
Deep Gaussian Mixture Ensembles

Yousef El-Laham¹

Niccolò Dalmaso¹

Elizabeth Fons¹

Svitlana Vyetenko¹

¹J.P. Morgan AI Research, New York, USA

Abstract

This work introduces a novel probabilistic deep learning technique called deep Gaussian mixture ensembles (DGMEs), which enables accurate quantification of both epistemic and aleatoric uncertainty. By assuming the data generating process follows that of a Gaussian mixture, DGMEs are capable of approximating complex probability distributions, such as heavy-tailed or multimodal distributions. Our contributions include the derivation of an expectation-maximization (EM) algorithm used for learning the model parameters, which results in an upper-bound on the log-likelihood of training data over that of standard deep ensembles. Additionally, the proposed EM training procedure allows for learning of mixture weights, which is not commonly done in ensembles. Our experimental results demonstrate that DGMEs outperform state-of-the-art uncertainty quantifying deep learning models in handling complex predictive densities.

1 INTRODUCTION

Uncertainty quantification plays a key role in the development and deployment of machine learning systems, especially in applications where user safety and risk assessment are of paramount importance [Abdar et al., 2021]. While deep learning (DL) has cemented its superiority in terms of raw predictive performance for a variety of applications, the principled incorporation of uncertainty quantification in DL models remains an open challenge. Since standard DL models are unable to properly quantify predictive uncertainty, one common challenge for deep learning models is detecting out-of-distribution (OOD) inputs. It is often the case that OOD inputs lead a DL model to making erroneous predictions [Ovadia et al., 2019]. Without uncertainty quantification, one cannot reason about whether an input is OOD

and this can be catastrophic in applications such as machine-assisted medical decision making [Begoli et al., 2019] or self-driving vehicles [Michelmore et al., 2018]. Moreover, uncertainty quantification can also be used as a means to select samples to label in active learning and for enabling exploration in reinforcement learning algorithms [Clements et al., 2019, Charpentier et al., 2022].

Uncertainty in machine learning models is derived from two different sources: aleatoric uncertainty and epistemic uncertainty [Kendall and Gal, 2017, Hüllermeier and Waegeman, 2019, Valdenegro-Toro and Mori, 2022]. Aleatoric uncertainty derives from measurement process of the data, while epistemic uncertainty derives from the uncertainty in the parameters of the machine learning model. A variety of approaches have been proposed to quantify both types of uncertainty in DL models from both a Bayesian and frequentist perspective; we refer the reader to Gawlikowski et al. [2021] for a comprehensive review. Under the Bayesian paradigm, the goal is to infer the posterior predictive density of the target variable given the input features and the training data, which encodes both types of uncertainty. Unfortunately, exact Bayesian inference algorithms (e.g., Neal 2012) cannot scale to the parameter space of modern DL architectures and one often has to resort to mini-batching [Chen et al., 2014] or forming a rough parametric approximation of the posterior distribution of the parameters, such as the Laplace approximation [Daxberger et al., 2021] or stochastic variational inference [Graves, 2011, Hoffman et al., 2013]. The drawback of the parametric approach is the inability to express more complex (e.g., heavy-tailed or multimodal) predictive distributions. As an example, approximations such as mean-field variational inference form a Gaussian predictive distribution that tends to underestimate the true uncertainty of more complex models.

In recent years, there have been developments in probabilistic DL which exploit the inherent stochasticity in learning to quantify predictive uncertainty. Examples include techniques such as probabilistic backpropagation [Hernández-Lobato and Adams, 2015], Monte Carlo dropout (MCD,

Gal and Ghahramani 2016), Monte Carlo batch normalization [Teye et al., 2018], deep ensembles (DEs, Lakshminarayanan et al. 2017) among others. MCD and DEs have emerged as state-of-the-art solutions for quantifying uncertainty in DL models due to their simplicity and effectiveness. MCD utilizes the inherent stochasticity of dropout (i.e., random masking of neural network weights) to form an ensemble-based approximation of the predictive distribution through multiple stochastic forward passes of the model to account for epistemic uncertainty. Aleatoric uncertainty is handled as a post-processing step under the assumption that the underlying data noise is homoscedastic. DEs, on the other hand, independently train a small ensemble of dual-output neural networks, where the outputs characterize the mean and variance of the predictive distribution. Each network in the ensemble is independently trained to maximize the likelihood of the data under the heteroscedastic Gaussian assumption. At test time, the networks are linearly combined into a single Gaussian approximation of the predictive distribution. Unfortunately, both MCD and DEs are not adequate solutions for modeling more complex data distributions (e.g., heavy-tailed or multimodal distributions).

Contributions. Our contribution are as follows:

1. We propose a novel probabilistic DL technique called *deep Gaussian mixture ensembles* (DGMEs) for jointly quantifying epistemic and aleatoric uncertainty. DGMEs train a weighted deep ensemble using the expectation maximization algorithm;
2. We show DGMEs optimizes the joint data likelihood directly, unlike deep ensembling that targets a lower bound of the data likelihood. As a consequence, DGMEs achieve a superior loss to deep ensembling, which is corroborated by our experimental results;
3. We empirically show that our model is more expressive than standard probabilistic DL approaches and can capture both heavy-tailedness and multimodality.

2 RELATED WORK

Mixture Density Networks. Mixture density networks (MDNs, Bishop 1994) use a deep neural network to simultaneously learn the means, variances and mixture weights of a Gaussian mixture model. MDNs have been successfully used in many machine learning applications, such as computer vision [Bazzani et al., 2016], speech synthesis [Zen and Senior, 2014], probabilistic forecasting [Zhang et al., 2020], astronomy [D’Isanto and Polsterer, 2018], chemistry [Unni et al., 2020] and epidemiology [Davis et al., 2020], among others. While MDNs are closely related to DGMEs in terms of uncertainty quantification, they are not an ensemble technique per se, as the epistemic and aleatoric uncertainty cannot be disentangled in MDNs. Moreover, without the ensemble structure, MDNs cannot easily be

trained in a distributed setting, whereas the training of DGMEs can trivially be parallelized.

