks.
2. In Section 5.2 we consider the frequentist analogue of Epistemic Variance, that is compatible with neural networks and does not make assumptions about any posterior.
3. In Section 5.3 and 5.4 we consider alternative derivations that are based on adversarial robustness and out-of-distribution detection and rely on even fewer assumptions.

All of the considered derivations yield Delta Variances with principled covariance matrices. For an overview consider Table 5.1, where we can observe that assuming a Bernstein-von Mises Posterior is computationally equivalent to performing OOD Detection. Interestingly the former makes strong assumptions about f_θ which typically do not apply to neural networks, while the later only requires that the covariance of gradients is finite and that θ converges locally. Due to their milder assumptions the frequentist interpretations can serve as fall-back interpretations if the stricter conditions on the Bayesian interpretations are not met.

5.1 Bayesian Interpretation

We begin with a derivation that gives rise to a bound on the approximation error. While requiring strong assumptions, it serves as a motivation and introduction. The error diminishes with the number of observed data-points N :

$$\underbrace{\mathbb{V}_{\theta \sim p(\theta|f, \mathcal{D})}[u_\theta(z)]}_{\text{Epistemic Variance}} = \underbrace{\Delta_{u(z)}^\top \Sigma \Delta_{u(z)}}_{\text{Delta Variance}} + O(N^{-1.5})$$

Bernstein-von Mises Motivation As a motivational introduction we will derive the approximation error when the Bernstein-von Mises conditions are met (e.g. differentiability and unique optimum – see van der Vaart (1998) for details). Under such conditions the posterior converges to a Gaussian

distribution centered around the maximum likelihood solution θ with a scaled inverse Fisher Information as covariance matrix.

$$P(\theta|\mathcal{D}) \rightarrow \mathcal{N}\left(\theta, \frac{1}{N}F_f^{-1}\right)$$

The Epistemic Variance can then be computed using the Delta Method resulting in Proposition 1.

Proposition 1. *For a normally distributed posterior with mean $\hat{\theta}$ and a covariance matrix Σ proportional to $\frac{1}{N}$ it holds:*

$$\underbrace{\mathbb{V}_{\theta \sim p(\theta|f, \mathcal{D})}[u_\theta(z)]}_{\text{Epistemic Variance}} = \underbrace{\Delta_{u(z)}^\top \Sigma \Delta_{u(z)}}_{\text{Delta Variance}} + O(N^{-1.5})$$

where $\Delta_{u(z)} := \nabla_\theta u_\theta(z)|_{\theta=\hat{\theta}}$ as usual.

Proof. See Schmitt, Shawe-Taylor, and van Hasselt (2025). \square

If the Bernstein-von Mises conditions are met Proposition 1 holds with $\Sigma = F_f^{-1}/N$.

Further Bayesian Interpretations Other Gaussian posterior approximations can be considered by plugging their respective posterior covariance matrix into Proposition 1: The misspecified Bernstein-von Mises theorem (see Kleijn and van der Vaart 2012) states that we obtain the sandwich covariance $H_f^{-1}F_fH_f^{-1} \times N$ if the model f_θ is misspecified (i.e. does not represent the data well). Proponents advocate that the sandwich estimate is more robust to heteroscedastic noise while others argue against it (Freedman 2006). Similarly a Laplace approximation (Laplace 1774; MacKay 1992b; Ritter, Botev, and Barber 2018) can be made resulting in a Delta Variance with H_f^{-1} . Again those choices of Σ are $\propto \frac{1}{N}$.

¹Only assumes differentiable u_θ and locally converged $\hat{\theta}$.

²Assumes: Bounded second derivatives of $\log f_\theta$ and u_θ , Hessian at $\hat{\theta}$ is invertible, and all from (1).

³Assumes: All above, well-specified model f_θ , $\hat{\theta}$ converged to a unique optimum of the posterior with locally quadratic shape. Often referred to as *Laplace Approximation* when some conditions are not met. See MacKay (1992b) for more details.

⁴Assumes: All above except the local shape requirement on the posterior in (3). Uniformly consistent maximum likelihood estimator and see van der Vaart (1998) for more conditions.

