

One Step Closer to Unbiased Aleatoric Uncertainty Estimation

Wang Zhang¹, Ziwen Martin Ma², Subhro Das³,
Tsui-Wei Lily Weng⁴, Alexandre Megretski¹, Luca Daniel¹, Lam M. Nguyen⁵

¹ Massachusetts Institute of Technology, Cambridge, MA, USA

² Harvard University, Cambridge, MA, USA

³ MIT-IBM Watson AI Lab, IBM Research, Cambridge, MA, USA

⁴ University of California San Diego, San Diego, CA, USA

⁵ IBM Research, Thomas J. Watson Research Center, Yorktown Heights, NY, USA

Abstract

Neural networks are powerful tools in various applications, and quantifying their uncertainty is crucial for reliable decision-making. In the deep learning field, the uncertainties are usually categorized into aleatoric (data) and epistemic (model) uncertainty. In this paper, we point out that the existing popular variance attenuation method highly overestimates aleatoric uncertainty. To address this issue, we propose a new estimation method by actively de-noising the observed data. By conducting a broad range of experiments, we demonstrate that our proposed approach provides a much closer approximation to the actual data uncertainty than the standard method.

1 Introduction

Neural networks (NN) are capable of performing various regression and classification tasks in computer vision and natural language processing, ranging from basic multi-layer perceptron (MLP) that approximate straightforward functions to more sophisticated structures like transformers. As the capabilities of neural networks continue to expand, it becomes increasingly important to acknowledge the limitations of machine learning techniques. For instance, how does the corrupted data affect the trained model? And when does the model lack confidence in its predictions? Since most the neural network approaches have failed to capture the prediction confidence, quantifying the uncertainties for the deep learning methods becomes a crucial research topic.

In the machine learning literature, the sources of uncertainty can be classified into two categories (Der Kiureghian and Ditlevsen 2009; Hüllermeier and Waegeman 2021; Gal et al. 2016): (1) aleatoric uncertainty, or data uncertainty, and (2) epistemic uncertainty, or model uncertainty. Aleatoric uncertainty is a measure of the inherent complexity or randomness that arises from the process of observing data from a true system, such as the noise in the measurements obtained from devices. This type of uncertainty is considered “*irreducible*” because no amount of additional data can eliminate the inherent stochasticity of the observation process. In contrast, epistemic uncertainty refers to the uncertainty that arises from the parameters of the model used

to analyze the data, and can be reduced by increasing the size of the training data, which in turn increases the confidence in the model. Distinguishing between aleatoric and epistemic uncertainties can be beneficial in a wide range of applications. For example, aleatoric uncertainty can be used to recover the measurement tolerance of an instrument, while epistemic uncertainty can guide optimization algorithms during training (Choi et al. 2021) or help to identify outliers during prediction (Seeböck et al. 2019).

There have been recent developments in deep learning research focused on techniques to capture aleatoric and epistemic uncertainties. To estimate the model uncertainty, one popular approach is to use Bayesian neural networks (Hernández-Lobato and Adams 2015; Blundell et al. 2015), which are neural networks that incorporate prior probability distributions over the model parameters. Another approach is to use ensemble methods (Parker 2013; Lakshminarayanan, Pritzel, and Blundell 2017), where multiple models are trained on the same dataset and their predictions are combined to produce a final prediction that incorporates epistemic uncertainties. Gal and Ghahramani (2016) showed the Monte Carlo dropout in deep learning approximates the Bayesian inference. A more recent branch directly estimates the aleatoric uncertainty together with the epistemic uncertainty (Lakshminarayanan, Pritzel, and Blundell 2017; Amini et al. 2020; Valdenegro-Toro and Mori 2022), by adding a parameterized variance term to the original stochastic prediction model. This addition aims to capture aleatoric uncertainty with a commonly used variance attenuation loss¹. Although it may appear to be capable to handle two sources of uncertainty, this method has several significant drawbacks. Seitzer et al. (2022); Stirn et al. (2023) observed a decrease in the prediction performance when training it together with the variance head, which is attributed to the loss function being heavily biased on well-predicted samples with small variances. The prediction model performance difference presents additional challenges to the fair evaluation of the variance approximation module. Additionally, as will be further discussed in this paper², the anticipated variance term under the ideal loss function exceeds solely the aleatoric uncertainty. Indeed, in the prior works

¹see Equation (2)

²see Section 2.2

on deep uncertainty estimation (Lakshminarayanan, Pritzel, and Blundell 2017; Amini et al. 2020; Valdenegro-Toro and Mori 2022; Seitzer et al. 2022; Stirn et al. 2023), the evaluation criteria are usually regression error or classification accuracy of the prediction model, rather than the quantitative difference between approximated uncertainty and true uncertainty. Therefore, gauging the effectiveness of uncertainty disentanglement becomes difficult.

The goal of this paper is to accurately identify the aleatoric uncertainty. We propose to split the prediction model training from the existing algorithm and develop a prediction-model-agnostic denoising approach that can better approximate the true noise level with a given trained prediction model. By augmenting a variance approximation module, this approach relies on the assumption that the distribution of the noise has zero mean, and seeks to remove the noise from the data in order to better capture the underlying signal. Doing so enables us to gain a more accurate understanding of the sources of uncertainty in the data, and leverage this information to improve our model and analysis.

