

# Jensen-Shannon Divergence Based Novel Loss Functions for Bayesian Neural Networks

Ponkrshnan Thiagarajan <sup>a</sup>, Susanta Ghosh <sup>b,c,\*</sup>

<sup>a</sup> *Hopkins Extreme Materials Institute, Johns Hopkins University, MD, USA*

<sup>b</sup> *Department of Mechanical Engineering-Engineering Mechanics, Michigan Technological University, Houghton, MI, USA*

<sup>c</sup> *The Center for Artificial Intelligence at the Institute of Computing and Cybersystems and The Center for Applied Mathematics and Statistics, Michigan Technological University, Houghton, MI, USA*

---

## Abstract

Bayesian neural networks (BNNs) are state-of-the-art machine learning methods that can naturally regularize and systematically quantify uncertainties using their stochastic parameters. Kullback-Leibler (KL) divergence-based variational inference used in BNNs suffer from unstable optimization and challenges in approximating light-tailed posteriors due to the unbounded nature of the KL divergence. To resolve these issues, we formulate a novel loss function for BNNs based on a new modification to the generalized Jensen-Shannon (JS) divergence, which is bounded. In addition, we propose a Geometric JS divergence-based loss, which is computationally efficient since it can be evaluated analytically. We found that the JS divergence-based variational inference is intractable, and hence employed a constrained optimization framework to formulate these losses. Our theoretical analysis and empirical experiments on multiple regression and classification data sets suggest that the proposed losses perform better than the KL divergence-based loss, especially when the data sets are noisy or biased. Specifically, there are approximately 5% and 8% improvements in accuracy for a noise-added CIFAR-10 dataset and a regression dataset, respectively. There is about 13% reduction in false negative predictions of a biased histopathology dataset. In addition, we quantify and compare the uncertainty metrics for the regression and classification tasks.

*Keywords:* Bayesian Neural Networks, Variational Inference, KL divergence,

---

\*Corresponding Author: susantag@mtu.edu

## 1. Introduction

Despite the widespread success of deep neural networks (DNNs) and convolutional neural networks (CNNs) in numerous applications [1, 2], they suffer from *overfitting* when the data set is small, noisy, or biased [3, 4]. Furthermore, due to deterministic parameters, CNNs cannot provide a robust measure of uncertainty. Without a measure of uncertainty in the predictions, erroneous predictions by these models may lead to catastrophic failures in applications that require high accuracy, such as autonomous driving and medical diagnosis. Several methods have been developed to provide prediction intervals as a measure of uncertainty in neural networks [5]. Amongst these, Bayesian methods have gained eminence due to their rigorous mathematical foundation for uncertainty quantification through their stochastic parameters [6, 5]. In addition, Bayesian posteriors are robust to adversarial attacks [7]. Due to these advantages, Bayesian Neural Networks have been widely adapted to perform uncertainty quantification in a broad range of fields including but not limited to cosmology [8], nuclear physics [9], material science [10], structural health monitoring [11] etc.

A Bayesian neural network (BNN) has stochastic parameters whose posterior distribution is learned through the Bayes rule [12, 13, 14, 15]. Since the posterior distribution of parameters is intractable, multiple approaches are proposed to approximate them [16]. The two most commonly used techniques to approximate the posterior distribution of parameters are variational inference (VI) [17, 18, 19, 20] and Markov chain Monte Carlo methods (MCMC) [21, 22]. MCMC methods comprise a set of algorithms to sample from arbitrary and intractable probability distributions. Inference of posterior using MCMC algorithms can be very accurate but they are computationally demanding [23]. An additional limitation of MCMC algorithms is that they do not scale well with the model size.

The VI is a technique for approximating an intractable posterior distribution by a tractable distribution called the variational distribution. The variational distribution is learned by minimizing an objective function derived from its dissimilarity with respect to the true posterior [24]. VI methods are efficient and scale well for larger networks and have gained significant popularity. Most of the VI techniques in the literature use the KL divergence as a measure of

the aforementioned dissimilarity. However, the KL divergence is unbounded which may lead to failure during training, poor approximations, and difficulties in approximating light-tailed posteriors as reported in Hensman et al. [25], Dieng et al. [26], Deasy et al. [27]. Therefore, it is imperative to explore alternative divergences that can alleviate these limitations.

In terms of exploring alternate divergences, Renyi’s  $\alpha$ -divergences have been introduced for VI in Li and Turner [28]. They proposed a family of variational methods that unified various existing approaches. A  $\chi$ -divergence-based VI has been proposed in Dieng et al. [26] that provides an upper bound of the model evidence. Additionally, their results have shown better estimates for the variance of the posterior. Along these lines, an f-Divergence based VI has been proposed in Wan et al. [29] to use VI for all f-divergences that unified the Reyni divergence [28] and  $\chi$  divergence [26] based VIs. While these recent works [28, 26, 29] mainly focused on obtaining a generalized/unified VI framework, the present work specifically attempts to alleviate the limitations due to the unboundedness of the KL divergence-based VI through the Jensen-Shanon (JS) divergence. As a result, two novel loss functions are proposed, which outperform the KL loss in applications that require regularization. The choice of JS divergence is motivated by the fact that it is bounded and it provides an interpolation between forward and reverse KL divergences. In addition, a modification to the skew-geometric Jensen-Shanon (JS) divergence has been proposed in Deasy et al. [27] to introduce a new loss function for Variational Auto Encoders (VAEs), which has shown a better reconstruction and generation as compared to existing VAEs. Due to these reasons, JS divergence is explored as an alternative to alleviate the limitations of KL divergence in this work.