5.2 Frequentist Interpretation

To better cater to complex function approximators such as neural networks this section discusses a frequentist derivation of the Delta Variance, which relies on milder assumptions: As it is frequentist it does not consider posterior distributions. This allows us to side-step any questions about the shape and tractability of posterior distributions for neural networks. It does not require global convexity or a unique optimum. Convergence of the parameters to some local optimum together with locally bounded second derivatives is sufficient. In Proposition 2 we observe that the Delta Variance computes an infinitesimal approximation to the leave-one-out variance (see Definition 5) for choice of $\Sigma = H_f^{-1} \hat{F}_f H_f^{-1}$:

$$\underbrace{\Delta_{u(z)}^\top \Sigma \Delta_{u(z)}}_{\text{Delta Variance}} = \underbrace{\mathbb{V}_{\theta \sim IJ} [u_\theta(z)]}_{\text{Infinitesimal LOO Variance}} \approx \underbrace{\mathbb{V}_{\theta \sim LOO} [u_\theta(z)]}_{\text{LOO Variance}}$$

The infinitesimal approximation to the leave-one-out variance (also known as the infinitesimal jackknife (Jaekel 1972)) is defined as follows:

Definition 4. Let θ_i be the parameters resulting from training f_θ on data \mathcal{D} with x_i down-weighted by ϵ (i.e. from weight 1 to $1 - \epsilon$), then the ϵ -Leave-One-Out Variance is defined as

$$\mathbb{V}_{\theta \sim \epsilon\text{-LOO}} [u_\theta(z)] := \frac{N - \epsilon}{\epsilon^2} \mathbb{V}_{i \sim U(1, \dots, N)} [u_{\theta_i}(z)]$$

Definition 5. With slight abuse of notation we define the Infinitesimal LOO Variance as the limit of the ϵ -Leave-One-Out Variance:

$$\mathbb{V}_{IJ} [u_\theta(z)] := \lim_{\epsilon \rightarrow 0} \mathbb{V}_{\theta \sim \epsilon\text{-LOO}} [u_\theta(z)]$$

Proposition 2. The Delta Variance equals the infinitesimal LOO Variance for $\Sigma = H_f^{-1} \hat{F}_f H_f^{-1} \times N$:

$$\mathbb{V}_{IJ} [u_\theta(z)] = \Delta_{u(z)}^\top \Sigma \Delta_{u(z)}$$

Proof. See Schmitt, Shawe-Taylor, and van Hasselt (2025). \square

5.3 Adversarial Data Interpretation

Sometimes it is of interest to quantify how much a prediction changes if the training dataset is subject to adversarial data injection. Intuitively this is connected to epistemic uncertainty: one may argue that predictions are more robust the more certain we are about their parameters and vice versa. In the appendix of Schmitt, Shawe-Taylor, and van Hasselt (2025) we show that this intuition also holds mathematically. In particular we observe that:

1. The Delta Variance with $\Sigma = H_f^{-1}$ computes how much a quantity of interest $u_\theta(z)$ changes if an adversarial data-point is injected.
2. This adversarial interpretation is technically equivalent to the Laplace Posterior approximation (from Section 5.1) – even though interestingly both start with different assumptions and objectives.

5.4 Out-of-Distribution Interpretation

We show that a large Delta Variance of $u_\theta(z)$ implies that its input z is out-of-distribution with respect to the training data. This relates to epistemic uncertainty intuitively: a model is likely to be uncertain about data-points that differ from its training data. The derivation in the appendix of Schmitt, Shawe-Taylor, and van Hasselt (2025) is based on the Mahalanobis Distance (Mahalanobis 1936) – a classic metric for out of distribution detection. It accounts for the possibility that $f_\theta \neq u_\theta$ and relies on minimal assumptions only requiring existence of gradients and that the training of f_θ has converged.

6 Experiments

To empirically study the Delta Variance we build on the state-of-the-art GraphCast weather forecasting system (Lam et al. 2023) which trains a neural network $f_\theta(x)$ to predict the weather 6 hours ahead. This $f_\theta(x)$ is then iterated multiple times to make predictions up to 10 days into the future. We define various quantities of interest u_θ such as the average rainfall in an area or the expected power of a wind turbine at a particular location and compute their Epistemic Variance. We assess the Epistemic Variance predictions on 5 years of hold-out data using multiple metrics such as the correlation between predicted variance and prediction error and the likelihood of the quantities of interest. Empirically Delta Variances with a diagonal Fisher approximation yield competitive results at lower computational cost – see Figure 3. Next we give an overview on the experimental methodology – please consider Schmitt, Shawe-Taylor, and van Hasselt (2025) for more technical details.