We summarize our main contribution as follows:

1. Through theoretical analysis and experimental validation, we point out that a popular aleatoric uncertainty estimator does not identify the true data uncertainty. In fact, its expectation consists of aleatoric uncertainty and the squared expectation of the difference between the true labels and prediction average.
2. Assuming the zero mean distribution of data noise, we propose a denoising method called Denoising Variance Attenuation (DVA)³, to actively infer the true data noise. We introduce additional parameterization on noise variance and normalized true noise, and optimize them through a customized projected gradient descent method. The method pipeline is illustrated in Figure 1 with a detailed explanation in Section 3.
3. From a theoretical standpoint, while our proposed method is still a biased estimator, it provides a more accurate approximation of the aleatoric uncertainty than the state-of-the-art alternatives. In practical applications, our approach yields estimates that are useful for capturing the true uncertainty of the data, and can be valuable to a wide range of problems.

2 Problem Formulation

We start with a regression problem where the system input is $x \in \mathcal{X} \subset \mathbb{R}^n$ and output is $y \in \mathbb{R}^m$, where \mathcal{X} is a compact set with positive volume. There exists a continuous mapping function $g(\cdot) : \mathcal{X} \rightarrow \mathbb{R}^m$, which serves as the target function to regress. Consider a distribution over input feature $x: p(x)$, the finite dataset $D = \{(x_i, y_i)\}_{i=1}^M$ denotes M sampled inputs from $p(x)$ and the corresponding outputs. When gathering data from the real world, it is common for the data to contain noise, which signifies the presence of aleatoric uncertainty. We denote the observed dataset by $\tilde{D} = \{(\tilde{x}_i, \tilde{y}_i)\}_{i=1}^M$, where $\tilde{x}_i = x_i + \xi_i$, $\tilde{y}_i = y_i + \epsilon_i$ are

³Source code available at <https://github.com/wz16/DVA>. Please refer to arXiv for full technical appendix.

the noisy data and $\{\xi_i, \epsilon_i\}$ are the i.i.d. realizations of the random noise variables ξ, ϵ . In this paper, we assume the random noise distribution has zero-mean distribution and finite variance, i.e. $\mathbb{E}[\xi] = 0, \mathbb{E}[\xi^2] = \sigma_\xi^2 < \infty, \mathbb{E}[\epsilon] = 0, \mathbb{E}[\epsilon^2] = \sigma_\epsilon^2 < \infty$. When the noise variance is uniform across the system space, the aleatoric uncertainty is called “*homoscedastic*”; and, if the uncertainty varies against the system space, e.g. $\sigma_\xi^2(x), \sigma_\epsilon^2(x)$, the aleatoric uncertainty is considered “*heteroscedastic*”. In the classification problem, the output has the form of categorical labels, the aleatoric uncertainty arises from mislabeling or ambiguity between classes (e.g. the label is a probability distribution rather than a one-hot encoding).

Epistemic uncertainty evaluates the level of confidence in a model’s prediction, which may be indirectly influenced by factors such as insufficient data, algorithmic randomness, or model structure choices. To capture the epistemic uncertainty, the model prediction is usually presented by a distribution rather than a deterministic value, e.g. $f_\theta(x) \sim N(\mu_\theta(x), s_\theta^2(x))$, indicating that the model output follows a Gaussian distribution with mean and variance parameterized by θ , and the variance identifies the epistemic uncertainty.

Epistemic uncertainty relies on many factors (e.g., model choice, training algorithms, etc.), whereas aleatoric uncertainty is primarily an inherent property of the data. To assess the effectiveness of separating the two uncertainties, it is natural to compare the approximated data uncertainty with the true noise level. However, previous studies that claim to disentangle the two uncertainties (Lakshminarayanan, Pritzel, and Blundell 2017; Amini et al. 2020; Seitzer et al. 2022; Valdenegro-Toro and Mori 2022; Stirn et al. 2023) often rely on the output regression error or classification accuracy as the benchmark criterion, rather than the direct comparison between the estimated aleatoric uncertainty and the true noise level. Our study aims to identify data noise (i.e., $\sigma_\xi^2, \sigma_\epsilon^2$) by quantitatively comparing the estimated noise levels with the true ones. As a byproduct, we develop a denoising strategy for noisy data, leading to improved prediction performance, evidenced by reduced regression error.

2.1 Prior Works on Disentangling Uncertainties

Mean squared error (MSE) is commonly used as a loss function to train regression tasks, i.e. $\frac{1}{M} \sum_{i=1}^M (\tilde{y}_i - f_\theta(x_i))^2$, where $f_\theta(\cdot)$ is a continuous regression function parameterized by θ and can be either deterministic or stochastic. From a probabilistic viewpoint, minimizing MSE is equivalent to minimizing the negative log-likelihood of the observed data sampled from a Gaussian distribution centered at the prediction value, assuming a uniform noise variance $\hat{\sigma}^2$,

$$\begin{aligned} & -\log(\prod_{i=1}^M p(\tilde{y}_i | f_\theta(x_i))) \\ &= \frac{1}{M} \sum_{i=1}^M \frac{(\tilde{y}_i - f_\theta(x_i))^2}{2\hat{\sigma}^2} + \frac{\log(\hat{\sigma}^2)}{2} + \text{constant}. \end{aligned} \quad (1)$$

Note that Equation (1) differs from the MSE loss by a coefficient, the variance term $\hat{\sigma}^2$. In most optimization formulation, this term is usually neglected and therefore it fails