### 1.1. Key contributions

In the present work, we propose two novel loss functions for BNNs based on: 1) a novel modified generalized JS divergence (denoted as JS-A) that is bounded and 2) the geometric JS divergence (denoted as JS-G), which can be analytically evaluated. The primary contribution of this work is that it resolves the instability in optimization caused by the unboundedness of the KL divergence used in BNNs by formulating a JS-A divergence-based loss function. We found that the JS divergence-based variational inference is intractable and therefore employ a constrained optimization framework to formulate these losses.

We show that these JS divergence-based loss functions are generalizations of the state-of-the-art KL divergence-based ELBO loss function. In addition, through numerical analysis, we explain why these loss functions should perform better. Further, we derive the conditions under which the proposed JS-G divergence-based loss function regularises better than that of the KL divergence. Furthermore, we perform experiments on multiple regression and classification datasets to show that the proposed loss functions perform better, especially for problems where the data set has noise or is biased towards a particular class. Such biased datasets are often encountered in biomedical image classification, for example, a histopathology dataset used in this work which is biased towards the negative class. Similarly, examples of datasets with noise are encountered in adversarial attacks. By implementing the proposed method, the RMSE for a standard regression dataset is reduced by 8%, in the best-case scenario, as compared to the KL loss. Similarly, there is a 12% reduction in false negative predictions for a binary classification problem that is biased towards the positive class. For a multi-class classification dataset with added noise, the proposed method improves the accuracy by about 5% as compared to the KL loss. These are significant improvements given that the KL loss is highly successful and widely used in Bayesian neural networks. In our work, we provide both closed-form and MC-based algorithms for implementing the two JS divergences. The MC implementation can include priors of any family.

The present work is different from the existing work on JS divergence-based VI Deasy et al. [27] for the following reasons: **(i)** The JS-G divergence proposed in the previous work is unbounded like KL. This issue is resolved by the JS-A divergence proposed in this work. **(ii)** Deasy et al. [27] introduced the JS-G divergence-based loss for variational autoencoders (VAEs). In the present work, the distributions of parameters of BNNs are learned, which are numerous, as opposed to a small number of latent factors typically found in VAEs. **(iii)** The previous work is restricted to Gaussian priors due to the closed-form implementation, which is not the case in this work due to our MC implementation.

## 2. Mathematical Background

### 2.1. Background: KL and JS divergences

The KL divergence between two random variables  $\mathcal{P}$  and  $\mathcal{Q}$  on a probability space  $\Omega$  is defined as  $\text{KL}[p \parallel q] = \int_{\Omega} p(x) \log [p(x)/q(x)] dx$ , where  $p(x)$  and

$q(x)$  are the probability distributions of  $\mathcal{P}$  and  $\mathcal{Q}$  respectively.

The KL divergence is widely used in literature to represent the dissimilarity between two probability distributions for applications such as VI. However, it has limitations such as the unboundedness property, i.e. the divergence is infinite when  $q(x) = 0$  and  $p(x) \neq 0$ . This may lead to difficulty in approximating posteriors with a shorter tail than Gaussian distribution as reported in Hensman et al. [25]. In addition, KL divergence is asymmetric, i.e.  $\text{KL}[p \parallel q] \neq \text{KL}[q \parallel p]$  and hence does not qualify as a metric.

To overcome these limitations a symmetric JS divergence can be employed which is defined as  $\text{JS}[p \parallel q] = \frac{1}{2}\text{KL}[p \parallel (p+q)/2] + \frac{1}{2}\text{KL}[q \parallel (p+q)/2]$ .

It can be further generalized as,

$$\text{JS}^{A_\alpha}[p \parallel q] = (1 - \alpha)\text{KL}(p \parallel A_\alpha) + \alpha\text{KL}(q \parallel A_\alpha) \quad (1)$$

where,  $A_\alpha$  is the weighted arithmetic mean of  $p$  and  $q$  defined as  $A_\alpha = (1 - \alpha)p + \alpha q$ . Although this JS divergence is symmetric and bounded, unlike the KL divergence its analytical expression cannot be obtained even when  $p$  and  $q$  are Gaussians. To overcome this difficulty a generalization of the JS divergence using the geometric mean was proposed in Nielsen [30]. By using the weighted geometric mean  $G_\alpha(x, y) = x^{1-\alpha}y^\alpha$ , where  $\alpha \in [0, 1]$ , for two real variables  $x$  and  $y$ , they proposed the following family of skew geometric divergence

$$\text{JS}^{G_\alpha}[p \parallel q] = (1 - \alpha)\text{KL}(p \parallel G_\alpha(p, q)) + \alpha\text{KL}(q \parallel G_\alpha(p, q)) \quad (2)$$

The parameter  $\alpha$  called a skew parameter, controls the divergence skew between  $p$  and  $q$ . However, the skew-geometric divergence in Eq. 2 fails to capture the divergence between  $p$  and  $q$  and becomes zero for  $\alpha = 0$  and  $\alpha = 1$ . To resolve this issue, Deasy et al. [27] used the reverse form of geometric mean  $G'_\alpha(x, y) = x^\alpha y^{1-\alpha}$ , with  $\alpha \in [0, 1]$  for JS divergence to use in variational autoencoders. Henceforth, only this reverse form is used for the geometric mean. The JS divergence with this reverse of the geometric mean is given by

$$\text{JS-G}[p \parallel q] = (1 - \alpha)\text{KL}(p \parallel G'_\alpha(p, q)) + \alpha\text{KL}(q \parallel G'_\alpha(p, q)) \quad (3)$$

This yields KL divergences in the limiting values of the skew parameter. Note that for  $\alpha \in [0, 1]$ ,  $\text{JS-G}(p \parallel q)|_\alpha = \text{JS-G}(p \parallel q)|_{1-\alpha}$  which is not symmetric. However, for  $\alpha = 0.5$ , the JS-G is symmetric with  $\text{JS-G}(p \parallel q)|_{\alpha=0.5} = \text{JS-G}(p \parallel q)|_{\alpha=0.5}$ . The geometric JS divergences,  $\text{JS}^{G_\alpha}$  and JS-G given in

Eq. 2 and Eq. 3 respectively, have analytical expressions when  $p$  and  $q$  are Gaussians. However, they are unbounded like the KL divergence. Whereas, the generalized JS divergence  $JS^{A\alpha}$  in Eq. 1 is both bounded and symmetric.