6.1 Weather Forecasting Benchmark

GraphCast Training We build on the state-of-the-art GraphCast weather prediction system. It trains a graph neural network to predict the global weather state 6 hours into the future. This step function $x_{t+1} = f_\theta(x_t)$ is then iterated to predict up to 10 days into the future. The global weather state x is represented as a grid with 5 surface variables and 6 atmospheric variables at 37 levels of altitude (see Lam et al. 2023, for details). The authors consider a grid-sizes of 0.25 degrees. To save resources we retrain the model for a grid size of 4 degrees and reduce the number of layers and latents each by factor a of 2. Finally we skip the fine-tuning curriculum for simplicity. Besides the graph neural network we also consider a standard convolutional neural network. Training data ranges from 1979-2013 with validation data from 2014-2017 and holdout data from 2018-2021 resulting in about 100 GB of weather data.

Quantities of Interest First we define 126 different quantities of interest u_θ based on 4 topics that we evaluate on the hold-out data (2018-2021) for two different neural network architectures: 1) Precipitation at various times into the future. 2) Inspired by wind turbine energy yield we measure the third power of wind-speed at various times into the future. 3) Inspired by flood risk we measure precipitation averaged over areas of increasing size five days into the future. 4) Inspired by health emergencies we predict the maximum temperature

Pareto Frontier: Improvement vs. Inference Compute

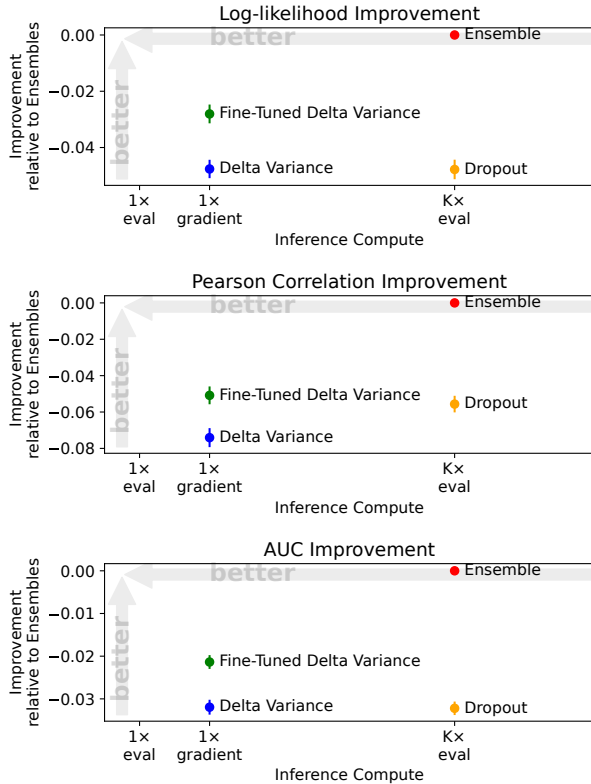


Figure 3: Comparison of variance estimators in terms of their inference cost and prediction quality. The quantities of interest are based on the GraphCast (Lam et al. 2023) weather prediction system that iterates a learned neural network dynamics model to form predictions. We evaluate the selected variance estimators based on three different evaluation criteria (Log-likelihood, correlation to prediction error and AUC akin to Van Amersfoort et al. (2020)). Lines indicate 2 standard errors. Delta Variances yield similar results as popular alternatives for lower computational cost. On average ensembles achieve the highest quality and Delta Variances the lowest computational overhead. See Section 7.1 for the fine-tuned Delta Variance.

maximized over areas of increasing size five days into the future. The first two quantities are predicted 1, . . . , 5 days ahead. The last two are measured 5 days ahead in quadratic areas with inradii ranging from 1 to 6. These measurements take place at 6 preselected capital cities. Finally note that we never train u_θ as it can be derived using f_θ .