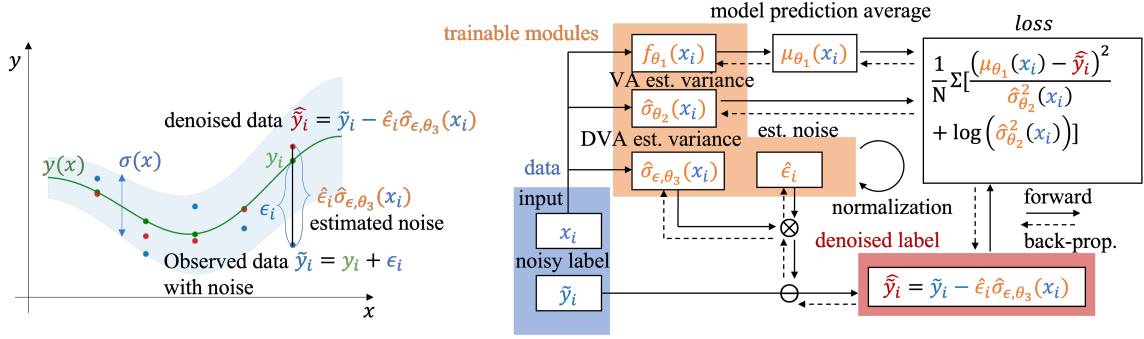


Figure 1: Illustrative diagram and pipeline of the denoising variance attenuation method

to capture the uncertainty. When $f_{\theta}(\cdot)$ is stochastic (e.g. Bayesian Neural network), we let $f_{\theta}(\cdot) = \mu_{\theta}(\cdot) + \epsilon_{\theta}(\cdot)$, where $\mu_{\theta}(\cdot)$ and $\epsilon_{\theta}(\cdot)$ denote the deterministic mean and epistemic random component with a variance of $s_{\theta}^2(\cdot)$. Inspired by Nix and Weigend (1994), a popular method (Lakshminarayanan, Pritzel, and Blundell 2017; Amini et al. 2020; Valdenegro-Toro and Mori 2022) replaces the fixed variance with a predicted one, to identify the aleatoric uncertainty. This modified loss function is called variance attenuation (VA):

$$L_{va} = \frac{1}{M} \sum_{i=1}^M \frac{(\tilde{y}_i - \mu_{\theta_1}(x_i))^2}{2\hat{\sigma}_{\theta_2}^2(x_i)} + \frac{\log(\hat{\sigma}_{\theta_2}^2(x_i))}{2}, \quad (2)$$

where $\hat{\sigma}_{\theta_2}(\cdot) > 0$ is a parameterized and continuous estimator for aleatoric uncertainty, and $\mu_{\theta_1}(\cdot)$ is the mean of stochastic prediction $f_{\theta_1}(\cdot)$. For notation clarity, we rewrite the parameter of the regression function and variance estimator as θ_1 and θ_2 , respectively.

Note that the above estimator implicitly assumes the input feature x has zero noise ($\sigma_{\xi}^2(x) \equiv 0$). Besides, the additional variance head in Equation (2) does not affect the original model $f_{\theta_1}(\cdot)$, the epistemic uncertainty is still incorporated in the distribution of $f_{\theta_1}(\cdot)$. There exist many quantification approaches for epistemic uncertainty. In Bayesian neural networks (Blundell et al. 2015), a posterior distribution of the neural network parameter θ is inferred during training, and given a testing input, the prediction distribution can be approximated by drawing samples of θ . For ensemble methods, multiple samplings are also required to approximate $f_{\theta_1}(\cdot)$. Amini et al. (2020) leverages Normal Inverse-Gamma (NIG) distribution and proposes Deep Evidential Regression (DER) to analytically output the mean and variance of the prediction distribution along with the aleatoric uncertainty. This approach streamlines the sampling process and enables efficient uncertainty learning.

Once $\mu_{\theta_1}(\cdot)$, $s_{\theta_1}^2(\cdot)$ are obtained from the prediction model, $\mu_{\theta_1}(\cdot)$ is incorporated into Equation (2) to help model to estimate $\hat{\sigma}_{\theta_2}^2(\cdot)$. This completes the framework of disentangling uncertainties, where $s_{\theta_1}^2(\cdot)$ and $\hat{\sigma}_{\theta_2}^2(\cdot)$ handles epistemic and aleatoric uncertainty respectively.

2.2 Limitations of the Aleatoric Estimator

In this section, we argue that the aleatoric estimator $\sigma_{\theta_2}(\cdot)$ in Equation (2) is imperfect and does not capture the exact aleatoric uncertainty. Consider a homoscedastic problem with uniform label noise variance $\mathbb{E}[(\tilde{y} - y)^2] = \sigma_{\epsilon}^2$, and $\hat{\sigma}_{\theta_2}^2$ being a single optimizable scalar parameter rather than a function. Fix the regression function parameter θ_1 and optimize Equation (2) over $\hat{\sigma}_{\theta_2}^2$, where the optimization parameter is a local minimum. Then, it reduces to,

$$\begin{aligned} \frac{dL_{va}}{d\hat{\sigma}_{\theta_2}^2} &= \sum_{i=1}^M \frac{-(\tilde{y}_i - \mu_{\theta_1}(x_i))^2}{2\hat{\sigma}_{\theta_2}^4} + \frac{1}{2\hat{\sigma}_{\theta_2}^2} = 0, \\ \implies \hat{\sigma}_{\theta_2}^2 &= \frac{1}{M} \sum_{i=1}^M (\tilde{y}_i - \mu_{\theta_1}(x_i))^2, \\ \implies \mathbb{E}[\hat{\sigma}_{\theta_2}^2] &= \mathbb{E}_{x_i, \epsilon_i} [(y_i + \epsilon_i - \mu_{\theta_1}(x_i))^2] \\ &= \mathbb{E}[(y - \mu_{\theta_1}(x))^2] + \sigma_{\epsilon}^2. \end{aligned} \quad (3)$$