## 2.2. Background: Variational inference

Given a set of training data  $\mathbb{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$  and test input,  $\mathbf{x} \in \mathbb{R}^p$ , we learn a data-driven model (e.g a BNN) to predict the probability  $P(\mathbf{y}|\mathbf{x}, \mathbb{D})$  of output  $\mathbf{y} \in \Upsilon$ , where  $\Upsilon$  is the output space. The posterior probability distribution ( $P(\mathbf{w}|\mathbb{D})$ ) of the parameters ( $\mathbf{w}$ ) of BNN, can be obtained using the Bayes' rule:  $P(\mathbf{w}|\mathbb{D}) = P(\mathbb{D}|\mathbf{w})P(\mathbf{w})/P(\mathbb{D})$ .

Where  $P(\mathbb{D}|\mathbf{w})$  and  $P(\mathbf{w})$  are the likelihood term and the prior distribution respectively. The term  $P(\mathbb{D})$ , called the evidence, involves marginalization over the distribution of weights:  $P(\mathbb{D}) = \int_{\Omega_{\mathbf{w}}} P(\mathbb{D}|\mathbf{w})P(\mathbf{w})d\mathbf{w}$ . Using the posterior distribution of weights, the predictive distribution of the output can be obtained by marginalizing the weights as  $P(\mathbf{y}|\mathbf{x}, \mathbb{D}) = \int_{\Omega_{\mathbf{w}}} P(\mathbf{y}|\mathbf{x}, \mathbf{w})P(\mathbf{w}|\mathbb{D})d\mathbf{w}$ .

The term  $P(\mathbb{D})$  in the Bayes' rule is intractable due to marginalization over  $\mathbf{w}$ , which in turn makes  $P(\mathbf{w}|\mathbb{D})$  intractable. To alleviate this difficulty, the posterior is approximated using variational inference.

In variational inference, the unknown intractable posterior  $P(\mathbf{w}|\mathbb{D})$  is approximated by a known simpler distribution  $q(\mathbf{w}|\boldsymbol{\theta})$  called the variational posterior having parameters  $\boldsymbol{\theta}$ .

The set of parameters  $\boldsymbol{\theta}$  for the model weights are learned by minimizing the divergence (e.g. KL divergence) between  $P(\mathbf{w}|\mathbb{D})$  and  $q(\mathbf{w}|\boldsymbol{\theta})$  as shown in Blundell et al. [24].

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} \text{KL} [q(\mathbf{w}|\boldsymbol{\theta}) || P(\mathbf{w}|\mathbb{D})] \quad (4)$$

$$= \arg \min_{\boldsymbol{\theta}} \int q(\mathbf{w}|\boldsymbol{\theta}) \left[ \log \frac{q(\mathbf{w}|\boldsymbol{\theta})}{P(\mathbf{w})P(\mathbb{D}|\mathbf{w})} + \log P(\mathbb{D}) \right] d\mathbf{w} \quad (5)$$

Note that the term  $\log P(\mathbb{D})$  in Eq. 4 is independent of  $\boldsymbol{\theta}$  and thus can be eliminated. The resulting loss function  $\mathcal{F}(\mathbb{D}, \boldsymbol{\theta})$ , which is to be minimised to learn the optimal parameters  $\boldsymbol{\theta}^*$  is expressed as:

$$\mathcal{F}_{KL}(\mathbb{D}, \boldsymbol{\theta}) = \text{KL} [q(\mathbf{w}|\boldsymbol{\theta}) || P(\mathbf{w})] - \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})}[\log P(\mathbb{D}|\mathbf{w})] \quad (6)$$

The negative of this loss is known as the evidence lower bound (ELBO) [19, 24].

### 3. Methods

In this section, we provide a modification to the generalized JS divergence, formulations of JS divergence-based loss functions for BNNs, and insights into the advantages of the proposed loss.

#### 3.1. Proposed modification to the generalized JS divergence

The generalised JS divergence given in Eq. 1 fails to capture the divergence between  $p$  and  $q$  in the limiting cases of  $\alpha$  since,

$$[\text{JS}^{\text{A}\alpha}(p \parallel q)]_{\alpha=0} = 0; \quad [\text{JS}^{\text{A}\alpha}(p \parallel q)]_{\alpha=1} = 0 \quad (7)$$

To overcome this limitation we propose to modify the weighted arithmetic mean as  $A'_\alpha = \alpha p + (1 - \alpha)q$ ,  $\alpha \in [0, 1]$  which modifies the generalized JS divergence as

$$\text{JS-A}[p \parallel q] = (1 - \alpha)\text{KL}(p \parallel A'_\alpha) + \alpha\text{KL}(q \parallel A'_\alpha) \quad (8)$$

Hence, this yields KL divergences in the limiting cases of  $\alpha$  as,

$$[\text{JS-A}(p \parallel q)]_{\alpha=0} = \text{KL}(p \parallel q) \quad (9)$$

$$[\text{JS-A}(p \parallel q)]_{\alpha=1} = \text{KL}(q \parallel p) \quad (10)$$

Eq. 8 ensures that  $\text{JS}(P \parallel q) = 0$  if and only if  $P = q$ ,  $\forall \alpha \in [0, 1]$ . This is necessary since the divergence is a metric to represent statistical dissimilarity.