Monte Carlo Dropout. MCD exploits the stochasticity of dropout training to quantify epistemic uncertainty in DL models. At test time, stochastic forward passes through a DL model with dropout produce “approximate” samples from the underlying posterior predictive distribution, which are typically summarized using first- and second-order moments (e.g, mean and variance of the samples). Aleatoric uncertainty is accounted for in a post-processing step, whereby the optimal homoscedastic variance that maximizes the evidence lower-bound is obtained via cross-validation. MCD’s popularity can be attributed to its simple implementation, as no changes to the standard DL training procedure are required. While vanilla MCD can yield favorable results in additive Gaussian settings, the method is less effective when dealing with more complex data generating processes (e.g., heavy-tailed or multimodal predictive densities). In this work, we incorporate MCD in our training procedure to account for epistemic uncertainty (see Section 4.3).

Deep Ensembles. DEs quantify both aleatoric and epistemic uncertainty by building an ensemble of independently trained models under different neural network weight initializations. Combined with adversarial training [Goodfellow et al., 2014], DEs achieve competitive or better performance than MCD in most settings in terms of calibration of predictive uncertainty and in terms of reasoning about OOD inputs. It has been argued in several works that DEs can be interpreted as a Bayesian approach, where the learned weights of each ensemble member correspond to a sample from the posterior distribution of the network weights [Wilson and Izmailov, 2020b]. In recent years, new variations of deep ensembles have been proposed, such as anchored DEs [Pearce et al., 2020], deep-split ensembles [Sarawgi et al., 2020], and hybrid training approaches that combine DEs with the Laplace approximation [Hoffmann and Elster, 2021]. We emphasize that a key distinction between DEs and DGMEs is that each sample in DEs is treated as an i.i.d. sample from a Gaussian distribution, whereas DGMEs assume that data are distributed according to a Gaussian mixture. This gives DGMEs the key advantage of being able to learn more complex data generating processes.

Neural Expectation Maximization. Neural EM is a differentiable clustering technique that combines the principles of the EM algorithm with neural networks for representation learning, particularly in the field of computer vision for perceptual grouping tasks [Greff et al., 2017]. The goal of neural EM is to group the individual entities (i.e., pixels) of a given input (i.e., image) that belong to the same object. To do this, a finite mixture model is used to construct a latent representation of each image, where each mixture component represents a distinct object. A neural network is then used to transform the parameters of that mixture model into pixel-wise distributions over the image, allowing for rea-

soning about which object each pixel in the image belongs to. While neural EM combines the ideas of EM with deep learning, we emphasize that this is different from our work which focuses on the accurate quantification of predictive uncertainty in the supervised learning setting.

3 PROBLEM FORMULATION

Consider a set of training data $\mathcal{D} = \{(x_n, y_n)\}_{n=1}^N$, where $x_n \in \mathbb{R}^{d_x}$ is the feature vector and y_n is the output, which can be real-valued if we are dealing with a regression task or integer-valued if we are dealing with a classification task. We would like to train a model that allows us to predict an output y given its corresponding input vector x .

From a probabilistic perspective, the goal is to determine the posterior predictive distribution $p(y|x, \mathcal{D})$. We assume a statistical model $p_\theta(y|x) \triangleq p(y|x, \theta)$ that relates each output to its corresponding feature vector through a set of parameters $\theta \in \Theta$. Then, the predictive distribution can be determined as

$$p(y|x, \mathcal{D}) = \int_{\Theta} p_\theta(y|x) p(\theta|\mathcal{D}) d\theta. \quad (1)$$

While this integral is generally intractable, it can be approximated using a Monte Carlo average, where samples are taken from the posterior $p(\theta|\mathcal{D})$. Let $Y = \{y_1, \dots, y_N\}$ and let $X = \{x_1, \dots, x_N\}$. According to Bayes theorem, the posterior distribution $p(\theta|\mathcal{D})$ is

$$p(\theta|\mathcal{D}) = \frac{p(Y|X, \theta)p(\theta)}{p(Y|X)}, \quad (2)$$

where $p(Y|X, \theta) = \prod_{n=1}^N p_\theta(y_n|x_n)$ is called the *data likelihood* under the i.i.d. assumption, $p(\theta)$ is the *prior distribution* of θ , and $p(Y|X) = \int_{\Theta} p(Y|X, \theta)p(\theta)d\theta$ is called the *marginal likelihood*. The posterior can only be computed analytically when $p(\theta)$ is a conjugate prior for the likelihood function $p(Y|X, \theta)$. For deep learning models, an analytical solution to the posterior cannot be determined and one must resort to an approximation of the predictive distribution.

Goal. The goal is to acquire an approximation of the posterior predictive distribution. Ideally, samples from the approximation form consistent estimators of key moments of the predictive distribution that allow one to (i) formulate predictions, (ii) identify the underlying stochastic risk associated with the prediction (e.g., aleatoric uncertainty), and (iii) reason about the model’s uncertainty in the presence of the OOD data (i.e., epistemic uncertainty).

4 DEEP GAUSSIAN MIXTURE ENSEMBLES

In this work, we propose DGMEs to effectively learn a mixture distribution that accurately represents the true con-

ditional density of the labels given the features. Since Gaussian mixtures are universal approximators for smooth probability density functions [Bacharoglou, 2010, Goodfellow et al., 2016], modeling the conditional density $p_\theta(y|x)$ as a Gaussian mixture allows for learning more complex distributions, such as skewed, heavy-tailed, and multimodal distributions. Under the assumption that our data follows a mixture distribution with K mixture components, the conditional density of a particular example (x, y) is given by:

$$p_\theta(y|x) = \sum_{k=1}^K \pi_k p_k(y|x, \theta_k), \quad (3)$$

where $\theta_k \in \Theta_k \subseteq \mathbb{R}^{d_\theta}$ denotes the underlying parameters of the k -th mixture and π_k denotes the weight of the k -th mixture and represents the probability that the example (x, y) is distributed according to $p_k(y|x, \theta_k)$. Throughout the rest of the text, we refer to all unknown parameters in the mixture as $\theta = \{\pi_1, \theta_1, \dots, \pi_K, \theta_K\}$. Hereafter, we consider the problem of learning the parameters of the mixture in (3) in the context of regression. We discuss a possible extension to classification in the Supplementary Material, Section C.