where, σ_{ϵ}^2 is the variance of the data noise, $\mu_{\theta_1}(x)$ is the mean of the epistemic regression model. Equation (3) says that when the loss function reaches optimal, the estimator $\hat{\sigma}_{\theta_2}^2(\cdot)$ overestimates the aleatoric uncertainty by an expectation of square difference between the clean data and mean of regression predictions. The above derivation is extendable to a heteroscedastic setup by parameterizing $\hat{\sigma}_{\theta_2}^2(\cdot)$ to a function of x and analyzing the necessary condition for the extreme points. The conclusions are similar for the heteroscedastic cases and the analysis is deferred to the Appendix A.

3 Proposed Method

In the previous section, we presented a theoretical argument highlighting the substantial bias present in the aleatoric uncertainty estimates obtained via the popular VA method. Here, we introduce a novel method that is developed to mitigate this bias and to offer more precise estimates⁴.

⁴Our method incorporates an additional module to improve variance estimation, rather than retrain the regression model $f_{\theta_1}(\cdot)$ from scratch.

3.1 Normalized Gradient Descent Denoising

When dealing with a dataset of a sufficiently large size, an intriguing opportunity arises to take advantage of the property that the mean value of noise realizations tends toward zero. By leveraging this characteristic of large datasets, we can design algorithms that recover individual noise values for each data point and the corresponding noise variance, providing a more precise estimate of the noise levels present in the data.

Assuming that the examples have only label noise, for each data point (x_i, \tilde{y}_i) where $\tilde{y}_i = y_i + \epsilon_i$, we augment a normalized estimated noise $\hat{\epsilon}_i$. For the noise variance, we employ an estimating function $\hat{\sigma}_\epsilon^2(\cdot)$ which can be either heteroscedastic or homoscedastic. We aim to approximate the true noise ϵ_i by $\hat{\epsilon}_i \hat{\sigma}_\epsilon(x_i)$, or estimate the clean data y_i by denoising the observed data $\tilde{y}_i - \hat{\epsilon}_i \hat{\sigma}_\epsilon(x_i)$.

In this paper, we propose a normalized gradient descent type algorithm, named Denoising Variance Attenuation (DVA), which is inspired by the Projected Gradient Descent (PGD) Attack (Madry et al. 2017; Kurakin, Goodfellow, and Bengio 2018). The PGD Attack is a technique used to generate robust adversarial examples by iteratively updating strong adversaries using the loss function gradient and projecting them into the \mathcal{L}_p neighborhood of the initial input. Our proposed method applies a similar process to denoise inputs or outputs, by optimizing over the estimated noise and increasing the log-likelihood of the denoised data while ensuring that the total of normalized noises has zero mean and unit variance. The illustrative diagram and pipeline of the DVA method are shown in Figure 1. Let the estimated variance function be parameterized by $\hat{\sigma}_{\epsilon, \theta_3}(\cdot)$, and we formulate the constrained optimization problem on the denoising variance attenuation loss:

$$\begin{aligned} & \min_{\theta_2, \theta_3, \{\hat{\epsilon}_i\}_{i=1}^M} L_{dva}, \text{ where} \\ L_{dva} &= \frac{1}{M} \sum_{i=1}^M \frac{(\tilde{y}_i - \hat{\epsilon}_i \hat{\sigma}_{\epsilon, \theta_3}(x_i) - \mu_{\theta_1}(x_i))^2}{2\hat{\sigma}_{\theta_2}^2(x_i)} + \frac{\log(\hat{\sigma}_{\theta_2}^2(x_i))}{2}, \\ \text{s.t.} \quad & \sum_{i=1}^M \hat{\epsilon}_i / M = 0, \sum_{i=1}^M \hat{\epsilon}_i^2 / M = 1. \end{aligned} \quad (4)$$

We implement the constrained optimization in a straightforward way. After each gradient update, we perform a normalization procedure for all the inferred noise, to ensure that the estimated normalized noises have zero mean and unit variance. The key steps of the algorithm are listed below, with the detailed algorithm being deferred to Appendix B.

$$\begin{aligned} \text{Gradient step: } \theta_2 &\leftarrow \theta_2 - \eta \frac{dL_{dva}}{d\theta_2}; \theta_3 \leftarrow \theta_3 - \eta \frac{dL_{dva}}{d\theta_3}; \\ \hat{\epsilon}_i &\leftarrow \hat{\epsilon}_i - \eta \frac{dL_{dva}}{d\hat{\epsilon}_i}, \text{ for } i = 1, 2, \dots, M; \end{aligned}$$

$$\begin{aligned} \text{Normalization step: } \mu &\leftarrow \sum_{i=1}^M \hat{\epsilon}_i / M; \\ \text{var} &\leftarrow \sum_{i=1}^M (\hat{\epsilon}_i - \mu)^2 / M; \\ \hat{\epsilon}_i &\leftarrow \frac{\hat{\epsilon}_i - \mu}{\sqrt{\text{var}}}, \text{ for } i = 1, 2, \dots, M. \end{aligned} \quad (5)$$

Note that this optimization process is not dependent on the training of the regression model, $f_{\theta_1}(\cdot)$. In practice, they can be trained simultaneously or separately.