Evaluation Methodology The data from 2018-2021 is held out for evaluation resulting in approximately 6×10^3 different (input, target-value) pairs for each of the 252 quantities of interest u_θ . For each pair we obtain a prediction error $|y - u_\theta(z)|$ and corresponding variance predictions $\nu(z)$. Unfortunately many practical applications do not admit ground truth values for Epistemic Variance that one could compare

variance estimators to. Instead there are multiple popular approaches in the literature relying on the prediction error, which is subject to both epistemic and aleatoric uncertainty. In Figure 3 we consider multiple such different criteria:

1. Akin to Van Amersfoort et al. (2020) AUC considers how fast the average L_1 error decreases when data-points are removed from the dataset – in the order of their largest predicted Epistemic Variance.
2. We consider the Pearson correlation between absolute error and predicted epistemic standard deviation: $corr(|y - u_\theta(z)|, \sqrt{\nu(z)})$.
3. To evaluate the Log-likelihood of observations y we interpret $u_\theta(z)$ as the mean of a Laplace distribution with variance derived from the predicted Epistemic Variance $\nu(z)$. We parameterize the Laplace distribution such that its variance decomposes in a constant α and the predicted Epistemic Variance $\nu(z)$ scaled by β . Intuitively α represents the aleatoric variance and $\beta\nu(z)$ represents the Epistemic Variance: $\text{Laplace}(\mu = u_\theta(z), 2b^2 = \alpha + \beta\nu(z))$. Both α and β are learned on the validation data (2014-2017) that is used for hyper-parameter selection. We then observe how well it models the actual observed target values y from the evaluation data (2018-2021).

Finally to reduce variance we define the *Improvement* of a variance estimator as the difference of its score to the score obtained by the ensemble estimator. Intuitively this indicates the loss in Quality when using an estimator in place of an ensemble. This procedure is repeated for each of the 252 quantities of interest u_θ .

7 Illustrations and Extensions

To highlight the generality of our approach we illustrate two extensions in this section.

1. By learning Σ to represent uncertainty well, we generalize the parametric form of Delta Variances beyond Fisher and Hessian matrices and observe improved results in the GraphCast benchmark – see Figure 3.
2. We consider an example where u_θ is not an explicit function but maps to a fixed-point of an iterative algorithm. We observe that it is possible to compute the Delta Variance of fixed-points using the implicit function theorem. Applied to an eigenvalue solver we observe empirically that the Delta Variance variance yields reasonable uncertainty estimates – see Figure 4.

7.1 Learning Σ

In Section 5 we observed that Delta Variances with special Σ such as the Fisher Information approximate theoretically established measures of uncertainty. In this section we observe that Σ may also be learned or fine-tuned. In an illustrative example we differentiate the Delta Variances with respect to Σ and use gradient descent to obtain an improved Σ . This may be helpful to improve the uncertainty prediction or to improve a downstream use-case if the variance is used in a larger system.

Fine-Tuning Σ Example We present a simple instance of fine-tuning a few parameters of Σ , which empirically yields improved results – see Figure 3. Note that Σ is approximated block-diagonally in most practical cases to limit the computational requirements – with one block for each weight vector in each neural network layer. Hence the Delta Variance splits into a sum of per-block Delta Variances derived from per-block gradients Δ_i :

$$\Delta_{f(x)}^\top \Sigma \Delta_{f(x)} = \sum_i \Delta_i^\top \Sigma_i \Delta_i$$

In this example we introduce a factor to rescale Σ within each block. Intuitively this adjusts the importance of each layer. Since only a few parameters need to be estimated we only need little fine-tuning data. This is applicable in situations where there is a small amount of training data for u_θ . In our experiments we optimize the coefficients of this linear combination using gradient descent to improve the log-likelihood or correlation on a small set of held-out validation data. Note that the per-layer variances can be cached which reduces the optimization problem significantly.

7.2 Epistemic Variance of Iterative Algorithms and Implicit Functions

So far we considered quantities of interest u_θ that are explicit functions of the parameters θ . Here we consider an example where the quantity of interest is an implicit function: u_θ maps to the fixed-point (solution) of an iterative algorithm for which there is no closed-form formula that we could differentiate to obtain its gradient.

