

Normalizing Flow Ensembles for Rich Aleatoric and Epistemic Uncertainty Modeling

Lucas Berry, David Meger

McGill University, Centre for Intelligent Machines, Montreal, Canada

lucas.berry@mail.mcgill.ca

Abstract

In this work, we demonstrate how to reliably estimate epistemic uncertainty while maintaining the flexibility needed to capture complicated aleatoric distributions. To this end, we propose an ensemble of Normalizing Flows (NF), which are state-of-the-art in modeling aleatoric uncertainty. The ensembles are created via sets of fixed dropout masks, making them less expensive than creating separate NF models. We demonstrate how to leverage the unique structure of NFs, base distributions, to estimate aleatoric uncertainty without relying on samples, provide a comprehensive set of baselines, and derive unbiased estimates for differential entropy. The methods were applied to a variety of experiments, commonly used to benchmark aleatoric and epistemic uncertainty estimation: 1D sinusoidal data, 2D windy grid-world (*Wet Chicken*), *Pendulum*, and *Hopper*. In these experiments, we setup an active learning framework and evaluate each model’s capability at measuring aleatoric and epistemic uncertainty. The results show the advantages of using NF ensembles in capturing complicated aleatoric while maintaining accurate epistemic uncertainty estimates.

Introduction

One common decomposition of uncertainty is aleatoric and epistemic (Hora 1996; Der Kiureghian and Ditlevsen 2009; Hüllermeier and Waegeman 2021). Aleatoric uncertainty refers to the inherent randomness in the outcome of an experiment, while epistemic uncertainty can be described as ignorance or a lack of knowledge. The important distinction between the two is that epistemic uncertainty can be reduced by the acquisition of more data while aleatoric cannot. Our goal in this paper is to learn aleatoric distributions in high dimensions, with arbitrary distributional form, while also tracking epistemic uncertainty due to non-uniform data sampling.

Normalizing Flows (NFs) have been shown to be effective at capturing highly expressive aleatoric uncertainty with little prior knowledge (Kingma and Dhariwal 2018; Rezende and Mohamed 2015). This is done by transforming a base distribution via a series of nonlinear bijective mappings, and can model complex heteroscedastic and multi-modal noise. Robotic systems display such noise, as robots interact with

nonlinear stochastic dynamics often in high dimensions. To further complicate the problem of dynamics modeling, data collection for many robot systems can be prohibitively expensive. Therefore, an active learning framework is usually adopted to iteratively collect data in order to most efficiently improve a model. In order to apply such a framework, ensembles have been employed to capture epistemic uncertainty for deep learning models (Gal, Islam, and Ghahramani 2017).

In this paper, we utilize NFs’ ability to capture rich aleatoric uncertainty and extend such models to epistemic uncertainty estimation with ensembles. We then use our NF models to tackle epistemic uncertainty for regression tasks. The contributions of this work are as follows:

- We develop two methods for estimating uncertainty for NFs, derive unbiased estimates for said models, and leverage the base distribution to reduce the sampling burden on the estimation of uncertainty.
- We leverage memory-efficient ensembling by creating ensembles via fixed dropout masks and apply them to NFs.
- We demonstrate the usefulness of uncertainty estimation on an array of previously proposed benchmarks (Deweg et al. 2018) and novel settings for active learning problems.

Problem Statement

Given a dataset $\mathcal{D} = \{x_i, y_i\}_{i=1}^N$, where $x_i \in \mathbb{R}^K$ and $y_i \in \mathbb{R}^D$, we wish to approximate the conditional probability $p_{Y|X}(y|x)$. Our models therefore take as input x and output the approximate conditional probability density function, $p_{Y|X}(y|x) = f_{\theta}(y, x)$, where θ is a set of parameters to be learned. In the experiments that follow, the ground truth distribution, $p_{Y|X}(y|x)$, is not assumed homoscedastic nor do we put restrictions on its shape.

In order to capture uncertainty, we utilize the information based criteria proposed in Houthby et al. (2011) to estimate uncertainty. Let $H(\cdot)$ denote the differential entropy of a random variable, $H(X) = E[-\log(p_X(x))]$, and $W \sim p(w)$ index different models in a Bayesian framework or an ensemble. Then one can define epistemic uncertainty, $I(y^*, W)$, as the difference of total, $H(y^*|x^*)$, and aleatoric,

$E_{p(w)} [H(y^*|x^*, w)]$, uncertainty,

$$I(y^*, W) = H(y^*|x^*) - E_{p(w)} [H(y^*|x^*, w)], \quad (1)$$

where $I(\cdot)$ is the mutual information and y^*, x^* denote values not seen during training.

These uncertainty estimates can be used for the problem of active learning, where we iteratively add data points to our training dataset in order to improve our model’s performance as much as possible with each data point selected (Settles 2009; MacKay 1992). When doing active learning, one attempts to find the x^* that maximizes Equation (1) and query new points in that region. Intuitively, we are looking for the point x^* which has high total uncertainty, $H(y^*|x^*)$, but each model has low uncertainty, $H(y^*|x^*, w)$.

Background

NFs were chosen to capture aleatoric uncertainty as they are capable of representing flexible distributions. Ensembles were chosen as they are a computationally efficient way of estimating epistemic uncertainty for deep learning.