3.2 Necessary Conditions for Constrained Optimization

We start with a uniform noise variance and assign a single optimizable parameter for both $\hat{\sigma}_{\theta_2}^2$ and $\hat{\sigma}_{\epsilon, \theta_3}^2$. Consider the problem in Equation (4), when the optimization parameter is a local minimum, we have the following necessary conditions for constrained optimization with multiplier λ_1, λ_2 :

$$\begin{aligned} & \frac{-2(\tilde{y}_i - \hat{\epsilon}_i \hat{\sigma}_{\epsilon, \theta_3} - \mu_{\theta_1}(x_i)) \hat{\sigma}_{\epsilon, \theta_3}}{2\hat{\sigma}_{\theta_2}^2} + \lambda_1 + 2\lambda_2 \hat{\epsilon}_i = 0, \\ & \text{for } i = 1, 2, \dots, M; \\ & \sum_{i=1}^M \frac{-2(\tilde{y}_i - \hat{\epsilon}_i \hat{\sigma}_{\epsilon, \theta_3} - \mu_{\theta_1}(x_i)) \hat{\epsilon}_i}{2\hat{\sigma}_{\theta_2}^2} = 0, \\ & \sum_{i=1}^M \frac{(\tilde{y}_i - \hat{\epsilon}_i \hat{\sigma}_{\epsilon, \theta_3} - \mu_{\theta_1}(x_i))^2}{M} = \hat{\sigma}_{\theta_2}^2. \end{aligned} \quad (6)$$

While the derivation of this constrained optimization is deferred to Appendix A.2, its solution yields,

$$\begin{aligned} \hat{\sigma}_{\epsilon, \theta_3}^2 &= \frac{1}{M} \sum_{i=1}^M (\tilde{y}_i - \mu_{\theta_1}(x_i))^2 - \frac{1}{M^2} \left[\sum_{i=1}^M (\tilde{y}_i - \mu_{\theta_1}(x_i)) \right]^2 \\ \mathbb{E}[\hat{\sigma}_{\epsilon, \theta_3}^2] &= \frac{M-1}{M} \sigma_\epsilon^2 + \frac{M-1}{M} \mathbb{E}[(y - \mu_{\theta_1}(x))^2] \\ &\quad - \frac{M-1}{M} \mathbb{E}^2[y - \mu_{\theta_1}(x)]. \end{aligned} \quad (7)$$

With $M \rightarrow \infty$, $\frac{M-1}{M} \rightarrow 1$,

$$\begin{aligned} \mathbb{E}[\hat{\sigma}_{\epsilon, \theta_3}^2] &\rightarrow \sigma_\epsilon^2 + \mathbb{E}[(y - \mu_{\theta_1}(x))^2] - \mathbb{E}^2[y - \mu_{\theta_1}(x)] \\ &= \sigma_\epsilon^2 + \text{var}[y - \mu_{\theta_1}(x)]. \end{aligned}$$

Our new estimator still overestimates the aleatoric uncertainty by $\text{var}[y - \mu_{\theta_1}(x)]$; however, it is smaller by a term of $\mathbb{E}^2[y - \mu_{\theta_1}(x)]$ compared to the expectation of the previous estimator in Equation (3). This brings us one step closer to capturing aleatoric uncertainty. Further theoretical analysis on other setups is provided in Appendix A.

3.3 Segmented Normalization for Heteroscedastic Setting

In estimating the true realization of the noise, we rely on the law of large numbers and unbiased estimation of the variance theorem (DeGroot and Schervish 2012) that the mean and variance of a large number of i.i.d. samples converge to the expected value of the underlying distribution. However, it should be noted that the normalization process in Equation (5) only ensures that the noise has zero mean and unit variance, without enforcing the independent and identically distributed property. In practice, if the normalization is performed globally, it is possible to encounter situations where the variances of the noise are dependent on the location (e.g. input x) and such imbalances are corrected by the variance estimator $\hat{\sigma}_{\epsilon, \theta_3}(x)$ leading to its inaccurate estimation. To eliminate such an effect, we propose segmented normalization (with details in Appendix B), where we segment the whole training dataset into multiple components and perform normalization individually to prevent the dependence on the local variance.

4 Experiments

To showcase and validate the effectiveness of the DVA technique, we start with a simple regression example. This example helps to investigate the effects of various noise types, such as input and label noise, as well as heteroscedastic and homoscedastic noise. Our approach exhibits strong performance when handling label noise. Detecting input noise typically poses a notable challenge due to the non-linear mapping of the prediction model. Nevertheless, we successfully apply our approach to a problem involving the identification of a dynamic system, in which both input and label are influenced by the same source of noise. To further demonstrate the applicability of our method, we then apply it to real-world tasks involving age prediction and depth estimation. The outcomes underscore the effectiveness of our strategy, even when dealing with large-scale problems.

4.1 Toy Regression Example

We begin with a simplified regression scenario where the underlying mapping function is $f(x) = x(1 + \sin(x))$. We examine various noise configurations and compare our denoising variance attenuation (DVA) technique with the traditional variance attenuation (VA) method. We investigate two different noise configurations:

- Homoscedastic label noise: $\sigma_{\xi}(x) = 0, \sigma_{\epsilon}(x) = a$,
- Heteroscedastic label noise: $\sigma_{\xi}(x) = 0, \sigma_{\epsilon}(x) = a(1 + 0.1x)$,

where $a > 0$ is a constant that represents the magnitude of the standard deviation of the noise. Although we do not make any assumptions on the noise distribution in our analysis, for this example we generate the noise by sampling from a Gaussian distribution with a specified variance (e.g. $N(0, \sigma_{\epsilon}^2(x))$).