Given some initial point w_0 the iteration $w_{k+1} = F_\theta(w_k)$ may converge to a fixed-point w_θ^* that depends on the parameters θ . To estimate $\nabla_\theta [w_\theta^*]$ we need to define u_θ as follows, which can not be differentiated with regular back-propagation due to the limit

$$u_\theta(w_0) := \lim_{k \rightarrow \infty} \underbrace{F_\theta \circ \dots \circ F_\theta}_{k \text{ times}}(w_0) = w_\theta^*$$

Implicit Epistemic Variance Calculation To compute the Delta Variance of an implicitly defined u_θ we need its gradient $\nabla_\theta u_\theta$. This can be obtained under mild conditions using the implicit function theorem. Let us denote $w_{k+1} = F_\theta(w_k)$ any fixed-point iteration converging to w^* with the corresponding non-linear equation $G_\theta(w) := F_\theta(w) - w = 0$. The implicit function theorem yields the gradient of u_θ by considering the Jacobian of G at the fixed-point w^* :

$$\Delta_{w^*} := \nabla_\theta u_\theta = -(\nabla_{w^*} G_\theta(w^*))^{-1} \nabla_\theta G_\theta(w^*)$$

whenever G is continuously differentiable and the inverse of $\nabla_{w^*} G_\theta(w^*)$ exists. Now we can compute the Epistemic Variance as $\nabla_\theta [w^*] \approx \Delta_{w^*}^\top \Sigma \Delta_{w^*}$.

Eigenvalue Example Eigenvalues are a quantity of interest in structural engineering. As an illustrative example we consider the eigenvalues $\lambda_i(A_\theta)$ of a finite element model matrix $A_\theta = M_\theta^{-1} K_\theta$ that indicates the stability of a physical structure. If the parameters θ are uncertain the eigenvalues will be uncertain as well. Recall that they are the solutions

to $\det(A_\theta - \lambda I) = 0$. We can estimate the Epistemic Variance of an eigenvalue $\nabla_\theta [\lambda_i(A_\theta)]$ using Delta Variances if we obtain the gradient of the eigenvalue $\nabla_\theta \lambda_i(A_\theta)$. To this extend we need the implicit function theorem as $\lambda_i(A_\theta)$ is an implicitly defined function – please consider Schmitt, Shawe-Taylor, and van Hasselt (2025) for technical details.

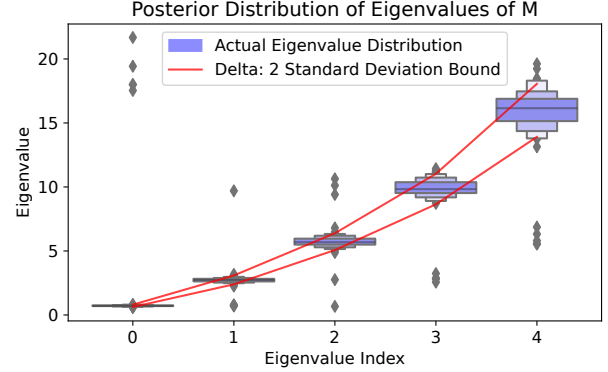


Figure 4: To investigate more intricate quantities of interest, we consider the mapping from a matrix A_θ to its eigenvalue $u_\theta = \lambda_i(A_\theta)$. This function is not explicit and computed using iterative algorithms, but we can use the implicit function approach to estimate the Delta Variance. Here A_θ is an illustrative finite-element problem with 11-dimensional parameters θ and 5 eigenvalues.

8 Related Work

The proposed Delta Variance family bridges and extends Bayesian, frequentist and heuristic notions of variance. Furthermore it generalizes related work by considering explicit and implicit quantities of interest other than the neural network itself $u_\theta \neq f_\theta$ and permits learning improved covariances Σ – see Section 7. Below we give a brief historic account of related methods that mostly consider the $u_\theta = f_\theta$ case.

Delta Method The Delta Method dates back to Cotes (1722), Lambert (1765) and Gauss (1823) in the context of error propagation and received a modern treatment by Kelley (1928); Wright (1934); Doob (1935); Dorfman (1938) – see Gorroochurn (2020) for a historical account. Denker and LeCun (1990) apply the Delta Method to the outputs of neural networks $f_\theta(x)$ and Nilsen et al. (2022) improves computational efficiency. When applied to neural networks the Delta Method requires strong assumptions about the posterior (e.g. unique optimum) or training process, which have not been proven to hold. Delta Variances – named after the Delta Method – provide multiple alternative theoretical justification through its unifying perspective. Furthermore Delta Variances generalize to the $u_\theta \neq f_\theta$ case and other Σ .