Normalizing Flows

NFs are non-parametric models that have been shown to be able to fit flexible multi-modal distributions (Tabak and Vanden-Eijnden 2010; Tabak and Turner 2013). They do so by transforming a simple continuous distribution (e.g. Gaussian, Beta, etc.) into a more complex one via the change of variable formula. These transformations make it so one can score and sample from the fitted distribution, thus allowing NFs to be applied to a multitude of problems. Let B be a D -dimensional continuous random vector with $p_B(b)$ as its density function and let $Y = g(B)$ where g is invertible, g^{-1} exists, and both g and g^{-1} are differentiable. Using the change of variable formula, we can write the distribution of Y as,

$$p_Y(y) = p_B(g^{-1}(y)) |\det(J(g^{-1}(y)))|, \quad (2)$$

where $J(\cdot)$ is the Jacobian and \det refers to the determinant. The first term in the product on the RHS of Equation (2) is what changes the shape of the distribution while $|\det(J(g^{-1}(y)))|$ normalizes it, forcing it to integrate to one.

NF models can be learned by making $g(\cdot)$ parameterized by θ , i.e., $g_\theta(\cdot)$, and then learned via the log likelihood. In addition, NFs can be made conditional (Winkler et al. 2019; Ardizzone et al. 2019). Following the framework of Winkler et al. (2019), Equation (2) becomes,

$$p_{Y|X}(y|x) = p_{B|X}(g_\theta^{-1}(y, x)) \times |\det(J(g_\theta^{-1}(y, x)))|, \quad (3)$$

$$\log(p_{Y|X}(y|x)) = \log(p_{B|X}(g_\theta^{-1}(y, x))) + \log(|\det(J(g_\theta^{-1}(y, x)))|). \quad (4)$$

Note that now $g_\theta : Y \times X \mapsto B$ and NF models commonly chain together multiple bijective mappings to make their models more flexible. When fitting a NF, one typically optimizes the negative log likelihood in Equation (4) over mini batches. For a complete overview of NFs, please refer to Papamakarios et al. (2021).

Ensembles

Ensembles use multiple models to obtain better predictive performance and to measure uncertainty. The conditional model can be written as,

$$f_\theta(y, x) = \sum_{w=1}^M \pi_w f_{\theta_w}(y, x), \quad (5)$$

where M and π_w are the number of model components and the component weights, respectively.

Ensembles are typically generated in one of two ways: randomization (Breiman 2001) and boosting (Freund and Schapire 1997). Randomization has been preferred method for deep learning models (Lakshminarayanan, Pritzel, and Blundell 2017). Each model is randomly initialized and then at each training step for a model w , a sample, with replacement, is drawn from a training set \mathcal{D} and then a step in the direction of the gradient is taken. This creates diversity as each ensemble component is exposed to a different portion of \mathcal{D} at each step in the gradient.

Normalizing Flows Ensembles

We propose two approaches for creating NF ensembles, Nflows Out and Nflows Base, both of which rely on neural spline bijective mappings, $g_\theta(\cdot)$, as they have been shown to be very flexible distribution approximators (Durkan et al. 2019a,b). Both bagging and random initialization are utilized in training our ensemble components. Each ensemble component is created via fixed dropout masks (Durasov et al. 2021), which reduces the memory and computation cost of our method.

Nflows Out

Nflows Out creates an ensemble in the nonlinear transformations g ’s,

$$p_{Y|X,W}(y|x, w) = f_{\theta_w}(y, x) = p_{B|X}(g_{\theta_w}^{-1}(y, x)) \times |\det(J(g_{\theta_w}^{-1}(y, x)))|. \quad (6)$$

The base distribution is static for each component and the bijective transformation is where the component variability lies. The network g_{θ_w} outputs the parameters of cubic spline and thus each ensemble component produces a different cubic spline. **By including the complex aleatoric uncertainty prediction of an NF as well as the ability of dropout ensembles to capture uncertainty over learned parameters, our method bridges the state-of-the-art of aleatoric representation and epistemic uncertainty estimation.** These new capabilities allows Nflows Out to be applied to new decision making tasks.

Using Nflows Out, we can approximate Equation (1). We employ Monte Carlo sampling to estimate the quantities of interest. Thus, total uncertainty is estimated as,