**Theorem 1:** Boundedness of the modified generalized JS divergence

*For any two distributions  $P_1(t)$  and  $P_2(t)$ ,  $t \in \Omega$ , the value of the JS-A divergence is bounded such that,*

$$\text{JS-A}(P_1(t) \parallel P_2(t)) \leq -(1 - \alpha) \log \alpha - \alpha \log(1 - \alpha), \quad \text{for } \alpha \in (0, 1) \quad (11)$$

The proof of Theorem 1 is presented in App. Appendix B. Due to this boundedness property of the JS-A divergence, the ensuing loss functions overcome the instability in optimization that is encountered in the KL divergence-based loss.

The unboundedness of the KL and JS-G divergence and the boundedness of JS-A divergence are depicted in Fig 1. Note that, when both the distributions  $P$  and  $q$  are heavy-tailed, such as the Gaussian distribution, KL divergence does not become singular. However, this restricts the choice of the prior significantly.

We provide a comparison of various divergences used in this work in App. Appendix A and Appendix B.

### 3.2. Intractability of the JS divergence-based loss functions formulated through the variational inference approach

In this subsection, we demonstrate that the JS divergence-based variational inference is intractable. If the JS-G divergence is used instead of the KL divergence in the VI setting (see Eq. 4), the optimization problem becomes,

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} \text{JS-G} [q(\mathbf{w}|\boldsymbol{\theta}) \parallel P(\mathbf{w}|\mathbb{D})] \quad (12)$$

The loss function can then be written as,

$$\begin{aligned} \mathcal{F}_{\text{JS-G}}(\mathbb{D}, \boldsymbol{\theta}) &= \text{JS-G} [q(\mathbf{w}|\boldsymbol{\theta}) \parallel P(\mathbf{w}|\mathbb{D})] \\ &= (1 - \alpha) \text{KL} (q \parallel G'_\alpha(q, P)) + \alpha \text{KL} (P \parallel G'_\alpha(q, P)) \end{aligned} \quad (13)$$

Where,  $G'_\alpha(q, P) = q(\mathbf{w}|\boldsymbol{\theta})^\alpha P(\mathbf{w}|\mathbb{D})^{(1-\alpha)}$ . Rewriting the first and the second term in Eq. 13 as,

$$T_1 = (1 - \alpha)^2 \int q(\mathbf{w}|\boldsymbol{\theta}) \log \left[ \frac{q(\mathbf{w}|\boldsymbol{\theta})}{P(\mathbf{w}|\mathbb{D})} \right] d\mathbf{w}; \quad (14)$$

$$T_2 = \alpha^2 \int P(\mathbf{w}|\mathbb{D}) \log \left[ \frac{P(\mathbf{w}|\mathbb{D})}{q(\mathbf{w}|\boldsymbol{\theta})} \right] d\mathbf{w} \quad (15)$$




Figure 1: Depiction of the unboundedness (denoted by  $\infty$ ) of the KL and JS-G divergence and the boundedness of the JS-A divergence. The distributions  $q$  and  $P$  are assumed Gaussian and Uniform respectively with  $q = \mathcal{N}(0, \sigma)$  and  $P = \mathcal{U}(-5, 5)$ .



A detailed derivation of terms  $T_1$  and  $T_2$  is given in App. Appendix C. Term  $T_1$  is equivalent to the loss function in Eq. 6 multiplied by a constant  $(1 - \alpha)^2$ .

The term  $P(\mathbf{w}|\mathbb{D})$  in  $T_2$  is intractable as explained in section 2.2. Therefore the JS-G divergence-based loss function given in Eq. 13 cannot be used to find the optimum parameter  $\boldsymbol{\theta}^*$  which contrasts the KL divergence-based loss function in Eq. 6. Similarly, the JS-A divergence-based loss function obtained through VI is also intractable. We address this issue of intractability in the following subsection.

### 3.3. Proposed JS divergence-based loss functions formulated through a constrained optimization approach

To overcome the intractability of the variational inference, we propose to use a constrained optimization framework, following Higgins et al. [31], Deasy et al. [27], to derive JS divergence-based loss functions for BNNs. We also show that such a loss function is a generalization of the loss function obtained through the variational inference.

Given a set of training data  $\mathbb{D}$ , we are interested in learning the probability distribution  $q(\mathbf{w}|\boldsymbol{\theta})$  of network parameters such that, the likelihood of observing the data given the parameters is maximized. Thus, the optimization problem can be written as

$$\max_{\boldsymbol{\theta}} \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} [\log P(\mathbb{D}|\mathbf{w})] \quad (16)$$

Where  $\boldsymbol{\theta}$  is a set of parameters of the probability distribution  $q(\mathbf{w}|\boldsymbol{\theta})$ . This optimization is constrained to make  $q(\mathbf{w}|\boldsymbol{\theta})$  similar to a prior  $P(\mathbf{w})$ . This leads to a constrained optimization problem as given below:

$$\max_{\boldsymbol{\theta}} \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} [\log P(\mathbb{D}|\mathbf{w})] \quad \text{subject to } D(q(\mathbf{w}|\boldsymbol{\theta}) || P(\mathbf{w})) < \epsilon \quad (17)$$

where  $\epsilon$  is a real number that determines the strength of the applied constraint and  $D$  is a divergence measure. Following the KKT approach, the Lagrangian function corresponding to the constrained optimization problem can be written as

$$\mathcal{L} = \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} [\log P(\mathbb{D}|\mathbf{w})] - \lambda(D(q(\mathbf{w}|\boldsymbol{\theta}) || P(\mathbf{w})) - \epsilon) \quad (18)$$

Since  $\epsilon$  is a constant it can be removed from the optimization. Also changing the sign of the above equations leads to the following loss function that needs

to be minimized.<sup>1</sup>

$$\tilde{\mathcal{F}}_D = \lambda D(q(\mathbf{w}|\boldsymbol{\theta}) \parallel P(\mathbf{w})) - \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} [\log P(\mathbb{D}|\mathbf{w})] \quad (19)$$

This loss function reproduces the ELBO loss [24] when KL divergence is used and  $\lambda$  is taken as 1.