Laplace Approximation Building on work by Gull (1989), MacKay (1992b) and Ritter, Botev, and Barber (2018) apply the Laplace approximation to neural networks. Approximating functions at an optimum by what should later be called

a Gaussian distribution dates back to Laplace (1774). While only applicable to a single optimum MacKay (1992b) heuristically argues for its applicability to posterior distributions of neural networks. Given such Gaussian posterior approximation they apply the Delta Method yielding a special instance of the Delta Variance family with $u_\theta = f_\theta$ and $\Sigma = H_f^{-1}$ – see Section 5.1.

Influence Functions and Jackknife Methods Influence functions were proposed in Hampel (1974) concurrently with the closely related Infinitesimal Jackknife by Jaeckel (1972) which approximates cross validation (Quenouille 1949). Koh and Liang (2017) apply the influence function analysis to neural networks to evaluate how training data influences predictions. In Sections 5.2 and 5.3 we apply similar techniques to general quantities of interest different from f_θ .

Uncertainty Estimation for Deep Neural Networks We focus our comparison on two popular methods: Lakshminarayanan, Pritzel, and Blundell (2017) train multiple neural networks to form an ensemble and Gal and Ghahramani (2016) which re-interprets Monte-Carlo dropout as variational inference. In Table 1 we compare their properties with Delta Variances observing that they come at larger inference cost. Osband et al. (2023) aims to reduce the training costs of ensemble methods. To this extent they change the neural network architecture and training procedure, however how to reduce the remaining k -fold inference cost and memory requirements remain open research questions. Other popular methods come with similar requirements to change the architecture or training procedure (Blundell et al. 2015; Van Amersfoort et al. 2020; Immer, Korzepa, and Bauer 2021), while approaches like Sun et al. (2022) are of non-parametric flavour exhibiting inference cost that increases with the dataset size. SWAG (Maddox et al. 2019) reduces the training and memory cost by considering an ensemble of parameters from a single learning trajectory with stochastic gradient descent and approximating it with a Gaussian posterior. For inference they employ expensive k -fold sampling. We note that it is natural to derive a SWAG-inspired Delta Variances that employs the Σ from SWAG inside the computationally efficient Delta Variance formula – we leave those considerations for future research. Finally Kallus and McInerney (2022) propose a Delta Method inspired approach to approximate Epistemic Variance with an ensemble of two predictors and Schnaus et al. (2023) learn scale parameters in Gaussian prior distributions for transfer learning.

9 Conclusion

We have addressed the question of how the uncertainty from limited training data affects the computation of downstream predictions in a system that relies on learned components. To this extent we proposed the Delta Variance family, which unifies and extends multiple related approaches theoretically and practically. We discussed how the Delta Variance family can be derived from six different perspectives (including Bayesian, frequentist, adversarial robustness and out-of-distribution detection perspectives) highlighting its wide theoretical support and providing a unifying view on those

perspectives. Next we presented extensions and applications of the Delta Variance family such as its compatibility with implicit functions and ability to be improved through fine-tuning. Finally an empirical validation on a state-of-the-art weather forecasting system shows that Delta Variances yield competitive results more efficiently than other popular approaches.

Acknowledgements

We would like to thank Remi Lam for his advice and support with GraphCast-related questions. Furthermore we would like to thank Mark Rowland, Wojciech M. Czarnecki, Danilo J Rezende, Iurii Kemaev and the anonymous AAAI 2025 reviewers for their valuable feedback.