To effectively model this mixture, we make the following assumptions:

Assumption 4.1. The mixture weights $(\pi_1, \dots, \pi_K) \in \mathcal{S}_K$ do not depend on the input features, where \mathcal{S}_K denotes the K -dimensional probability simplex.

Assumption 4.2. The conditional density $p_k(y|x, \theta_k)$ is a Gaussian distribution whose parameters are modeled via parameterized functions (neural networks) dependent on x :

$$p_k(y|x, \theta_k) = \mathcal{N}(y; \mu_{\theta_k}(x), \sigma_{\theta_k}^2(x)), \quad (4)$$

where θ_k denote the parameters of functions $\mu_{\theta_k}(\cdot)$ and $\sigma_{\theta_k}^2(\cdot)$ that output the mean and variance of the k -th mixture, respectively. Importantly, these functions are assumed to share parameters, just as in the original work on DEs [Lakshminarayanan et al., 2017].

Under the above assumptions, learning the mixture representation of $p_\theta(y|x)$ is equivalent to learning the parameters θ to maximize the data likelihood of the training examples $\mathcal{D} = \{(x_n, y_n)\}_{n=1}^N$.

4.1 LEARNING THE MIXTURE PARAMETERS

We obtain the maximum likelihood (ML) estimate or maximum a posteriori (MAP) estimate of the unknown parameters θ using the EM algorithm. Let $Y = \{y_1, \dots, y_N\}$ and $X = \{x_1, \dots, x_N\}$. Furthermore, let $Z = \{z_1, \dots, z_N\}$, where each $z_n \in \{1, \dots, K\}$ is a latent variable that denotes membership assignment of the training example (x_n, y_n) to a particular mixture component, where $\pi_k \triangleq P_\theta(z_n = k)$ is the probability that the example (x_n, y_n) belongs to the

k -th component. Assuming that the training examples are independent and identically distributed, we can write the joint likelihood as

$$p_\theta(Y, Z|X) = \prod_{n=1}^N \prod_{k=1}^K \pi_k^{\mathbb{I}(z_n=k)} \mathcal{N}(y_n; \mu_{\theta_k}(x_n), \sigma_{\theta_k}^2(x_n))^{\mathbb{I}(z_n=k)},$$

with corresponding log-likelihood of

$$\log p_\theta(Y, Z|X) = \sum_{n=1}^N \sum_{k=1}^K \mathbb{I}(z_n = k) (\log \pi_k + \ell_{\theta_k}(x_n, y_n))$$

where $\ell_{\theta_k}(x, y) = \log(\mathcal{N}(y; \mu_{\theta_k}(x), \sigma_{\theta_k}^2(x)))$. Our goal is to solve the following optimization problem:

$$\theta^* = \arg \max_{\theta} \log p_\theta(Y|X) \quad (5)$$

$$= \arg \max_{\theta} \log (\mathbb{E}_{Z|X, Y, \theta} [p_\theta(Y, Z|X)]), \quad (6)$$

which we numerically solve using the EM algorithm. In the following, we describe both the expectation step (E-Step) and maximization step (M-Step) as it relates to our model. As a note, all results presented hereafter also apply to the more general problem of obtaining the MAP estimate of the parameters θ .¹

E-Step: We update the posterior probabilities of each z_n given the parameters θ and the example (x_n, y_n) for each n , denoted by $\gamma_{k,n} \triangleq P_\theta(z_n = k|x_n, y_n)$. This can be done directly using Bayes' theorem:

$$\gamma_{k,n} = \frac{p_k(y_n|x_n, \theta_k) P_\theta(z_n = k)}{\sum_{j=1}^K p_j(y_n|x_n, \theta_j) P_\theta(z_n = j)} \quad (7)$$

$$= \frac{\pi_k \mathcal{N}(y_n; \mu_{\theta_k}(x_n), \sigma_{\theta_k}^2(x_n))}{\sum_{j=1}^K \pi_j \mathcal{N}(y_n; \mu_{\theta_j}(x_n), \sigma_{\theta_j}^2(x_n))} \quad (8)$$

M-Step: The parameters θ are updated in the maximization step by maximizing the expected joint log-likelihood $Q(\theta, \theta') \triangleq \mathbb{E}_{Z|X, Y, \theta'} [\log p_\theta(Y, Z|X)]$ given the previous parameter values θ' , which is equivalent to doing lower-bound maximization on the true log-likelihood [Minka, 1998]. The function $Q(\theta, \theta')$ can be readily determined as:

$$Q(\theta, \theta') = \sum_{n=1}^N \sum_{k=1}^K \gamma_{k,n} (\log(\pi_k) + \ell_{\theta_k}(x_n, y_n)). \quad (9)$$

The optimization of the mixture weights (π_1, \dots, π_K) can be carried out analytically and done independently of optimizing the mixture parameters $\{\theta_1, \dots, \theta_K\}$:

$$(\pi_1^*, \dots, \pi_K^*) = \arg \max_{(\pi_1, \dots, \pi_K) \in \mathcal{S}_K} Q(\theta, \theta'), \quad (10)$$

¹That is, the maximizer of $\log p(Y, Z|X) = \log p_\theta(Y|X) + \log p(\theta)$, where $p(\theta)$ is the prior distribution of the mixture parameters.

Algorithm 1 Deep Gaussian Mixture Ensembles (DGMEs)

1: Inputs:

- Training dataset $\mathcal{D} = \{(x_n, y_n)\}_{n=1}^N$
- Number of mixture components K
- Number of EM steps J

2: Initialize mixture parameters:

- Sample $\theta_k^{(0)} \sim p(\theta)$ for all k .
- Set $\pi_k^{(0)} = \frac{1}{K}$ for all k .