In the following, we obtain loss functions based on the JS-A and JS-G divergences.

### 3.3.1. Geometric JS divergence

Using the modified geometric JS divergence (JS-G) as the measure of divergence in Eq. 19 leads to the following loss function:

$$\begin{aligned} \tilde{\mathcal{F}}_{JSG} &= \lambda \text{JS-G}(q(\mathbf{w}|\boldsymbol{\theta}) \parallel P(\mathbf{w})) - \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} [\log P(\mathbb{D}|\mathbf{w})] \\ &= \lambda(1 - \alpha) \text{KL}(q \parallel G'_\alpha(q, P_w)) + \lambda\alpha \text{KL}(P_w \parallel G'_\alpha(q, P_w)) - \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} [\log P(\mathbb{D}|\mathbf{w})] \end{aligned} \quad (20a)$$

Note,

$$\begin{aligned} \text{KL}(q \parallel G'_\alpha(q, P_w)) &= \int q(\mathbf{w}|\boldsymbol{\theta}) \log \frac{q(\mathbf{w}|\boldsymbol{\theta})}{q(\mathbf{w}|\boldsymbol{\theta})^\alpha P(\mathbf{w})^{1-\alpha}} d\mathbf{w} \\ &= (1-\alpha) \int q(\mathbf{w}|\boldsymbol{\theta}) \log \frac{q(\mathbf{w}|\boldsymbol{\theta})}{P(\mathbf{w})} d\mathbf{w} \\ \text{KL}(P_w \parallel G'_\alpha(q, P_w)) &= \int P(\mathbf{w}) \log \frac{P(\mathbf{w})}{q(\mathbf{w}|\boldsymbol{\theta})^\alpha P(\mathbf{w})^{1-\alpha}} d\mathbf{w} \\ &= \alpha \int P(\mathbf{w}) \log \frac{P(\mathbf{w})}{q(\mathbf{w}|\boldsymbol{\theta})} d\mathbf{w} \end{aligned}$$

Hence, the loss function can be written as,

$$\begin{aligned} \tilde{\mathcal{F}}_{JSG} &= \lambda(1 - \alpha)^2 \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} \left[ \log \frac{q(\mathbf{w}|\boldsymbol{\theta})}{P(\mathbf{w})} \right] \\ &\quad + \lambda\alpha^2 \mathbb{E}_{P(\mathbf{w})} \left[ \log \frac{P(\mathbf{w})}{q(\mathbf{w}|\boldsymbol{\theta})} \right] - \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} [\log P(\mathbb{D}|\mathbf{w})] \end{aligned} \quad (21)$$

In Eq. 21, the first term is the *mode seeking* reverse KL divergence  $\text{KL}(q(\mathbf{w}|\boldsymbol{\theta}) \parallel P(\mathbf{w}))$  and the second term is the *mean seeking* forward KL

---

<sup>1</sup> The constrained optimization approach-based loss functions are marked by an overhead tilde.

divergence  $\text{KL}(P(\mathbf{w})||q(\mathbf{w}|\boldsymbol{\theta}))$ . Therefore, the proposed loss function offers a weighted sum of the forward and reverse KL divergences in contrast to only the reverse KL divergence in ELBO. Whereas the likelihood part remains identical. The relative weighting between the forward and the reverse KL divergences can be controlled by the parameter  $\alpha$ . The proposed loss function would ensure better regularisation by imposing stricter penalization if the posterior is away from the prior distribution which will be demonstrated in detail in Sec. 5.0.1. The parameters  $\lambda$  and  $\alpha$  can be used to control the amount of regularisation <sup>2</sup>.

### 3.3.2. Modified generalized JS divergence

Using the modified Generalised JS divergence (JS-A) as the measure of divergence in Eq. 19 leads to the following loss function:

$$\begin{aligned}\tilde{\mathcal{F}}_{JSA} &= \lambda \text{JS-A}(q(\mathbf{w}|\boldsymbol{\theta}) || P(\mathbf{w})) - \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} [\log P(\mathbb{D}|\mathbf{w})] \\ &= \lambda(1 - \alpha) \text{KL}(q || A'_\alpha(q, P_w)) + \lambda\alpha \text{KL}(P_w || A'_\alpha(q, P_w)) - \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} [\log P(\mathbb{D}|\mathbf{w})]\end{aligned}\quad (22)$$

Where,  $A'_\alpha(q, P_w) = \alpha q + (1 - \alpha)P_w$ . The above equation, Eq. 22, can be expanded as,

$$\begin{aligned}\tilde{\mathcal{F}}_{JSA} &= \lambda(1 - \alpha)\mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} \left[ \log \frac{q(\mathbf{w}|\boldsymbol{\theta})}{A'_\alpha(q, P_w)} \right] \\ &\quad + \lambda\alpha\mathbb{E}_{P(\mathbf{w})} \left[ \log \frac{P(\mathbf{w})}{A'_\alpha(q, P_w)} \right] - \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} [\log P(\mathbb{D}|\mathbf{w})]\end{aligned}\quad (23)$$

Note that the proposed loss functions in Eq. 21 and Eq. 23 yield the ELBO loss for  $\alpha = 0$  and  $\lambda = 1$ . The minimization algorithms for the loss functions Eq. 21 and Eq. 23 are given in the following section.