$$H(y^*|x^*) = -E [\log(p_{Y|X}(y^*|x^*))] \quad (7)$$

$$\approx -\frac{1}{N} \sum_{n=1}^N \log(p_{Y|X}(y_n|x^*)), \quad (8)$$

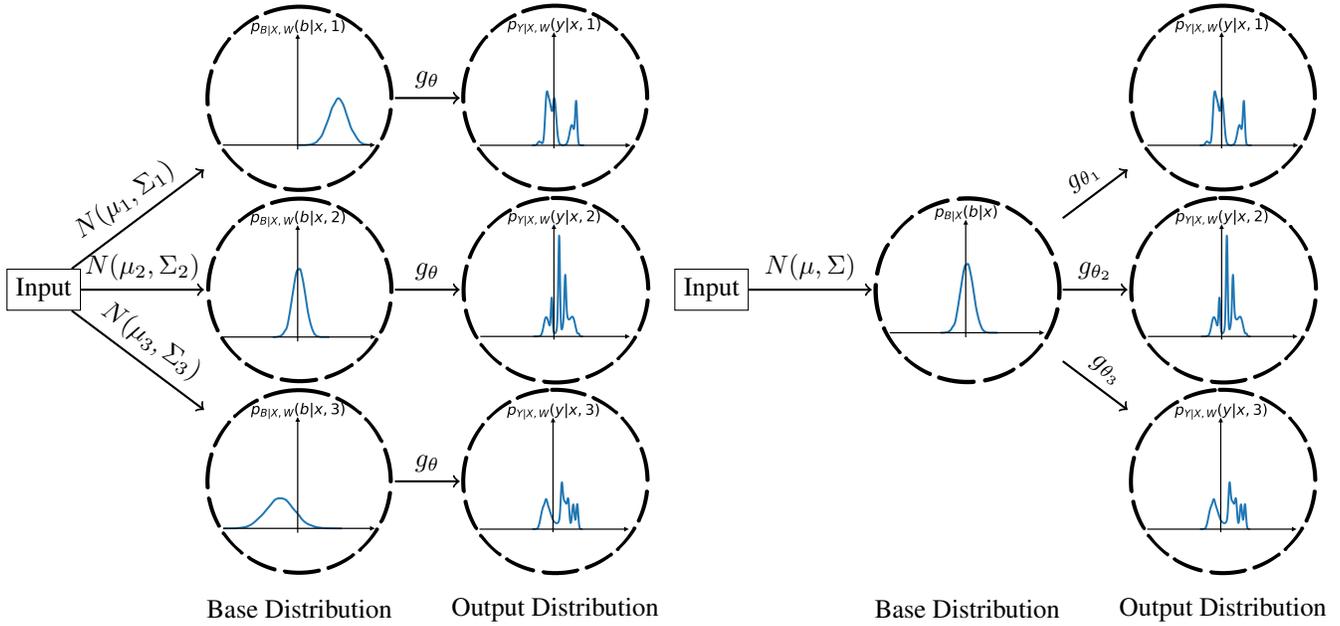


Figure 1: The left figure illustrates Nflows Base and the right figure depicts Nflows Out on an ensemble of 3 components with one bijective transformation.

where N is the number of samples drawn. For a given x^* we sample N points from $p_{B|X}$ and then randomly select an ensemble component g_w to transform each point. The aleatoric uncertainty, $E_{p(w)}[H(y^*|x^*, w)]$, of Equation (1) is calculated in a similar fashion,

$$\begin{aligned}
 & E_{p(w)}[H(y^*|x^*, w)] \\
 &= -\frac{1}{M} \sum_{w=1}^M E[\log(p_{Y|X,W}(y|x^*, w))] \quad (9) \\
 &\approx -\frac{1}{M} \sum_{w=1}^M \frac{1}{N_w} \left(\sum_{n_w=1}^{N_w} \log(p_{Y|X,W}(y_{n_w}|x^*, w)) \right). \quad (10)
 \end{aligned}$$

For a given x^* and w , we sample N_w points from $p_{B|X}$ and then transform the samples according to g_{θ_w} . Monte Carlo estimation can suffer from the curse of dimensionality (Rubinstein and Glynn 2009). Experiments and an intuitive proof detailing Monte Carlo’s estimation limitations in high dimensions is contained in the Appendix A.8.

Nflows Base

In an attempt to alleviate the model’s reliance on sampling, we propose a second ensembling technique for NFs, Nflows Base. The ensembles are created in the base distribution,

$$p_{Y|X,W}(y|x, w) = f_{\theta_w}(y, x) = p_{B|X,W}(g_{\theta}^{-1}(y, x)) \times |\det(J(g_{\theta}^{-1}(y, x)))|, \quad (11)$$

where $p_{B|X,W}(b|x, w) = N(\mu_w, \Sigma_w)$, μ_w and Σ_w are estimated via a neural network with fixed dropout masks. In

Nflows Base, the bijective mapping is static for each component while the base distribution varies. Figure 1 depicts Nflows Base on the left and Nflows Out on the right.

Nflows Base has advantages when estimating uncertainty as the component variability is contained in the base distribution and thus analytical formulae can be used to approximate aleatoric uncertainty. Equation (10) becomes,

$$\approx \frac{1}{M} \sum_{w=1}^M \frac{1}{2} \log(\det(2\pi\Sigma_w)). \quad (12)$$

These advantages are seen in memory reduction, computationally efficiency, and estimation error. Let N_x denote the number of x^* to estimate uncertainty for. Then Nflows Out needs to sample TS points where $TS = N_x N_w M$ and the samples are also used to estimate Equation (8) ($N_{out} = N_w M$). On the other hand, for Nflows Base, $TS = N_x N_{base}$ points are needed to estimate uncertainty, where $N_{base} = \frac{N_{out}}{M}$. This reduces the number of samples needed to drawn by a factor M . Please refer to the Appendix A.8 for analysis of time savings. In addition to reducing the sampling required, there is less estimation error as component-wise entropy can be computed directly from Equation (12) instead of through sampling. For nonparametric models, it is uncommon to estimate aleatoric uncertainty without sampling as the output distribution does not hold a parametric form.

Baseline Models

We have included several baselines and compared each method’s ability to measure aleatoric and epistemic uncertainty. These baselines are detailed below.