3: for $j = 1, \dots, J$ do

- 4: E-Step:** Update posterior probabilities $\gamma_{k,n}^{(j)}$ according to (8) with mixture weights $\pi_k^{(j-1)}$ and mixture parameters $\theta_k^{(j-1)}$ for all k and n .
- 5: M-Step:** Update mixture weights $\pi_k^{(j)}$ and parameters $\theta_k^{(j)}$ for all k as

$$\pi_k^{(j)} = \frac{1}{N} \sum_{n=1}^N \gamma_{k,n}^{(j)}$$

and

$$\theta_k^{(j)} = \arg \max_{\theta_k \in \Theta_k} \sum_{n=1}^N \gamma_{k,n}^{(j)} \ell_{\theta_k}(x_n, y_n)$$

6: end for

- 7: Return:** $\pi_k^* = \pi_k^{(J)}$ and $\theta_k^* = \theta_k^{(J)}$ for all k .
-

where for each k

$$\pi_k^* = \frac{1}{N} \sum_{n=1}^N \gamma_{k,n}. \quad (11)$$

Since the mixture parameters are assumed to be parameterised by neural networks, their optimization must be carried out using stochastic optimization. It is easy to see that the optimization of each θ_k can be done independently as:

$$\theta_k^* = \arg \max_{\theta_k \in \Theta_k} \sum_{n=1}^N \gamma_{k,n} \ell_{\theta_k}(x_n, y_n) \quad (12)$$

$$= \arg \min_{\theta_k \in \Theta_k} \sum_{n=1}^N \gamma_{k,n} \left(\log \sigma_{\theta_k}^2(x_n) + \frac{(y_n - \mu_{\theta_k}(x_n))^2}{\sigma_{\theta_k}^2(x_n)} \right) \quad (13)$$

This optimization step can be thought as training a deep ensemble, where each sample (x_n, y_n) is weighted by $\gamma_{k,n}$ in its negative log-likelihood contribution.

4.2 IMPLEMENTATION DETAILS

Our implementation of DGMEs trained via the EM algorithm is summarized in Algorithm 1. To initialize the ensemble, the parameters of each network in the ensemble are

randomly initialized, while the mixture weights are assumed to be equal. The algorithm is run for J steps or alternatively until some stopping criterion is met. The E-Step for updating the posterior probabilities is computed directly for each sample in the training set. In the M-Step, the updates for the mixture weights are also carried out analytically, but for mixture component parameters θ_k we use the stochastic optimization to numerically solve for the updates, as an analytical solution is not available. At round j , we initialize each network to $\theta_k^{(j-1)}$ and then run the Adam optimizer for E epochs to minimize the weighted negative log-likelihood in (13), where the weights are given by $\gamma_{k,n}^{(j)}$ for all n . Finally, we note that the computational complexity of each EM step is equivalent to that of DEs and the overall time complexity scales linearly with the number of EM steps.

4.3 QUANTIFYING EPISTEMIC UNCERTAINTY

It is important to highlight that up until this point, we have not explicitly considered epistemic uncertainty in DGMEs. This is because the operation of training DGMEs according to Algorithm 1 yields a single set of parameters of the assumed Gaussian mixture model.² To account for model uncertainty, we need to account for the uncertainty in the parameters of the mixture (i.e., the mixture weights and/or the weights of the ensemble neural networks). One simple way to do this is to apply MCD to the training procedure of DGMEs — although we emphasize other techniques can be applied to account for epistemic uncertainty (e.g., Laplace approximation or a variational approximation to the posterior parameters).

Let $a_k = [a_{k,1}, \dots, a_{k,d_\theta}]^\top \in \{0, 1\}^{d_\theta}$ denote a random binary vector of the same size as each θ_k and let $p_d \in [0, 1]$ denote a fixed dropout probability. Also, let $\theta^* = \{\pi_1^*, \theta_1^*, \dots, \pi_K^*, \theta_K^*\}$ denote the parameters learned by running Algorithm 1 with dropout incorporated in the training in the M-Step. For a given mixture component k , samples from the approximate posterior distribution of θ_k learned via dropout can be obtained via the following procedure:

$$a_{k,i} \sim \text{Bernoulli}(p_d), \quad i = 1, \dots, d_\theta,$$

$$\theta_k = a_k \odot \theta_k^*,$$

where \odot denotes a Hadamard (or element-wise) product. It follows that a sample from the predictive distribution can

²This point highlights the intrinsic difference in training DEs versus training DGMEs. DGMEs do not have a ‘‘Bayesian’’ interpretation, because the EM algorithm used to train them only outputs a single set of possible parameters for the DGMEs (i.e., the corresponding posterior distribution of the weights is a Dirac measure centered at the learned parameter values).

directly be obtained as follows:

$$k \sim \text{Categorical}(\pi_1, \dots, \pi_K), \quad (14)$$

$$a_{k,i} \sim \text{Bernoulli}(p_d), \quad i = 1, \dots, d_\theta, \quad (15)$$

$$\theta_k = a_k \odot \theta_k^*, \quad (16)$$

$$y \sim p_k(y|x, \theta_k). \quad (17)$$

In this procedure, one first samples the mixing component k via (14). Then, one draws a sample from the approximate posterior distribution of the parameters of the k -th mixture via (15)–(16). Finally, a prediction can be sampled via (17). We refer the reader to the Supplementary Material, Section E for details on the validity of this sampling procedure.

4.4 THEORETICAL INSIGHTS

In this section we provide insights into the connection between DGMEs and DEs, along with general results on convergence of our training procedure using DGMEs. We refer the reader to the Supplementary Material, Section A for details on the proofs of each propositions.

Proposition 4.3 shows that maximizing the data likelihood directly as in DGMEs achieves an equal or better likelihood than maximizing each ensemble member’s likelihood separately as in DEs.

Proposition 4.3. *Under the assumption that $\pi_i = 1/K$ for $i = 1, \dots, K - 1$, maximizing the Gaussian mixture data likelihood directly achieves better or equal joint likelihood than maximizing each ensemble member’s likelihood separately.*

Proof Sketch. The result can be obtained by using Jensen’s inequality on the joint log-likelihood of equation (5) along with the assumption.

Next, Proposition 4.8 combines recent results on neural network convergence in regression by Arora et al. [2019] and Farrell et al. [2021] with classical EM analysis [Wu, 1983] to give intuition on why DGMEs should converge towards the maximum of the data likelihood³.