References

- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3): 235–256.
- Bernstein, S. N. 1917. *The Theory of Probabilities*.
- Blundell, C.; Cornebise, J.; Kavukcuoglu, K.; and Wierstra, D. 2015. Weight Uncertainty in Neural Network. In Bach, F.; and Blei, D., eds., *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, 1613–1622. Lille, France: PMLR.
- Cook, R. D.; and Weisberg, S. 1982. *Residuals and Influence in Regression*. Monographs on Statistics & Applied Probability. Chapman & Hall. ISBN 9780412242809.
- Cotes, R. 1722. *Harmonia Mensurarum*. Robert Smith.
- de Veaux, R. D.; Schumi, J.; Schweinsberg, J.; and Ungar, L. H. 1998. Prediction Intervals for Neural Networks via Nonlinear Regression. *Technometrics*, 40(4): 273–282.
- Denker, J.; and LeCun, Y. 1990. Transforming Neural-Net Output Levels to Probability Distributions. In Lippmann, R.; Moody, J.; and Touretzky, D., eds., *Advances in Neural Information Processing Systems*, volume 3. Morgan-Kaufmann.
- Doob, J. L. 1935. The limiting distributions of certain statistics. *Ann. Math. Stat.*, 6(3): 160–169.
- Dorfman, R. 1938. A note on the delta-method for finding variance formulae. *Biometric Bulletin*.
- Duff, M. 2002. *Optimal Learning: Computational procedures for Bayes-adaptive Markov decision processes*. Ph.D. thesis, University of Massachusetts Amherst.
- Freedman, D. A. 2006. On The So-Called “Huber Sandwich Estimator” and “Robust Standard Errors”. *The American Statistician*, 60(4): 299–302.
- Gal, Y.; and Ghahramani, Z. 2016. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In Balcan, M. F.; and Weinberger, K. Q., eds., *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, 1050–1059. New York, New York, USA: PMLR.
- Gauss, C. 1823. *Theoria combinationis observationum erroribus minimis obnoxiae*. H. Dieterich.