## 4. Minimisation of the proposed loss functions

Note that when training using mini-batches, the divergence part of the loss function is normalized by the size of the minibatch ( $M$ ). Therefore, loss for the minibatch  $i = 1, 2, 3, \dots, M$  can be written as,

$$\tilde{\mathcal{F}}_i = \frac{\lambda}{M} \text{D}(q(\mathbf{w}|\boldsymbol{\theta}) || P(\mathbf{w})) - \mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} [\log P(\mathbb{D}_i|\mathbf{w})] \quad (24)$$

---

**Algorithm 1** Minimization of the JS-G loss function: Closed-form evaluation of the divergence

---

**Initialize**  $\boldsymbol{\mu}, \boldsymbol{\rho}$

**Evaluate** JS-G term of Eq. 20a analytically using Eq. 25

**Evaluate**  $\mathbb{E}_{q(\mathbf{w}|\boldsymbol{\theta})} [\log P(\mathbb{D}|\mathbf{w})]$  term of Eq. 20a

Sample  $\boldsymbol{\varepsilon}_i \sim \mathcal{N}(0, 1); i = 1, \dots, \text{No. of samples}$

$\mathbf{w}_i \leftarrow \boldsymbol{\mu} + \log(1 + \exp(\boldsymbol{\rho})) \circ \boldsymbol{\varepsilon}_i.$

$f_1 \leftarrow \sum_i \log P(\mathbb{D}|\mathbf{w}_i)$

**Loss:**

$$F \leftarrow \lambda \text{JS-G} - f_1$$

**Gradients:**

$$\begin{aligned} \frac{\partial F}{\partial \boldsymbol{\mu}} &\leftarrow \sum_i \frac{\partial F}{\partial \mathbf{w}_i} + \frac{\partial F}{\partial \boldsymbol{\mu}} \\ \frac{\partial F}{\partial \boldsymbol{\rho}} &\leftarrow \sum_i \frac{\partial F}{\partial \mathbf{w}_i} \frac{\boldsymbol{\varepsilon}_i}{1 + \exp(-\boldsymbol{\rho})} + \frac{\partial F}{\partial \boldsymbol{\rho}} \end{aligned}$$

**Update:**

$$\boldsymbol{\mu} \leftarrow \boldsymbol{\mu} - \beta \frac{\partial F}{\partial \boldsymbol{\mu}}; \quad \boldsymbol{\rho} \leftarrow \boldsymbol{\rho} - \beta \frac{\partial F}{\partial \boldsymbol{\rho}}$$


---

#### 4.1. Evaluation of the JS-G divergence in a closed-form

In this subsection, we describe the minimization of the JS-G divergence-based loss function by evaluating the divergence in closed form for Gaussian priors. Assuming the prior and the likelihood are Gaussians, the posterior will also be a Gaussian. Let the prior and posterior be diagonal multivariate Gaussian distributions denoted by  $P_N(\mathbf{w}|\boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2^2)$  and  $q_N(\mathbf{w}) = \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1^2)$  respectively<sup>3</sup>. The closed-form expression of the JS-G divergence

---

<sup>2</sup> $\lambda$  is taken as 1 for  $\tilde{\mathcal{F}}_{JS-G}$  in this work unless otherwise stated.

<sup>3</sup> Where the subscript  $(\cdot)_N$  indicates Gaussian distribution.  $\boldsymbol{\mu}_1$  and  $\boldsymbol{\mu}_2$  are n-dimensional vectors and  $\boldsymbol{\Sigma}_1^2, \boldsymbol{\Sigma}_2^2$  are assumed to be diagonal matrices such that  $\boldsymbol{\mu}_1 = [\mu_{11}, \mu_{12}, \dots, \mu_{1n}]^T$

between  $q_N(\mathbf{w})$  and  $P_N(\mathbf{w}|\boldsymbol{\theta})$  can be written as,

$$\begin{aligned} \text{JS-G}(q_N||P_N) = \frac{1}{2} \sum_{i=1}^n & \left[ \frac{(1-\alpha)\sigma_{1i}^2 + \alpha\sigma_{2i}^2}{\sigma_{\alpha i}^2} + \log \frac{(\sigma'_{\alpha i})^2}{\sigma_{1i}^{2(1-\alpha)}\sigma_{2i}^{2\alpha}} \right. \\ & \left. + (1-\alpha)\frac{(\mu'_{\alpha i} - \mu_{1i})^2}{(\sigma'_{\alpha i})^2} + \frac{\alpha(\mu'_{\alpha i} - \mu_{2i})^2}{(\sigma'_{\alpha i})^2} - 1 \right] \end{aligned} \quad (25)$$

where,

$$(\sigma'_{\alpha i})^2 = \frac{\sigma_{1i}^2\sigma_{2i}^2}{(1-\alpha)\sigma_{1i}^2 + \alpha\sigma_{2i}^2}; \quad \mu'_{\alpha i} = (\sigma'_{\alpha i})^2 \left[ \frac{\alpha\mu_{1i}}{\sigma_{1i}^2} + \frac{(1-\alpha)\mu_{2i}}{\sigma_{2i}^2} \right]$$

Therefore, the divergence term of the proposed loss function, the first term in Eq. 20a, can be evaluated by this closed-form expression given in Eq. 25. The expectation of the log-likelihood, the second term in Eq. 20a, can be approximated by a Monte-Carlo sampling <sup>4</sup>. The details of the minimization process are given in Algorithm 1. Note, for sampling  $w_i$  the reparametrization trick is used to separate the deterministic and the stochastic variables.

#### 4.2. Evaluation of divergences via a Monte Carlo sampling

In this subsection, we describe the minimization of the JS divergence-based loss functions by evaluating the divergences using the Monte Carlo sampling technique. The algorithm provided in this subsection is more general as it applies to both the JS-G and the JS-A divergences with no restrictions on the priors. The loss functions given in Eq. 21 and Eq. 23 can be approximated using Monte Carlo samples from the corresponding distributions as shown in Algorithm 2.