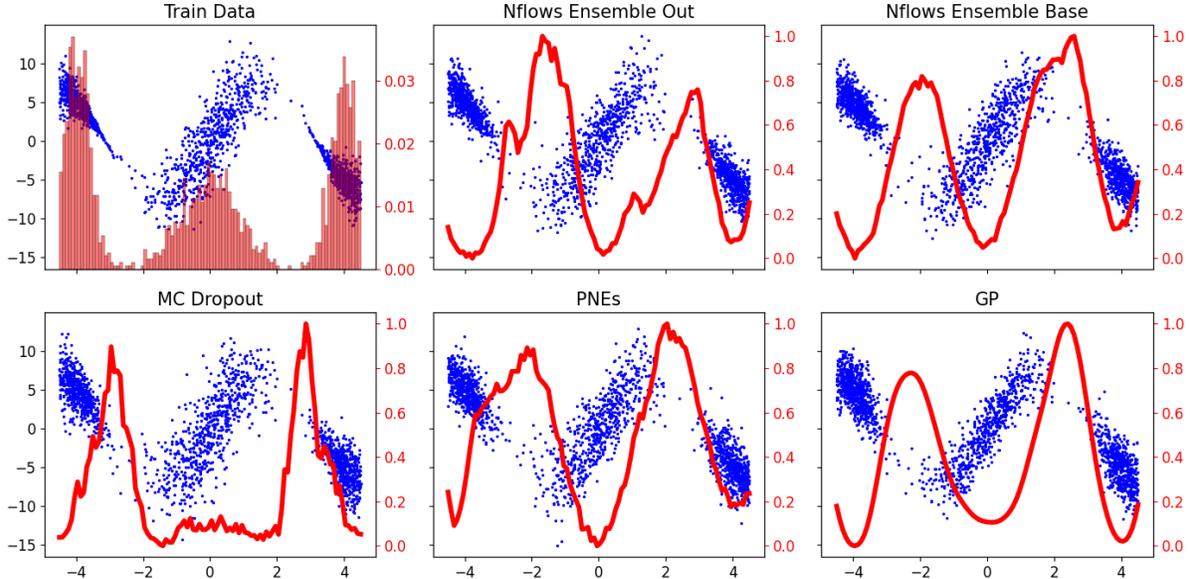


Figure 2: Samples from each model at the corresponding x -values on the *Hetero* Environment are represented by the blue dots and left y -axes. Epistemic uncertainty is represented by the red curve and the right y -axes. Note that the probability density is in the top left graph instead of epistemic uncertainty

Probabilistic Network Ensembles

Probabilistic network ensembles (PNEs) have been shown to be a powerful tool to measure uncertainty for neural networks (Chua et al. 2018; Kurutach et al. 2018). We are particularly interested in capturing their capabilities at measuring aleatoric and epistemic uncertainty in supervised learning tasks. PNEs were created with fixed dropout masks with each component modeling a Gaussian,

$$p_{Y|X}(y|x) = \frac{1}{M} \sum_{w=1}^M p_{Y|X,W}(y|x, w). \quad (13)$$

The model is then trained via negative log likelihood, with randomly initialized weights and bootstrapped samples from the training set. We estimate epistemic uncertainty for PNEs via the same method for Nflows Base, Equations (8) and (12).

Monte Carlo Dropout

In addition to ensembles, other Bayesian approximations exist in the deep learning community. MC dropout is one of the more prevalent and commonly used Bayesian approximations (Gal and Ghahramani 2016; Gal, Islam, and Ghahramani 2017; Kirsch, Van Amersfoort, and Gal 2019). MC dropout creates ensembles via dropout during training and uses dropout at test time to estimate uncertainty. The output distribution is therefore similar to PNEs in Equation (13). However, we cannot sample each mask at test time, and thus a random sample of masks needs to be drawn. Therefore, when estimating uncertainty for MC dropout we first sample a set of masks and then sample each Gaussian corresponding to each mask. After which, we apply Equations (8) and

(12) to measure uncertainty. Note the each mask has equal probability, as the dropout probability was set to 0.5.

Gaussian Processes

Gaussian Processes (GPs) are Bayesian models widely used to quantify uncertainty (Rasmussen 2003). A GP model can be fully defined by its mean function $m(\cdot)$ and a positive semidefinite covariance function/kernel $k(\cdot, \cdot)$ of a real process $f(x)$,

$$m(x) = E[f(x)],$$

$$k(x, x') = E[(f(x) - m(x))(f(x') - m(x')))]. \quad (14)$$

Choosing the mean and covariance function allows a practitioner to input prior knowledge into the model. From these choices, the predictive posterior becomes,

$$E_f[x^*] = m(x^*) = k_*^T (K + \sigma_w^2 I)^{-1} y \quad (15)$$

$$\text{var}_f[x^*] = k_{**} - k_*^T (K + \sigma_w^2 I)^{-1} k_*, \quad (16)$$

where $k_* = k(X, x^*)$, $k_{**} = k(x^*, x^*)$, K is the kernel matrix with entries $K_{ij} = k(x_i, x_j)$, and σ_w^2 is a noise variance hyperparameter. Note that X and y refer to the complete set of training data. GPs place a probability distribution over functions that are possible fits over the data points. This distribution over functions is used to express uncertainty and is used to quantify epistemic uncertainty.