Assumption 4.4 (Non-flatness of the weighted log-likelihood). Given a DGMEs with K mixtures, in each EM round t there exists an $\epsilon_{t,k}$ such that:

$$\sum_{n=1}^N \gamma_{k,n} (\ell_{\theta^*}(x_n, y_n) - \ell_{\theta^{(t)}}(x_n, y_n)) \geq \frac{\epsilon_{t,k}}{K},$$

³We note Proposition 4.8 covers a specific setup, in which mean and variance function estimation is performed separately by using a shared pre-trained feature extraction layer and that the true data generating process is identifiable with a mixture model to begin with. A more thorough investigation on both using a separate neural network from mean and variance, as well as the under- or over-specified case, is outside of the scope of this paper.

where $\theta_k^* = \arg \max_{\theta \in \Theta} \sum_{n=1}^N \gamma_{k,n} \ell_{\theta}(x_n, y_n)$. Let $\epsilon = \min_{t \in T, k \in K} \epsilon_{t,k}$.

Assumption 4.5 (Smoothness of the true mean function). Let $\mu(x) : \mathcal{X} \rightarrow \mathbb{R}$ be the true mean function and let $X \subset \mathcal{X}$. Assume there exists some $\beta \in \mathbb{N}^+$ such that $\mu(x) \in \mathcal{W}^{\beta, \infty}(X)$, where $\mathcal{W}^{\beta, \infty}(X)$ is a (β, ∞) -Sobolev ball.

Assumption 4.6 (Smoothness of the true variance function). Let $\sigma(x) : \mathcal{X} \rightarrow \mathbb{R}^+$ be the true variance function and let $X \subset \mathcal{X}$. Let \mathbf{H}^{∞} be the Graham matrix as defined by Arora et al. [2019, Equation 12], and assume that there exists an $M \in \mathbb{R}$ such that $\sigma(x)^T (\mathbf{H}^{\infty}) \sigma(x)^T \leq M$ for some $M \in \mathbb{R}$.

Assumption 4.7 (Non-degenerate weights). In each EM iteration, the weights are positive and bounded away from zero, e.g., $\pi_i^{(t)} > \xi_i^{(t)} > 0$.

Proposition 4.8. *Under assumptions 4.4, 4.5, 4.6 and 4.7, let the mean and variance in each ensemble model be estimated via a separate 2-layer deep ReLU network from a common feature extraction layer. Then the DGMEs EM algorithm converges to a non-stationary point that maximizes the data likelihood with high probability.*

Proof Sketch. The result follows if one shows that $Q(\theta; \theta^{(j)})$ is an increasing function of the EM steps j (Wu 1983), for parameter values $\theta^{(j)}$ that are not stationary points of $Q(\theta; \theta^{(j)})$. In the DGMEs case, this corresponds to proving that the weighted log-likelihood in each ensemble increases at every round j . The result follows by combining assumptions on the non-flatness of the weighted log-likelihood (A.3), the smoothness of true mean function (A.4) and the smoothness of the true variance function (A.5) with the results obtained about convergence of deep ReLU networks by Farrell et al. [2021] and Arora et al. [2019] respectively.

Finally, Proposition 4.9 connects DGMEs and DEs, showing DEs is equivalent to a single-EM-step of DGMEs under specific neural network weights initialization. As shown in Proposition 4.8, the EM training of DGME improves the function Q at each iteration t , i.e., $Q(\theta^{(t+1)}, \theta^{(t)}) \geq Q(\theta^{(t)}, \theta^{(t)})$. Hence, the final joint DGME likelihood will be larger or equal to the joint likelihood achieved by DE.

Proposition 4.9. *If the weights of each ensemble member are initialized to 0 with fixed bias terms, a single EM step for DGMEs is equivalent to perform DEs.*

Proof Sketch. The initialization schema implies that mixture membership is equal across samples in the first expectation round of the EM. Hence, the first M-step consists in training K separate networks with each log-likelihood contribution being weighted equally.

5 EXPERIMENTS

We evaluate the empirical performance of DGMEs via three different numerical experiments. We compare our method to the MDNs [Bishop, 1994], MCD [Gal and Ghahramani, 2016], and DEs [Lakshminarayanan et al., 2017]. MCD and DEs are widely considered to be state-of-the-art solutions for quantifying predictive uncertainty in deep learning models and have repeatedly been used as baselines for developing new techniques. Additional results and figures can be found in the Supplementary Material, Section B. A summary table qualitatively comparing DGMEs to the benchmarks can be found in the Supplementary Material, Section D.

5.1 TOY REGRESSION

Consider the following model:

$$y_n = u_n x_n^3 + \epsilon_n, \quad (18)$$

where $u_n \in \{-1, 1\}$ with $p_u \triangleq P(u_n = -1)$ and $\epsilon_n \sim p(\epsilon)$ for all $n = 1, \dots, N$. We generate $N = 800$ training samples from this model for the training set, where the input values x_n range from -4 to 4. For each considered setting, we use a learning rate of $\eta = 0.01$, a batch size of 32, and $E = 80$ epochs to resolve the stochastic optimization problem in the M-step. For each method, we utilize a dropout probability $p_d = 0.1$ to account for epistemic uncertainty. Additionally, we generate data from this toy model under three different noise settings to demonstrate the flexibility and expressive power of DGMEs as compared to other baselines. Unless otherwise stated, we assume $K = 5$ networks in each mixture model-based approaches (i.e., MDNs, DEs, and DGMEs). Experimental results are described below for each noise scenario. Additional experimental results and ablation studies are provided in the Supplementary Material, Section B.1.

Case 1 - Gaussian Noise: We set $p_u = 0$ and assume that the noise is zero-mean and Gaussian distributed with variance of 9. This is analogous to the setup utilized in Hernández-Lobato and Adams [2015]. Figure 3 (Supplementary Material, Section B.1.5) shows the performance of DGMEs as compared to the baselines, where we observe that it outperforms MDNs and obtains comparable results to MCD and DEs.

Case 2 - Heavy-tailed Noise: We set $p_u = 0$ and assume that the noise distributed according to a zero-mean Student-t distribution with $\nu = 3$ degrees of freedom with variance of 9. Figure 1 shows the histogram of samples from the predictive distribution of both a training and a test input with their corresponding sample (excess) kurtosis. We observe that on the training examples (i.e., purple histograms), only MDNs and DGMEs are able to learn the heavy-tailedness of the noise, as both MCD and DEs obtain a kurtosis close




Figure 1: Histogram of samples from the predictive distributions for a single training example (top panel) and for a single test example (bottom panel) from the heavy-tailed toy regression example, shown with corresponding sample kurtosis value κ . DGMEs generally estimate heavier tailed predictions for both training and test samples, while baseline approaches samples are closer to following a Gaussian distribution.