### 5. Insights into the proposed JS divergence-based loss functions

To better understand the proposed JS divergence-based loss functions, we use a contrived example to compare them against the conventional KL divergence-based loss function. In the following, we explore the regularization ability of the proposed loss functions. Further insights on Monte Carlo estimates are given in App. Appendix D

---

**Algorithm 2** Minimization of the JS-G and JS-A loss functions: Monte Carlo approximation of the divergence

---

**Initialize**  $\mu, \rho$

**Approximate**  $\mathbb{E}_{q(\mathbf{w}|\theta)}$  terms of Eq. 21 or Eq. 23

Sample  $\epsilon_i^q \sim \mathcal{N}(0, 1); i = 1, \dots, \text{No. of samples}$

$\mathbf{w}_i^q \leftarrow \mu + \log(1 + \exp(\rho)) \circ \epsilon_i^q$ .

Evaluate first and third terms of Eq. 21:

$$f_1 \leftarrow \sum_i c_1 \log q(\mathbf{w}_i^q | \theta) - c_1 \log P(\mathbf{w}_i^q) - \log P(\mathbb{D} | \mathbf{w}_i^q)$$

where,  $c_1 = \lambda(1 - \alpha)^2$

(or)

Evaluate first and third terms of Eq. 23:

$$f_1 \leftarrow \sum_i c_1 \log q(\mathbf{w}_i^q | \theta) - c_1 \log A_\alpha(\mathbf{w}_i^q) - \log P(\mathbb{D} | \mathbf{w}_i^q)$$

where,  $c_1 = \lambda(1 - \alpha)$

**Approximate**  $\mathbb{E}_{P(\mathbf{w})}$  terms of Eq. 21 or Eq. 23

Sample  $\mathbf{w}_j^p \sim P(\mathbf{w}); j = 1, \dots, \text{No. of samples}$

Evaluate second term of Eq. 21:

$$f_2 \leftarrow \sum_j c_2 \log P(\mathbf{w}_j^p) - c_2 \log q(\mathbf{w}_j^p | \theta)$$

where,  $c_2 = \lambda\alpha^2$

(or)

Evaluate second term of Eq. 23:

$$f_2 \leftarrow \sum_j c_2 \log A'_\alpha(\mathbf{w}_j^p) - c_2 \log q(\mathbf{w}_j^p | \theta)$$

where,  $c_2 = \lambda\alpha^2$

**Loss:**

$$F \leftarrow f_1 + f_2$$

**Gradients:**

$$\begin{aligned} \frac{\partial F}{\partial \mu} &\leftarrow \sum_i \frac{\partial F}{\partial \mathbf{w}_i^q} + \frac{\partial F}{\partial \mu} \\ \frac{\partial F}{\partial \rho} &\leftarrow \sum_i \frac{\partial F}{\partial \mathbf{w}_i^q} \frac{\epsilon_i}{1 + \exp(-\rho)} + \frac{\partial F}{\partial \rho} \end{aligned}$$

**Update:**

14

$$\mu \leftarrow \mu - \beta \frac{\partial F}{\partial \mu}; \quad \rho \leftarrow \rho - \beta \frac{\partial F}{\partial \rho}$$


---

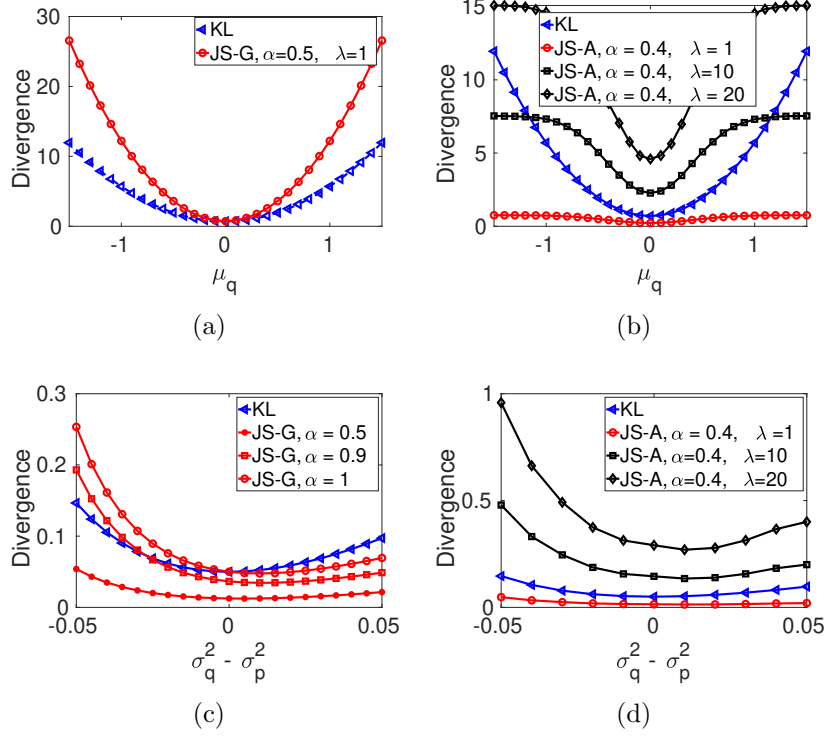


Figure 2: Comparison of the KL and the JS divergences of distributions  $P$  and  $q$ . (a) and (b)  $\sigma_q^2, \mu_p, \sigma_p^2$  are fixed and  $\mu_q$  is varied. (c) and (d)  $\mu_q, \mu_p, \sigma_p^2$  are fixed and  $\sigma_q^2$  is varied. The fixed values of the parameters are  $\mu_q = 0.1, \sigma_q^2 = 0.01, \mu_p = 0, \sigma_p^2 = 0.1$

### 5.0.1. Regularisation performance of JS divergences

Let two Gaussian distributions  $q = \mathcal{N}(\mu_q, \sigma_q^2)$  and  $P = \mathcal{N}(\mu_p, \sigma_p^2)$  represent the posterior and the prior distribution of a parameter in a BNN. The KL, JS-A, and JS-G divergences are evaluated by varying the mean and variance of the distribution  $q$ . This emulates the learning of the network parameter during training. Fig. 2 shows that as the posterior distribution ( $q$ ) moves away from the prior distribution ( $P$ ), the JS divergences increase more rapidly than the KL divergence. In the case of the JS-A divergence in Fig.2(b) and 2(d), this is achieved by a higher value of  $\lambda$ . This implies that a greater

and  $\Sigma_1^2 = \text{diag}(\sigma_{11}^2, \sigma_{12}^2, \dots, \sigma_{1n}^2)$  (similarly for  $\mu_2$  and  $\Sigma_2^2$ ).