Experiments

We first evaluate each method on two 1D environments, previously proposed (Depeweg et al. 2018), to compare whether each method can capture multi-modal and heteroscedastic noise while measuring epistemic uncertainty. In addition,

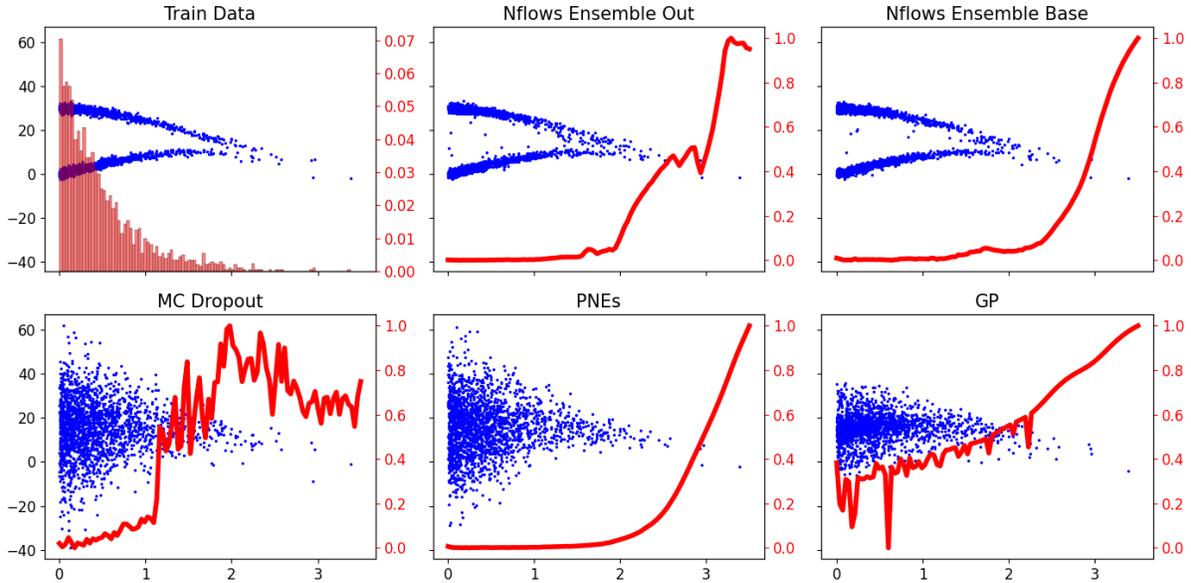


Figure 3: Samples from each model at the corresponding x-values on the *Bimodal* Environment are represented by the blue dots and left y-axes. Epistemic uncertainty is represented by the red curve and the right y-axes. Note that the probability density is in the top left graph instead of epistemic uncertainty

we provide 5 active learning problems, both 1D and multi-dimensional. Both openai gym and nflows libraries were utilized with minor changes (Brockman et al. 2016; Durkan et al. 2020). The model from Depeweg et al. (2016, 2018) is Bayesian Neural Network that is limited to mixture Gaussians and is thus not as expressive NFs for aleatoric uncertainty. In addition, there is no source code for it and was not included as a baseline, though we have included MC dropout as a close comparison. For all model hyper-parameters please refer to the Appendix A.1 and the code can be found at https://github.com/nwafpt23/nflows_epistemic.

Data

In order to assess our uncertainty estimation, we evaluate on two 1D benchmarks, *Hetero* and *Bimodal* (Depeweg et al. 2018). The *Hetero* data can be seen in Figure 2 in the top left pane. There are two valleys of low data density displayed in the red bar chart where one would expect high epistemic uncertainty. The *Bimodal* data can be seen in Figure 3 in the top left pane. Density of data drops off when moving right along the x-axis, thus you would expect epistemic uncertainty to grow as x does. For complete details involved in generating the data, please refer to the Appendix A.2. Both 1D environments were included to provide proof of concept for uncertainty estimation and visualizations.

In addition to the 1D environments, we validated our methods across three multi-dimensional environments. Trajectories were gathered from an agent and then the transition function for each environment was estimated, $f(s_t, a_t) = s_{t+1}$. The first, *Wet Chicken* (Trep. 1994), is commonly used to evaluate a model’s capacity to fit multi-modal and heteroscedastic noise (Depeweg et al. 2018, 2016). It simulates

a canoe approaching the edge of a waterfall. The paddlers are enticed to the edge of the waterfall as that is the region with the most fish, but as they get closer the probability increases that they fall over the edge and start over. Hence towards the edge of the waterfall the transitions become bimodal. The dynamics are naturally stochastic and are governed by the equations contained in the Appendix A.3. *Wet Chicken* was included to assess uncertainty estimation on an intrinsically stochastic multi-dimensional environment.

Moreover, we evaluated all methods on *Pendulum* (Brockman et al. 2016) and *Hopper* (Todorov, Erez, and Tassa 2012). These environments are included because they are commonly used in benchmarking and provide us a higher dimensional output space to validate our methods. These environments are inherently nonstochastic and thus noise is injected into the dynamics in order to produce multi-modal noise. The noise is applied to each action, $a'_t = a_t + a_{max}\epsilon$, where epsilon is drawn from a multi-modal distribution and a_{max} refers to the maximum value in the action space. Note that the parameters used to create the noise distribution are included in the Appendix A.4 and that a_t is recorded in the replay buffer, not a'_t .