- Gorroochurn, P. 2020. Who invented the delta method, really? *Math. Intelligencer*, 42(3): 46–49.
- Gull, S. F. 1989. *Developments in Maximum Entropy Data Analysis*, 53–71. Dordrecht: Springer Netherlands. ISBN 978-94-015-7860-8.
- Hampel, F. R. 1974. The Influence Curve and Its Role in Robust Estimation. *Journal of the American Statistical Association*, 69(346): 383–393.
- Heger, M. 1994. Consideration of Risk in Reinforcement Learning. In *Machine Learning: Proceedings of the 11th International Conference*, 105–111. Morgan Kaufmann Publishers, San Francisco, CA.
- Hodges, J. L. 1967. Efficiency in normal samples and tolerance of extreme values for some estimates of location. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, 163–186. Berkeley: University of California Press.
- Hwang, J. T. G.; and Ding, A. A. 1997. Prediction Intervals for Artificial Neural Networks. *Journal of the American Statistical Association*, 92(438): 748–757.
- Immer, A.; Korzepa, M.; and Bauer, M. 2021. Improving predictions of Bayesian neural nets via local linearization. In Banerjee, A.; and Fukumizu, K., eds., *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, 703–711. PMLR.
- Jaekel, L. 1972. The Infinitesimal Jackknife. *Bell Lab. Memorandum*, MM72-1215-11.
- Kallus, N.; and McInerney, J. 2022. The Implicit Delta Method. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 37471–37483. Curran Associates, Inc.
- Kelley, T. L. 1928. *Crossroads in the mind of man; a study of differentiable mental abilities*. Palo Alto: Stanford Univ. Press.
- Kleijn, B.; and van der Vaart, A. 2012. The Bernstein-Von-Mises theorem under misspecification. *Electronic Journal of Statistics*, 6(none): 354 – 381.
- Koh, P. W.; and Liang, P. 2017. Understanding Black-box Predictions via Influence Functions. In Precup, D.; and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 1885–1894. PMLR.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Lam, R.; Sanchez-Gonzalez, A.; Willson, M.; Wirnsberger, P.; Fortunato, M.; Alet, F.; Ravuri, S.; Ewalds, T.; Eaton-Rosen, Z.; Hu, W.; Merose, A.; Hoyer, S.; Holland, G.; Vinyals, O.; Stott, J.; Pritzel, A.; Mohamed, S.; and Battaglia, P. 2023. Learning skillful medium-range global weather forecasting. *Science*, 382(6677): 1416–1421.
- Lambert, J. H. 1765. *Beyträge zum Gebrauche der Mathematik und deren Anwendung*, volume 1, chapter 13 of *Beyträge zum Gebrauche der Mathematik und deren Anwendung*. Verlag des Buchladens der Realschule.
- Laplace, P. S. 1774. Mémoire sur la probabilité des causes par les événements. *Mémoires de Mathématique et de Physique*, 6.
- Le Cam, L. 1953. *On some asymptotic properties of maximum likelihood estimates and related Baye’s estimates*. University of California Press, Berkeley.
- MacKay, D. J. C. 1992a. Information-based objective functions for active data selection. *Neural Computation*, 4(2): 550–604.
- MacKay, D. J. C. 1992b. A Practical Bayesian Framework for Backpropagation Networks. *Neural Computation*, 4: 448–472.
- Maddox, W. J.; Izmailov, P.; Garipov, T.; Vetrov, D. P.; and Wilson, A. G. 2019. A Simple Baseline for Bayesian Uncertainty in Deep Learning. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Mahalanobis, P. C. 1936. On The Generalized Distance in Statistics. *Sankhyā: The Indian Journal of Statistics, Series A (2008-)*, 80: pp. S1–S7.
- Martens, J. 2014. New perspectives on the natural gradient method. *CoRR*, abs/1412.1193.
- Martens, J.; and Grosse, R. B. 2015. Optimizing Neural Networks with Kronecker-factored Approximate Curvature. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, 2408–2417.
- Miller, R. G. 1974. The Jackknife—A Review. *Biometrika*, 61(1): 1–15.
- Nilsen, G. K.; Munthe-Kaas, A. Z.; Skaug, H. J.; and Brun, M. 2022. Epistemic uncertainty quantification in deep learning classification by the Delta method. *Neural Networks*, 145: 164–176.
- Osband, I.; Wen, Z.; Asghari, S. M.; Dwaracherla, V.; IBRAHIMI, M.; Lu, X.; and Van Roy, B. 2023. Epistemic Neural Networks. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 2795–2823. Curran Associates, Inc.
- Quenouille, M. H. 1949. Approximate Tests of Correlation in Time-Series. *Journal of the Royal Statistical Society. Series B (Methodological)*, 11(1): 68–84.
- Ritter, H.; Botev, A.; and Barber, D. 2018. A Scalable Laplace Approximation for Neural Networks. In *International Conference on Learning Representations*.
- Schmitt, S.; Shawe-Taylor, J.; and van Hasselt, H. 2025. General Uncertainty Estimation with Delta Variances. *CoRR*.
- Schnaus, D.; Lee, J.; Cremers, D.; and Triebel, R. 2023. Learning Expressive Priors for Generalization and Uncertainty Estimation in Neural Networks. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on*

Machine Learning, volume 202 of *Proceedings of Machine Learning Research*, 30252–30284. PMLR.

Sun, Y.; Ming, Y.; Zhu, X.; and Li, Y. 2022. Out-of-Distribution Detection with Deep Nearest Neighbors. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 20827–20840. PMLR.

Tibshirani, R. 1996. A Comparison of Some Error Estimates for Neural Network Models. *Neural Computation*, 8(1): 152–163.

Tishby, N.; Levin, E.; and Solla, S. 1989. Consistent inference of probabilities in layered networks: Predictions and generalization. In Anon, ed., *IJCNN Int Jt Conf Neural Network*, 403–409. Publ by IEEE. IJCNN International Joint Conference on Neural Networks ; Conference date: 18-06-1989 Through 22-06-1989.

Tukey, J. W. 1958. Bias and confidence in not-quite large samples (abstract). *J-ANN-MATH-STAT*, 29(2): 614–614.

Van Amersfoort, J.; Smith, L.; Teh, Y. W.; and Gal, Y. 2020. Uncertainty estimation using a single deep deterministic neural network. In *Proceedings of the 37th International Conference on Machine Learning*, ICML'20. JMLR.org.

van der Vaart, A. W. 1998. *Asymptotic Statistics*. Cambridge University Press.

von Mises, R. 1931. *Wahrscheinlichkeitsrechnung und ihre Anwendung in der Statistik und theoretischen Physik*, volume 1. Franz Deuticke.

Wright, S. 1934. The method of path coefficients. *Ann. Math. Stat.*, 5(3): 161–215.