The parameterization of both VA variance ($\hat{\sigma}_{\theta_2}^2$) and DVA variance ($\hat{\sigma}_{\epsilon, \theta_3}^2$) follow the noise setup. If the noise is homoscedastic, we use a single scalar parameter with global

normalization of noises, and if the noise is heteroscedastic, we use a neural network to parameterize the variance dependence on x and perform local segmented normalization with 10 segments.

To capture the epistemic uncertainty, we employ two stochastic prediction models: an ensemble model with 5 base learners and Bayesian neural network (BNN, Hernández-Lobato and Adams 2015) with a sampling size of 5. For a fair comparison between the VA and DVA, the prediction model is pre-trained with MSE loss ahead of the variance training. A total of 1000 samples are drawn from $x \in [1, 9]$ for training. Each experiment is performed on 5 random seeds for reproducibility. For more details on the experimental setup, please refer to Appendix C.

In the homoscedastic problem, we increase the global label noise variance from 0.5 to 8.0, and evaluate the performance across different experimental setups. The results, as shown in Table 1, indicate that our method still tends to overestimate the true variance. However, the gap between the estimated and true values is considerably smaller than that of the VA method across all experiment setups.

Method	$a^2 = 0.5$	1.0	8.0
VA+BNN	2.41±0.20	3.20±0.36	10.35±0.32
DVA+BNN	1.12±0.20	1.87±0.25	9.03±0.65
VA+ensemble	2.53± 1.21	2.60±0.68	10.54±1.09
DVA+ensemble	1.77±0.81	2.28±0.61	9.34±0.30

Table 1: Aleatoric uncertainty estimation under homoscedastic label noise (closer to a^2 is better)

In the heteroscedastic problem, we assign a x -dependent noise variance $\sigma_{\epsilon}(x) = a(1 + 0.1x)$ and use squared mean error between the estimation and the true variance as a comparison criterion. Our proposed DVA estimation is closer to the true label variance across almost all experiment setups. The results of a typical experiment when $a = 1$ are shown in Figure 2. The left figure displays the data and model uncertainties, the middle figure visualizes the estimated data uncertainty, and the right figure demonstrates the comparison between the noisy data observations and their corresponding denoised counterparts. Specifically, the VA method tends to overestimate the data uncertainty across the whole input domain, while the DVA method provides a more accurate estimation. The DVA method exhibits a flat line around $x = 8$, this corresponds to the inferior denoising performance in that specific region, which is attributed to the bias in the prediction model.

Method	$a^2 = 0.5$	1.0	8.0
VA + BNN	0.65±0.17	0.32±0.06	0.51±0.18
DVA + BNN	0.45±0.26	0.40±0.61	0.37±0.17
VA + ensemble	1.97± 0.86	1.30± 0.44	1.32±0.32
DVA + ensemble	1.42±1.17	0.82±0.42	1.23±0.39

Table 2: Aleatoric uncertainty estimation under heteroscedastic label noise, squared mean difference to true variance (smaller is better)

In addition to the label noise cases, we also perform




Figure 2: Visualization of the toy example experiment under heteroscedastic label noise

experiments on scenarios with input noise ($\sigma_\xi(x) = a, \sigma_\epsilon(x) = 0$) and defer the results to Appendix D. In general, we observe a trend of increasing noise, but the estimated noise tends to be exaggerated. This is due to the denoised data inferred from neural networks starting to converge towards regions outside the known distribution, as the networks’ ability to accurately predict beyond existing data diminishes.

4.2 Dynamical System Noise Quantification with Neural ODE Identification

For general regression problems with both input and output noise, our method is not effective because it cannot distinguish the weights of the noise sources⁵. However, for dynamical systems which take initial conditions as inputs and trajectory observations as labels, the inputs and labels are subject to the same sources of inherent uncertainty. As a result, we can effectively model this uncertainty using the same set of parameters. We conducted experiments with a synthetic dynamical system ($\dot{x} = x(1 + \sin(x))$) and sampled 100 trajectories with 51 points/trajectory and homoscedastic Gaussian noise as training data. We learn the dynamical system by employing neural ODE (Chen et al. 2018) with Bayesian parameters and estimate the uncertainty by direct MSE loss and DVA method. Table 3 shows the DVA estimation is much closer than the MSE estimation.

Method	$a^2=0.5$	1.0	8.0
MSE+BNN	1.87±0.56	2.69±0.87	11.63±4.82
DVA+BNN	1.06±0.04	1.66±0.10	10.55±0.45

Table 3: Dynamical system aleatoric uncertainty estimation under homoscedastic label noise (closer to a^2 is better)

4.3 Depth Estimation

We demonstrate the applicability of our DVA method to a depth estimation task. The NYU Depth v2 dataset (Silberman et al. 2012) contains 27k RGB-depth image pairs. These images represent indoor scenes such as kitchens, and the pixel labels share the same dimensions as the input. We utilize the VPD (Visual Perception with a pre-trained Diffusion

⁵Consider an identity mapping problem where $y = x$ with $\sigma_\xi^2 = \sigma_\epsilon^2 = 2$. If we apply our denoising method on both input and labels, we may end up with any arbitrary combination of $(\hat{\sigma}_\xi^2, \hat{\sigma}_\epsilon^2)$ that satisfies $\hat{\sigma}_\xi^2 + \hat{\sigma}_\epsilon^2 = 4$.

model, (Zhao et al. 2023)) model trained on the dataset as a prediction model. We add different levels of homoscedastic synthetic Gaussian noise to the label pixels and apply our DVA method to detect the noise level. The proposed DVA method better estimates the aleatoric uncertainty.