1D Fits

First, we turn to our 1D environments to give motivation and empirical proof of the epistemic uncertainty estimation. In Figure 2, we can see that each model does a reasonable job of capturing the narrow and widening of the ground truth distribution for the *Hetero* environment. Moreover, Figure 2 displays each model’s capacity to estimate epistemic uncertainty in the red curve and right y-axis. Each model does a good job of capturing higher epistemic uncertainty in the

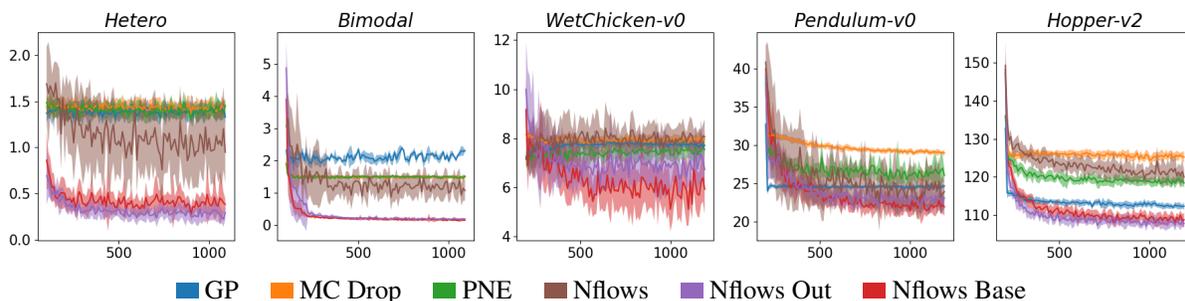


Figure 4: Mean KL divergence on 50 randomly sampled test set inputs as data was added to the training sets ($\frac{1}{50} \sum_{i=1}^{50} D_{KL}(P_i \parallel Q_i)$, where P_i is the ground truth conditional distribution and Q_i is the model’s conditional distribution conditioned on x_i).

two regions with fewer data points. It is not surprising to see GPs, PNEs, and MC dropout perform well on *Hetero*, as each method has been shown to accurately fit heteroscedastic Gaussian noise.

While the baseline methods were sufficiently capable of capturing the *Hetero* data, this not the case for the *Bimodal* setting. Figure 3 shows that none of the non-NF models can fit both modes. This is to be expected as GPs, PNEs and MC dropout fit a Gaussian at each x . While ensembles such as PNEs and MC dropout can theoretically capture multiple modes, there is no guarantee of this. To guarantee capturing multiple modes, deep learning engineering is required. One can either manipulate of the loss function or separate the training data according to each mode, both of which may require prior knowledge. On the other hand, NFs are able to fit multiple modes directly via the log likelihood. Note that some fine tuning of the number of transformations and transformation types is required.

In addition to the aleatoric uncertainty estimation, Figure 3 displays each model’s ability to capture epistemic uncertainty on the *Bimodal* data with the red curve on the right y-axis. **Nflows Out and Nflows Base are the only models to capture both modes while maintaining accurate epistemic uncertainty estimates.** Each model shows the pattern of increasing uncertainty where the data is more scarce with MC dropout having the most trouble displaying this pattern. MC dropout had the most difficulty to estimate epistemic uncertainty on most tasks. This is likely the case because, as opposed to the PNEs or the NFs methods, the number of ensemble components is generally quite larger (2^n where n is the number of neurons in your network). Therefore, when estimating Equation (1), a subset of masks needs to be sampled, leading to less stable estimates of uncertainty.

Active Learning

In order to assess each model’s capability at utilizing epistemic and capturing flexible aleatoric uncertainty, we provide an active learning experimental setup. For both 1D and multi-dimensional experiments, each model started with 100 and 200 data points, respectively. At each epoch, the models sampled 1000 new inputs and kept the 10 with the highest epistemic uncertainty. The models are then evaluated

by sampling 50 inputs from the test set and averaging the Kullback-Liebler (KL) divergence for the ground truth distribution at those points and model’s distributions,

$$\frac{1}{50} \sum_{i=1}^{50} D_{KL}(P_i \parallel Q_i) \quad (17)$$

Note that the KL divergences reported were estimated via samples using the k -nearest-neighbor method (Wang, Kulkarni, and Verdu 2009). KL divergence was chosen as an evaluation metric as we are most interested in distributional fit. In order to ensure variation, the training and test were gathered via different processes; refer to the Appendix A.5 for more details. All experiments were run across 10 seeds and their mean and standard deviation are reported.

Figure 4 displays the performance of each method as the training size increases. **For each data setting, the NF ensemble models reach lower KLs, thus verifying that they can leverage epistemic uncertainty estimation to learn more expressive aleatoric uncertainty faster.** In some cases other models provided better results with small number of data points, this information is conveyed in Table 1, with the best-performing models in bold at different acquisition epochs. Note that in addition to the baselines discussed, we included an NF with no ensembles, using total entropy as the acquisition function.

Related Work

Using Bayesian methods, researchers have developed information based criterion for the problem of active learning using Gaussian Processes (GPs) on classification problems (Houlsby et al. 2011). Researchers have leveraged said information based criterion for uncertainty estimation with Bayesian neural networks (Gal, Islam, and Ghahramani 2017; Kendall and Gal 2017; Kirsch, Van Amersfoort, and Gal 2019). These works extended previous epistemic uncertainty estimation by leveraging dropout to estimate uncertainty on image classification for neural networks. In contrast, our work estimates epistemic uncertainty on a harder output space. The experiments contained in this paper were conducted on regression problems where the output is drawn from continuous distributions in 1-11 dimensions, whereas