Figure 2: Predictive distribution plots for the bimodal Gaussian toy regression example. DEs cannot capture the multimodality of the noise, while MDNs and DGMEs can. Furthermore, DGMEs approximate the mixture weights of the noise accurately (ground truth: $\pi_1 = 0.7$ and $\pi_2 = 0.3$).

to 0. Unlike the baseline approaches, which are unable to learn the tail behavior in the test example, we observe that DGMEs is the best method at capturing the heavy-tailedness of the test examples, as it gives the largest corresponding kurtosis.

Case 3 - Bimodal Gaussian Noise: We set $p_u = 0.3$ and assume that the noise is zero-mean and Gaussian distributed with variance of 9. For this example, we only compare the mixture-based approaches assuming $K = 2$ components. Figure 2 shows the predictive density for the corresponding 99% credible interval for each mixture in each approach, where for DGMEs we also show the learned mixture weights of each component. We observe that only MDNs and DGMEs are able to capture the bimodality of the data, with DGMEs also accurately capturing the mixture weight

proportions. DEs instead overestimates the heteroscedastic variance in each network. This is due to the fact that DEs train each ensemble member independently under the assumption of Gaussian likelihood. We also show that DGMEs can robustly estimate this bimodality, even if the assumed number of mixture components is larger than 2 (see Supplementary Material, Section B.1.3).

5.2 REGRESSION ON REAL DATASETS

We evaluate the performance of DGMEs in regression against MDNs, MCD and DEs on a set of UCI regression benchmark datasets [Dua and Graff, 2017]; see Supplementary Material, Section B.2, for further details on the datasets. We use the experimental setup used in Hernández-Lobato and Adams [2015], with each dataset split into 20 train-test

TEST RMSE							
Dataset	MDNs	MCD	DEs	DGMEs (J=1)	DGMEs (J=2)	DGMEs (J=5)	DGMEs (J=10)
Boston housing	2.79 ± 0.84	2.97 ± 0.85	3.28 ± 1.00	3.11 ± 0.94	3.00 ± 0.90	2.87 ± 0.86	2.83 ± 0.91
Concrete	5.21 ± 0.56	5.23 ± 0.53	6.03 ± 0.58	5.67 ± 0.57	5.36 ± 0.51	5.20 ± 0.59	5.14 ± 0.58
Energy	0.71 ± 0.14	1.66 ± 0.19	2.09 ± 0.29	2.01 ± 0.29	1.79 ± 0.24	1.22 ± 0.25	1.07 ± 0.41
Kin8nm	0.08 ± 0.00	0.10 ± 0.00	0.09 ± 0.00	0.08 ± 0.00	0.08 ± 0.00	0.07 ± 0.00	0.07 ± 0.00
Power plant	4.12 ± 0.17	4.02 ± 0.18	4.11 ± 0.17	4.12 ± 0.16	4.10 ± 0.15	4.07 ± 0.15	4.05 ± 0.13
Wine	0.66 ± 0.04	0.62 ± 0.04	0.64 ± 0.04	0.63 ± 0.04	0.64 ± 0.04	0.64 ± 0.04	0.66 ± 0.05
Yacht	0.96 ± 0.36	1.11 ± 0.38	1.58 ± 0.48	0.98 ± 0.38	0.85 ± 0.36	0.83 ± 0.40	0.70 ± 0.26

Table 1: Average RMSE of test examples for regression experiments on real datasets. DGMEs obtain competitive or better performance in terms of RMSE on the majority of datasets as compared to the baselines.

TEST NLL							
Dataset	MDNs	MCD	DEs	DGMEs (J=1)	DGMEs (J=2)	DGMEs (J=5)	DGMEs (J=10)
Boston housing	2.62 ± 0.43	2.46 ± 0.25	2.41 ± 0.25	2.34 ± 0.19	2.33 ± 0.22	2.41 ± 0.25	2.46 ± 0.31
Concrete	3.11 ± 0.26	3.04 ± 0.09	3.06 ± 0.18	3.04 ± 0.11	3.00 ± 0.12	2.95 ± 0.13	2.94 ± 0.14
Energy	1.18 ± 0.30	1.99 ± 0.09	1.38 ± 0.22	1.71 ± 0.19	1.48 ± 0.15	1.20 ± 0.23	1.20 ± 0.40
Kin8nm	-1.18 ± 0.04	-0.95 ± 0.03	-1.20 ± 0.02	-1.20 ± 0.02	-1.23 ± 0.03	-1.24 ± 0.02	-1.25 ± 0.02
Power plant	2.81 ± 0.04	2.80 ± 0.05	2.79 ± 0.04	2.82 ± 0.03	2.81 ± 0.03	2.81 ± 0.03	2.79 ± 0.03
Wine	1.01 ± 0.10	0.93 ± 0.06	0.94 ± 0.12	0.95 ± 0.11	0.96 ± 0.11	0.96 ± 0.12	1.10 ± 0.09
Yacht	1.18 ± 0.17	1.55 ± 0.12	1.18 ± 0.21	1.07 ± 0.22	0.75 ± 0.22	0.60 ± 0.29	0.49 ± 0.29

Table 2: Average NLL of test examples for regression experiments on real datasets. DGMEs obtain competitive or better performance in terms of NLL on the majority of datasets as compared to the baselines.

olds. We use the same network architecture across each dataset: an MLP with a single hidden layer and ReLU activations, containing 50 hidden units. For each dataset we train for $E = 40$ total epochs with a batch size of 32 and a learning rate of $\eta = 0.001$. To be consistent with previous evaluations, we used $K = 5$ networks in our ensemble and provide results for DGMEs for different numbers of EM steps $J \in \{1, 2, 5, 10\}$. Our results are shown in Tables 1 and 2, where we evaluate the root-mean-squared error (RMSE) and the negative log-likelihood (NLL) on the test set averaged over the different folds, respectively. In the same table, we also report the results for MDNs, MCD and DEs. Experimental results for MCD and DEs can also be found in their respective papers [Gal and Ghahramani, 2016, Lakshminarayanan et al., 2017]. Note that in this experiment, we do not apply dropout to MDNs and DGMEs and only account for the uncertainty obtained from training the models to maximize the NLL of the samples according to the Gaussian mixture assumption in (3).