Figure 3: Sample examples under $a^2 = 0.25$

Method	$(a^2)=0.01$	0.25	1.0
VA+VPD	0.11± 1e-3	0.34 ± 1e-3	1.09± 1e-3
DVA+VPD	0.02 ± 1e-3	0.28 ± 1e-3	1.06 ± 6e-3

Table 4: Aleatoric uncertainty estimation under homoscedastic input noise in depth estimation

4.4 Age Prediction

We further extend our DVA method to real-world data by performing an age prediction task by utilizing the APPA-REAL database (Agustsson et al. 2017) (Figure 4 shows two samples) consisting of 7591 images. Each image in the database has been assigned an apparent age label based on an average of 38 individual votes per image. We employ a ResNext-50 (32×4d) pretrained on ImageNet, replace the last layer and fine-tune it on the APPA-REAL database with averaged apparent age as labels.

To estimate uncertainty, we append an additional linear layer to the second last layer of the network, which maps its outputs (with a dimension of 2048) to a logarithmic uncertainty measure. We then fine-tune only this linear layer. In Figure 4, we present the VA and DVA estimates of the uncertainty associated with the averaged apparent age, as well as the standard deviation of the voters’ responses for the entire training dataset. Viewing the voters as prediction models that receive images without necessarily reflecting




Figure 4: Left: samples, Right: apparent age std. (from voters) vs estimated data uncertainty

the true underlying age, their forecasted variability encompasses two components: aleatoric uncertainty (discrepancy between the image-based age and the image itself) and epistemic uncertainty (disparity between image-based age and the individual predicted age by the voter). It is worth noting that the epistemic gap between individual voting and group-averaged voting has already been eliminated due to the averaging of labels. Consequently, the aleatoric uncertainty of the averaged labels should be lower than the standard deviation of the voting results indicated by the blue dots in the figure. Our proposed DVA method better approximates this lower bound, particularly for children (whose ages are generally easier to determine accurately than that of adults).

5 Discussion

5.1 Significance of Aleatoric Uncertainty Quantification

Through quantifying aleatoric uncertainty, researchers attain a deeper awareness of the data’s inherent fluctuations within a system or the observation process. This knowledge subsequently empowers researchers to construct more resilient models capable of accommodating and flexibly responding to conditions with built-in uncertainty. This becomes particularly pronounced in the realm of forecasting models, such as dynamical systems, which heavily rely on observational measurements as input. The ability to precisely quantify measurement uncertainty equips researchers with the means to not only evaluate the data trustworthiness but also to effectively incorporate the appropriate level of variability into their forecasts. In essence, aleatoric uncertainty analysis serves as an essential part when dealing with uncertainty-laden systems, fostering both the accuracy and reliability of scientific and engineering procedures.

5.2 Intuitive Explanation of DVA Mechanism

Our DVA method is built upon the principle of the sample variance estimation (DeGroot and Schervish 2012) describing that for a large number of i.i.d. random variables from a distribution with finite variance, the sample variance converges to the true variance of the distribution as the sample size increases. For instance, in the scenario of the zero mean noise case (σ_ϵ^2) with the sampled dataset $\tilde{D} = \{(x_i, y_i + \epsilon_i)\}_{i=1}^M$, the mean and variance of normalized noise $\{\epsilon_i/\sigma_\epsilon\}_{i=1}^M$ should converge to 0 and 1 when $M \rightarrow \infty$. This illustrates the concept that the larger the dataset becomes, the closer the normalized noise’s average and variance align with these specific values.

Keeping that in mind, we parameterize the normalized noises ($\{\hat{\epsilon}_i\}$) and the variance ($\hat{\sigma}_\epsilon(\cdot)$) separately, where the former can be globally normalized to capture the random realization, and the latter accounts for the noise magnitude. After training, the learned variance function can be used for characterizing the noise dependence on x . By introducing the denoised label ($\tilde{y}_i - \hat{\epsilon}_i \hat{\sigma}_\epsilon(x_i)$) into the DVA loss function in Equation (4), we aim to find a set of inferred noise that maximizes the likelihood of the denoised data matching the prediction. However, since the loss function involves the prediction model, the accuracy of the noise recovery depends on the prediction model’s precision. As discussed in Section 3, our approach yields a biased estimator but brings us “one step closer” to the true aleatoric uncertainty compared to the state-of-the-art VA method.

One might argue that the separate parameterization of the normalized noises and variance is over-parameterized, as we can directly model the unnormalized noise. Such approach is trivial because the denoised label will easily converge to the prediction assuming full flexibility of the variance. The neural network parameterization inherently regulates the smoothness of the variance function (Wang, Prakriya, and Jha 2022) and limited a finite number of limited linear regions (Takai, Sannai, and Cordonnier 2021), preventing trivial solutions and encouraging the learned variance to represent the true uncertainties.