Env	Acq Batch	GP	PNEs	MC Drop	Nflows	Nflows Out	Nflows Base
<i>Hetero</i>	10	1.43±0.12	1.39±0.06	1.43±0.1	1.54±0.23	0.48±0.09	0.51±0.17
	25	1.43±0.11	1.44±0.1	1.43±0.08	1.3±0.45	0.31±0.09	0.43±0.11
	50	1.41±0.07	1.39±0.07	1.46±0.1	1.39±0.21	0.27±0.08	0.36±0.09
	100	1.33±0.06	1.45±0.08	1.44±0.07	0.95±0.33	0.3±0.08	0.38±0.06
<i>Bimodal</i>	10	2.23±0.18	1.49±0.06	1.49±0.03	1.26±0.81	0.74±0.76	0.36±0.07
	25	2.02±0.07	1.46±0.05	1.47±0.04	1.21±0.6	0.24±0.04	0.2±0.03
	50	1.97±0.07	1.5±0.03	1.49±0.02	1.2±0.45	0.22±0.03	0.18±0.03
	100	2.3±0.03	1.51±0.05	1.49±0.05	1.07±0.32	0.18±0.02	0.14±0.02
<i>Wet Chicken</i>	10	7.61±0.21	7.14±0.24	7.93±0.27	8.04±1.07	7.83±1.2	7.67±1.36
	25	7.73±0.17	7.49±0.47	8.02±0.3	8.19±1.03	6.64±0.79	6.25±0.98
	50	7.81±0.12	7.61±0.44	7.95±0.18	8.12±0.78	6.51±0.56	5.86±0.92
	100	7.71±0.18	7.55±0.46	7.97±0.28	8.06±0.7	6.73±1.06	5.93±1.01
<i>Pendulum-v0</i>	10	24.56±0.21	27.29±0.73	31.04±0.41	26.45±4.61	27.62±1.87	26.07±2.22
	25	24.52±0.3	26.43±1.05	30.11±0.29	24.86±3.66	24.13±1.18	23.97±2.16
	50	24.68±0.26	26.47±1.19	29.6±0.26	24.44±3.04	22.86±1.63	22.45±1.29
	100	24.67±0.17	26.04±0.94	29.0±0.39	23.9±0.93	23.09±1.56	21.93±1.17
<i>Hopper-v2</i>	10	114.8±0.97	122.42±1.22	125.66±0.98	126.87±2.83	112.79±1.18	114.6±2.14
	25	113.29±0.62	120.3±1.56	125.99±0.84	123.93±1.84	109.31±2.0	110.64±1.59
	50	112.98±0.71	119.82±1.66	125.65±0.88	122.36±1.36	108.59±1.04	109.44±1.87
	100	112.27±1.0	118.4±1.21	125.21±1.36	119.97±1.91	107.74±0.99	108.71±1.83

Table 1: KL Divergence of 50 randomly sampled test inputs ground truth distribution and the corresponding model distribution. Experiments were across ten different seeds and the results are expressed as mean plus minus one standard deviation. Best models are bolded.

the previous works applied their methods to classification problems, a 1D categorical output.

NFs have been shown to be poor estimators of epistemic uncertainty (Kirichenko, Izmailov, and Wilson 2020; Zhang, Goldstein, and Ranganath 2021). Researchers have argued that NFs, inability to differentiate between in and out distribution samples via their likelihood is a methodological shortcoming. Some have found workarounds to this problem, specifically in the sphere of ensembles (Choi, Jang, and Alemi 2018). Ensembles have been shown to be a powerful tool in the belt of a machine learning practitioner by leveraging similar uncertainty quantification benefits to their Bayesian cousins but at a smaller computational footprint (Lakshminarayanan, Pritzel, and Blundell 2017). The work regarding NFs and uncertainty have focused on image generation and unsupervised learning. Our methods differ, as we consider supervised learning problems (Winkler et al. 2019). In addition, the ensembles created in Choi, Jang, and Alemi (2018) contrast with ours as we leverage the base distribution to estimate our uncertainty, use mutual information instead of WAIC and create our ensembles with less memory.

In addition to the examples discussed, work has been done to quantify epistemic uncertainty for regression problems (Depeweg et al. 2018; Postels et al. 2020). Depeweg et al. (2016) method’s relied on Bayesian approximations to neural networks which modeled mixture of Gaussians and demonstrated their ability to capture uncertainty on three environments. Our work expands on this, by developing more expressive NF models for uncertainty estimation. Postels

et al. (2020) develop theory to show how latent representations of a proxy network can be used to estimate uncertainty. They use the proxy network’s latent representation as their conditioner. In contrast, we show to how to do this with one NF model and how to leverage the base distribution to be more sample efficient. This study provided an array of multi-modal problems while Postels et al. (2020) considered a single uni-modal problem. Our work expands on both these papers by providing a comprehensive analysis of different baseline methods, comparing their uncertainty quantification and including higher dimensional data. In addition, we provide a full active learning experimental setup and develop new NF frameworks for measuring uncertainty that are more sample efficient and have lower memory costs.

Conclusion

In this paper, we introduced NF ensembles via fixed dropout masks and demonstrated how they can be used efficiently to quantify uncertainty. In doing so, we show how to leverage the base distribution to estimate uncertainty more sample efficiently. Moreover, Nflows Base shows that one can accurately measure uncertainty in the base distribution space. We empirically show that our models outperform the state-of-the-art in capturing the combination of aleatoric and epistemic uncertainty on 5 regression tasks. This paper shows that NF ensembles are an expressive model for aleatoric uncertainty while keeping the benefits of previous methods for capturing epistemic uncertainty.

Acknowledgements

Prof. Meger was supported by the National Sciences and Engineering Research Council (NSERC), through the NSERC Canadian Robotics Network (NCRN).