We observe that in this experiment DGMEs are able to obtain competitive (or better) performance with respect to the baseline methods. For certain datasets, we observe that increasing the number of EM steps greatly improves the performance (e.g., Concrete, Energy, Power Plant, and Yacht). We can see that this is not generally true for all datasets: for example, for the Boston housing dataset, increasing the number of EM steps begins to degrade the performance of the model in terms of NLL. We emphasize that performance

can further be improved by incorporating dropout in the training procedure, where the dropout probability p_d can be selected using cross-validation on each train-test split.

5.3 FINANCIAL TIME SERIES FORECASTING

For the final experiment, we focus on the task of one-step-ahead forecasting for financial time series. In particular, using historical daily price data from Yahoo finance⁴, we formulate a one-step ahead forecasting problem using a long short-term memory (LSTM) network [Hochreiter and Schmidhuber, 1997]. The input to the network is a time series that represents the closing price of a particular stock over the past 30 trading days. The target output is the next trading day’s closing price. We assess performance of the model using two metrics: (1) the NLL of the test set, and (2) the RMSE score on the test set. We evaluate each method on three different datasets:

- **GOOG - stable market regime:** We use training data from the Google (GOOG) stock from the period of Jan 2019 - July 2022 and test on GOOG stock data from the period of August 2022 - January 2023.
- **RCL - market shock regime:** We use training data from the Royal Caribbean (RCL) stock from the period of Jan 2019 - April 2020 and test on RCL stock data from the period of May 2020 - September 2020.

⁴<https://finance.yahoo.com/>

Test RMSE					
Dataset	MDNs	MCD	DEs	MultiSWAG	DGMEs
GOOG	2.74 ± 0.06	3.86 ± 0.16	2.73 ± 0.03	2.71 ± 0.05	2.71 ± 0.04
RCL	15.01 ± 4.71	16.19 ± 10.18	14.92 ± 1.44	11.73 ± 0.45	14.49 ± 2.73
GME	11.14 ± 7.75	2.70 ± 0.47	3.21 ± 0.46	2.00 ± 0.06	3.19 ± 0.33

Table 3: Average RMSE of the test examples for the financial forecasting experiment.

Test NLL					
Dataset	MDNs	MCD	DEs	MultiSWAG	DGMEs
GOOG	2.46 ± 0.03	2.98 ± 0.01	2.44 ± 0.01	2.54 ± 0.00	2.43 ± 0.02
RCL	18.83 ± 17.82	6.12 ± 3.93	5.94 ± 0.80	6.21 ± 0.18	5.00 ± 0.76
GME	6.01 ± 3.85	2.46 ± 0.16	2.66 ± 0.13	2.14 ± 0.03	2.61 ± 0.31

Table 4: Average NLL of the test examples for the financial forecasting experiment.

- **GME - high volatility regime:** We use the training data from the Gamestop (GME) stock during the “bubble” period of Nov 2020 - Jan 2022 and test on GME stock data following that period.

We ran each of the previously tested baselines and DGMEs on the three scenarios previously described. Additionally, we also test the MultiSWAG approach highlighted in Wilson and Izmailov [2020a], due to its effectiveness in quantifying epistemic uncertainty, which is of particular importance for the market shock regime [Chandra and He, 2021]. We train each model on each dataset for 5 independent runs and report the mean and standard error of both the test NLL and the test RMSE in Tables 3 and 4, where we have bolded the best performing method in each experiment according to the mean value of the metric. For details on selection of hyperparameters of each of the methods, we refer the reader to the Supplementary Material, Section B.3.

The results indicate that for the GOOG dataset, DGMEs achieve, on average, the best NLL and RMSE score. In the case of the RCL dataset, we observe an interesting result. DGMEs attain the best performance in terms of NLL, but MultiSWAG does best in terms of RMSE. We believe DGMEs outperform in terms of NLL because the likelihood function assumed by DGMEs is a true Gaussian mixture, while the MultiSWAG approach is applying stochastic weight averaging Gaussian (SWAG) independently on multiple networks under the Gaussian likelihood assumption. This gives DGMEs the advantage in terms of learning the complex nature of the RCL dataset. On the other hand, we have found that since MultiSWAG is accounting for uncertainty using SWAG, it appears to make model training more stable (hence the smaller standard error on each of the metrics) and better accounts for epistemic uncertainty. This could possibly explain why the RMSE score is lower than that of DGMEs and with smaller standard error. For the GME dataset, MultiSWAG outperforms DGMEs consistently, and with tighter standard error bars. As a final

remark, we emphasize that in our paper, we have accounted for epistemic uncertainty in DGMEs using dropout, but other methods could have been used (such as variational inference, Laplace approximation, or SWAG). Based on the results of this experiment, we highlight the possibility of incorporating SWAG in the training of DGMEs as a better way to account for epistemic uncertainty (as opposed to dropout).

6 CONCLUSIONS

This paper proposes DGMEs, a novel probabilistic DL ensemble method for jointly quantifying epistemic and aleatoric uncertainty. Unlike deep ensembling, DGMEs optimizes the data likelihood directly and is able to capture complex behavior in the predictive distribution (e.g., heavy-tailedness and multimodality) by modeling the conditional distribution of the data as a Gaussian mixture. Our experiments show that DGMEs can capture more complex distributional properties than a variety of probabilistic DL baselines in regression settings and obtain competitive performance on detecting OOD samples in classification settings. As next steps, alternative mechanisms for handling the epistemic uncertainty can be considered. For example, one can instead form a variational approximation to the posterior of each mixture component, thereby forming a Gaussian mixture approximation to the posterior parameters of the ensemble. Additionally, a more thorough analysis on the classification setting can be considered. Rather than using a mixture of categorical distributions to model the predictive density, one can use a mixture of Dirichlet distributions to account for uncertainty in the class probabilities, similar in line to the work of Hobbhahn et al. [2022]. Finally, DGMEs can be applied to improve the efficiency of active learning algorithms and exploration strategies in reinforcement learning.

Acknowledgments. This paper was prepared for informational purposes by the Artificial Intelligence Research group of JPMorgan Chase & Co. and its affiliates (“JP Morgan”), and is not a product of the Research Department of JP Morgan. JP Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful.