5.3 Limitations

Despite the benefits, our approach has certain limitations. First, we introduce additional variables, requiring more memory and computational cost⁶. Second, DVA assumes the zero mean of the sampled noise with the law of large numbers which requires a large sample size. When addressing heteroscedastic noise, we employ a segmented normalization technique to minimize the impact of data-dependent variance on the estimated normalized noise. Alternatively, the noise parameterization should be matched to the intrinsic heteroscedasticity of the system. In addition, this paper primarily concentrates on regression tasks, and the adaptation to classification tasks is beyond the scope of this work.

6 Conclusion

In this study, we point out that the variance attenuation method widely used in deep learning, overestimates the aleatoric uncertainty. We propose a novel approach for estimating aleatoric uncertainty from data by considering the inferred noise and variance separately. Although our method still represents a biased estimation, our theoretical analysis demonstrates that it offers an improvement over the prior method in capturing the true aleatoric uncertainty, and we validate the improvement by quantitative experiments. In conclusion, our work presents a valuable contribution to the field of uncertainty estimation for deep neural networks and has the potential to serve as a standard procedure in aleatoric uncertainty quantification for machine learning tasks.

⁶The computation time increases slightly (34.9s/epoch for DVA vs 34.7s/epoch for VA for training the age prediction task) as most of the back-propagation operations are shared.

Acknowledgements

This research is based upon the work supported by MIT-IBM Watson AI Lab. The authors thank Ching-Yun Irene Ko and Nima Dehmamy for the helpful discussions.

References

- Agustsson, E.; Timofte, R.; Escalera, S.; Baro, X.; Guyon, I.; and Rothe, R. 2017. Apparent and real age estimation in still images with deep residual regressors on appa-real database. In *2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)*, 87–94. IEEE.
- Amini, A.; Schwarting, W.; Soleimany, A.; and Rus, D. 2020. Deep evidential regression. *Advances in Neural Information Processing Systems*, 33: 14927–14937.
- Blundell, C.; Cornebise, J.; Kavukcuoglu, K.; and Wierstra, D. 2015. Weight uncertainty in neural network. In *International conference on machine learning*, 1613–1622. PMLR.
- Chen, R. T.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018. Neural ordinary differential equations. *Advances in neural information processing systems*, 31.
- Choi, J.; Elezi, I.; Lee, H.-J.; Farabet, C.; and Alvarez, J. M. 2021. Active learning for deep object detection via probabilistic modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10264–10273.
- DeGroot, M. H.; and Schervish, M. J. 2012. *Probability and statistics*. Pearson Education.
- Der Kiureghian, A.; and Ditlevsen, O. 2009. Aleatory or epistemic? Does it matter? *Structural safety*, 31(2): 105–112.
- Gal, Y.; and Ghahramani, Z. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, 1050–1059. PMLR.
- Gal, Y.; et al. 2016. Uncertainty in deep learning.
- Gelfand, I. M.; Silverman, R. A.; et al. 2000. *Calculus of variations*. Courier Corporation.
- Hernández-Lobato, J. M.; and Adams, R. 2015. Probabilistic backpropagation for scalable learning of bayesian neural networks. In *International conference on machine learning*, 1861–1869. PMLR.
- Hüllermeier, E.; and Waegeman, W. 2021. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110: 457–506.
- Kurakin, A.; Goodfellow, I. J.; and Bengio, S. 2018. Adversarial examples in the physical world. In *Artificial intelligence safety and security*, 99–112. Chapman and Hall/CRC.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30.
- Liberzon, D. 2011. *Calculus of variations and optimal control theory: a concise introduction*. Princeton university press.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.
- Nix, D.; and Weigend, A. 1994. Estimating the mean and variance of the target probability distribution. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*, volume 1, 55–60 vol.1.
- Parker, W. S. 2013. Ensemble modeling, uncertainty and robust predictions. *Wiley Interdisciplinary Reviews: Climate Change*, 4(3): 213–223.
- Seeböck, P.; Orlando, J. I.; Schlegl, T.; Waldstein, S. M.; Bogunović, H.; Klimescha, S.; Langs, G.; and Schmidt-Erfurth, U. 2019. Exploiting epistemic uncertainty of anatomy segmentation for anomaly detection in retinal OCT. *IEEE transactions on medical imaging*, 39(1): 87–98.
- Seitzer, M.; Tavakoli, A.; Antic, D.; and Martius, G. 2022. On the pitfalls of heteroscedastic uncertainty estimation with probabilistic neural networks. *arXiv preprint arXiv:2203.09168*.
- Silberman, N.; Hoiem, D.; Kohli, P.; and Fergus, R. 2012. Indoor segmentation and support inference from rgbd images. *ECCV (5)*, 7576: 746–760.
- Stirn, A.; Wessels, H.; Schertzer, M.; Pereira, L.; Sanjana, N.; and Knowles, D. 2023. Faithful Heteroscedastic Regression with Neural Networks. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*.
- Takai, Y.; Sannai, A.; and Cordonnier, M. 2021. On the number of linear functions composing deep neural network: Towards a refined definition of neural networks complexity. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*.
- Valdenegro-Toro, M.; and Mori, D. S. 2022. A deeper look into aleatory and epistemic uncertainty disentanglement. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 1508–1516. IEEE.
- Wang, Z.; Prakriya, G.; and Jha, S. 2022. A Quantitative Geometric Approach to Neural-Network Smoothness. In *Advances in Neural Information Processing Systems*.
- Zhao, W.; Rao, Y.; Liu, Z.; Liu, B.; Zhou, J.; and Lu, J. 2023. Unleashing Text-to-Image Diffusion Models for Visual Perception. *arXiv preprint arXiv:2303.02153*.