References

- Ardizzone, L.; Lüth, C.; Kruse, J.; Rother, C.; and Köthe, U. 2019. Guided image generation with conditional invertible neural networks. *arXiv preprint arXiv:1907.02392*.
- Breiman, L. 2001. Random forests. *Machine learning*, 45(1): 5–32.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. Openai gym. *arXiv preprint arXiv:1606.01540*.
- Choi, H.; Jang, E.; and Alemi, A. A. 2018. Waic, but why? generative ensembles for robust anomaly detection. *arXiv preprint arXiv:1810.01392*.
- Chua, K.; Calandra, R.; McAllister, R.; and Levine, S. 2018. Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models. In *Advances in Neural Information Processing Systems*, volume 31.
- Depeweg, S.; Hernández-Lobato, J. M.; Doshi-Velez, F.; and Udluft, S. 2016. Learning and policy search in stochastic dynamical systems with bayesian neural networks. *arXiv preprint arXiv:1605.07127*.
- Depeweg, S.; Hernandez-Lobato, J.-M.; Doshi-Velez, F.; and Udluft, S. 2018. Decomposition of uncertainty in Bayesian deep learning for efficient and risk-sensitive learning. In *International Conference on Machine Learning*, 1184–1193. PMLR.
- Der Kiureghian, A.; and Ditlevsen, O. 2009. Aleatory or epistemic? Does it matter? *Structural safety*, 31(2): 105–112.
- Durasov, N.; Bagautdinov, T.; Baque, P.; and Fua, P. 2021. Masksembles for uncertainty estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13539–13548.
- Durkan, C.; Bekasov, A.; Murray, I.; and Papamakarios, G. 2019a. Cubic-Spline Flows. In *Workshop on Invertible Neural Networks and Normalizing Flows, International Conference on Machine Learning*.
- Durkan, C.; Bekasov, A.; Murray, I.; and Papamakarios, G. 2019b. Neural Spline Flows. In *Advances in Neural Information Processing Systems*, volume 32.
- Durkan, C.; Bekasov, A.; Murray, I.; and Papamakarios, G. 2020. nflows: normalizing flows in PyTorch. <https://doi.org/10.5281/zenodo.4296287>. Accessed: 2021-09-01.
- Freund, Y.; and Schapire, R. E. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1): 119–139.
- 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.; Islam, R.; and Ghahramani, Z. 2017. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, 1183–1192. PMLR.
- Hora, S. C. 1996. Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management. *Reliability Engineering & System Safety*, 54(2-3): 217–223.
- Houlsby, N.; Huszár, F.; Ghahramani, Z.; and Lengyel, M. 2011. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*.
- Hüllermeier, E.; and Waegeman, W. 2021. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110(3): 457–506.
- Kendall, A.; and Gal, Y. 2017. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30.
- Kingma, D. P.; and Dhariwal, P. 2018. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31.
- Kirichenko, P.; Izmailov, P.; and Wilson, A. G. 2020. Why normalizing flows fail to detect out-of-distribution data. *Advances in neural information processing systems*, 33: 20578–20589.
- Kirsch, A.; Van Amersfoort, J.; and Gal, Y. 2019. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. *Advances in neural information processing systems*, 32.
- Kurutach, T.; Clavera, I.; Duan, Y.; Tamar, A.; and Abbeel, P. 2018. Model-ensemble trust-region policy optimization. *arXiv preprint arXiv:1802.10592*.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30.
- MacKay, D. J. 1992. Information-based objective functions for active data selection. *Neural computation*, 4(4): 590–604.
- Papamakarios, G.; Nalisnick, E. T.; Rezende, D. J.; Mohamed, S.; and Lakshminarayanan, B. 2021. Normalizing Flows for Probabilistic Modeling and Inference. *J. Mach. Learn. Res.*, 22(57): 1–64.
- Postels, J.; Blum, H.; Strümler, Y.; Cadena, C.; Siegwart, R.; Van Gool, L.; and Tombari, F. 2020. The hidden uncertainty in a neural networks activations. *arXiv preprint arXiv:2012.03082*.
- Rasmussen, C. E. 2003. Gaussian processes in machine learning. In *Summer school on machine learning*, 63–71. Springer.
- Rezende, D.; and Mohamed, S. 2015. Variational Inference with Normalizing Flows. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, 1530–1538. Lille, France.

- Rubinstein, R. Y.; and Glynn, P. W. 2009. How to deal with the curse of dimensionality of likelihood ratios in Monte Carlo simulation. *Stochastic Models*, 25(4): 547–568.
- Settles, B. 2009. Active learning literature survey. *Synthesis Lectures on Artificial Intelligence and Machine Learning*.
- Tabak, E. G.; and Turner, C. V. 2013. A family of nonparametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2): 145–164.
- Tabak, E. G.; and Vanden-Eijnden, E. 2010. Density estimation by dual ascent of the log-likelihood. *Communications in Mathematical Sciences*, 8(1): 217–233.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, 5026–5033. IEEE.
- Trep., V. 1994. The wet game of chicken. Technical report, Siemens AG.
- Wang, Q.; Kulkarni, S. R.; and Verdu, S. 2009. Divergence Estimation for Multidimensional Densities Via k -Nearest-Neighbor Distances. *IEEE Transactions on Information Theory*, 55(5): 2392–2405.
- Winkler, C.; Worrall, D.; Hoogeboom, E.; and Welling, M. 2019. Learning likelihoods with conditional normalizing flows. *arXiv preprint arXiv:1912.00042*.
- Zhang, L.; Goldstein, M.; and Ranganath, R. 2021. Understanding failures in out-of-distribution detection with deep generative models. In *International Conference on Machine Learning*, 12427–12436. PMLR.