References

- Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezazadegan, Li Liu, Mohammad Ghavamzadeh, Paul Fieguth, Xiaochun Cao, Abbas Khosravi, U Rajendra Acharya, et al. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion*, 76:243–297, 2021.
- Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332. PMLR, 2019.
- Athanassia Bacharoglou. Approximation of probability distributions by convex mixtures of gaussian measures. *Proceedings of the American Mathematical Society*, 138(7):2619–2628, 2010.
- Loris Bazzani, Hugo Larochelle, and Lorenzo Torresani. Recurrent mixture density network for spatiotemporal visual attention. *arXiv preprint arXiv:1603.08199*, 2016.
- Edmon Begoli, Tanmoy Bhattacharya, and Dimitri Kusnezov. The need for uncertainty quantification in machine-assisted medical decision making. *Nature Machine Intelligence*, 1(1):20–23, 2019.
- Christopher M Bishop. Mixture density networks. 1994.
- Rohitash Chandra and Yixuan He. Bayesian neural networks for stock price forecasting before and during covid-19 pandemic. *Plos one*, 16(7):e0253217, 2021.
- Bertrand Charpentier, Ransalu Senanayake, Mykel Kochenderfer, and Stephan Günnemann. Disentangling epistemic and aleatoric uncertainty in reinforcement learning. *arXiv preprint arXiv:2206.01558*, 2022.
- Tianqi Chen, Emily Fox, and Carlos Guestrin. Stochastic gradient hamiltonian monte carlo. In *International conference on machine learning*, pages 1683–1691. PMLR, 2014.
- William R Clements, Bastien Van Delft, Benoît-Marie Robaglia, Reda Bahi Slaoui, and Sébastien Toth. Estimating risk and uncertainty in deep reinforcement learning. *arXiv preprint arXiv:1905.09638*, 2019.
- Christopher N Davis, T Deirdre Hollingsworth, Quentin Caudron, and Michael A Irvine. The use of mixture density networks in the emulation of complex epidemiological individual-based models. *PLoS computational biology*, 16(3):e1006869, 2020.
- Erik Daxberger, Agustinus Kristiadi, Alexander Immer, Runa Eschenhagen, Matthias Bauer, and Philipp Hennig. Laplace redux-effortless bayesian deep learning. *Advances in Neural Information Processing Systems*, 34:20089–20103, 2021.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Antonio D’Isanto and Kai Lars Polsterer. Photometric redshift estimation via deep learning-generalized and pre-classification-less, image based, fully probabilistic redshifts. *Astronomy & Astrophysics*, 609:A111, 2018.
- Max H Farrell, Tengyuan Liang, and Sanjog Misra. Deep neural networks for estimation and inference. *Econometrica*, 89(1):181–213, 2021.
- Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016.
- Jakob Gawlikowski, Cedrique Rovile Njieutcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. A survey of uncertainty in deep neural networks. *arXiv preprint arXiv:2107.03342*, 2021.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Alex Graves. Practical variational inference for neural networks. *Advances in neural information processing systems*, 24, 2011.
- Klaus Greff, Sjoerd van Steenkiste, and Jürgen Schmidhuber. Neural expectation maximization. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- José Miguel Hernández-Lobato and Ryan Adams. Probabilistic backpropagation for scalable learning of bayesian neural networks. In *International conference on machine learning*, pages 1861–1869. PMLR, 2015.
- Marius Hobbhahn, Agustinus Kristiadi, and Philipp Hennig. Fast predictive uncertainty for classification with bayesian deep networks. In *Uncertainty in Artificial Intelligence*, pages 822–832. PMLR, 2022.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9:1735–80, 12 1997.
- Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 2013.

- Lara Hoffmann and Clemens Elster. Deep ensembles from a bayesian perspective. *arXiv preprint arXiv:2105.13283*, 2021.
- Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning: A tutorial introduction. 2019.
- Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30, 2017.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- Rhiannon Michelmoré, Marta Kwiatkowska, and Yarin Gal. Evaluating uncertainty quantification in end-to-end autonomous driving control. *arXiv preprint arXiv:1811.06817*, 2018.
- Thomas Minka. Expectation-maximization as lower bound maximization. *Tutorial published on the web at <http://www-white.media.mit.edu/tpminka/papers/em.html>*, 7(2), 1998.
- Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, 32, 2019.
- Tim Pearce, Felix Leibfried, and Alexandra Brintrup. Uncertainty in neural networks: Approximately bayesian ensembling. In *International conference on artificial intelligence and statistics*, pages 234–244. PMLR, 2020.
- Utkarsh Sarawgi, Wazeer Zufikar, Rishab Khincha, and Patie Maes. Why have a unified predictive uncertainty? disentangling it using deep split ensembles. *arXiv preprint arXiv:2009.12406*, 2020.
- Mattias Teye, Hossein Azizpour, and Kevin Smith. Bayesian uncertainty estimation for batch normalized deep networks. In *International Conference on Machine Learning*, pages 4907–4916. PMLR, 2018.
- Rohit Unni, Kan Yao, and Yuebing Zheng. Deep convolutional mixture density network for inverse design of layered photonic structures. *ACS Photonics*, 7(10):2703–2712, 2020.
- Matias Valdenegro-Toro and Daniel Saromo Mori. A deeper look into aleatoric and epistemic uncertainty disentanglement. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1508–1516. IEEE, 2022.
- Andrew G Wilson and Pavel Izmailov. Bayesian deep learning and a probabilistic perspective of generalization. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 4697–4708. Curran Associates, Inc., 2020a.
- Andrew G Wilson and Pavel Izmailov. Bayesian deep learning and a probabilistic perspective of generalization. *Advances in neural information processing systems*, 33:4697–4708, 2020b.
- CF Jeff Wu. On the convergence properties of the em algorithm. *The Annals of statistics*, pages 95–103, 1983.
- Heiga Zen and Andrew Senior. Deep mixture density networks for acoustic modeling in statistical parametric speech synthesis. In *2014 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 3844–3848. IEEE, 2014.
- Hao Zhang, Yongqian Liu, Jie Yan, Shuang Han, Li Li, and Quan Long. Improved deep mixture density network for regional wind power probabilistic forecasting. *IEEE Transactions on Power Systems*, 35(4):2549–2560, 2020.