<sup>4</sup>The approximations to the loss functions are denoted by F

penalization is offered by JS divergences than the KL divergence as the posterior deviates away from the prior (also seen in App. 7.6). Thus, by assuming small values for the means of prior distributions we can regularize better by the proposed JS divergences. In practice, zero mean Gaussian priors are widely accepted for BNNs. For such priors, higher penalization of the loss function implies pushing the parameters' mean closer to zero while learning the complexity of the data. In doing this, we can achieve better regularization. This regularization process requires finding optimal values of  $\alpha$  and  $\lambda$  through hyperparameter optimization. In the following subsection, we theoretically analyze the regularization performance of the JS-G divergence.

### 5.0.2. Condition for better regularisation of $\tilde{\mathcal{F}}_{JSG}$

The above example shows that the JS-G divergence is greater than the KL for the given Gaussian distributions. To generalize it further, we propose the following theorems that hold for any two arbitrary distributions.

**Theorem 2.** *For any two arbitrary distributions  $P$  and  $q$  such that  $P \neq q$ ,  $\tilde{\mathcal{F}}_{JSG} > \mathcal{F}_{KL}$  if and only if  $\alpha > \frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} \in (0, \infty)$*

**Proof:** Assuming,  $\tilde{\mathcal{F}}_{JSG} - \mathcal{F}_{KL} > 0$  and from Eq. 6 and Eq. 21 we have,

$$\begin{aligned} (1 - \alpha)^2 \text{KL}(q||P) + \alpha^2 \text{KL}(P||q) - \text{KL}(q||P) &> 0 \\ (\alpha^2 - 2\alpha) \text{KL}(q||P) + \alpha^2 \text{KL}(P||q) &> 0 \end{aligned}$$

$$\text{This leads to, } \alpha > \frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)}$$

This proves that if  $\tilde{\mathcal{F}}_{JSG} > \mathcal{F}_{KL}$  then  $\alpha > \frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)}$ . The converse can be proved similarly. A detailed proof is shown in the App. Appendix E.

**Theorem 3.** *If  $P = \mathcal{N}(\mu_p, \sigma_p^2)$  and  $q = \mathcal{N}(\mu_q, \sigma_q^2)$  are Gaussian distributions and  $P \neq q$ , then  $\frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} < 1$  if and only if  $\sigma_p^2 > \sigma_q^2$ .*

**Proof:** Assuming  $\frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} < 1$ , we get

$$\text{KL}(P||q) > \text{KL}(q||P) \tag{26}$$

Since  $P = \mathcal{N}(\mu_p, \sigma_p^2)$  and  $q = \mathcal{N}(\mu_q, \sigma_q^2)$ , Eq. (26) can be written as,

$$\ln \frac{\sigma_q^2}{\sigma_p^2} + \frac{\sigma_p^2 + (\mu_q - \mu_p)^2}{\sigma_q^2} - 1 > \ln \frac{\sigma_p^2}{\sigma_q^2} + \frac{\sigma_q^2 + (\mu_p - \mu_q)^2}{\sigma_p^2} - 1$$



Denoting  $\gamma = \frac{\sigma_p^2}{\sigma_q^2}$ , we get,

$$\gamma - \frac{1}{\gamma} + \ln \frac{1}{\gamma} - \ln \gamma + \frac{(\mu_q - \mu_p)^2}{\sigma_q^2} - \frac{(\mu_p - \mu_q)^2}{\gamma \sigma_q^2} > 0$$

$$\text{or, } \ln \left[ \frac{1}{\gamma^2} \exp \left( \gamma - \frac{1}{\gamma} \right) \right] + \frac{(\mu_q - \mu_p)^2}{\sigma_q^2} \left( 1 - \frac{1}{\gamma} \right) > 0 \quad (27)$$

This condition Eq. 27 is satisfied only when  $\gamma > 1$ , which implies  $\sigma_p^2 > \sigma_q^2$ . Thus if  $\frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} < 1$  then  $\sigma_p^2 > \sigma_q^2$ . This result is also observed in Fig. 2(c). The converse can be proved similarly as shown in App. Appendix F.

**Corollary:** From Theorem 2 and 3:  $\tilde{\mathcal{F}}_{JSG} > \mathcal{F}_{KL}$  if  $\sigma_p^2 > \sigma_q^2$  and  $\forall \alpha \in (0, 1]$  such that  $\alpha > \frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)}$ . Where,  $P$  and  $q$  are Gaussians and  $P \neq q$ .

## 6. Experiments

In order to demonstrate the advantages of the proposed losses in comparison to the KL loss, we performed experiments. We have implemented the divergence part of the JS-G loss and the JS-A loss via a closed-form expression and a Monte-Carlo method respectively, in these experiments.

### 6.1. Data sets

The experiments were performed on multiple regression and classification datasets. Regression experiments were performed on six data sets (Airfoil, Aquatic, Building, Concrete, Real Estate, and Wine) from the UCI Machine Learning Repository [32].

For classification two data sets are considered: the CIFAR-10 data set [33] and a histopathology data set [34, 35, 36]. To demonstrate the effectiveness of regularisation, varying levels of Gaussian noise were added to the normalized CIFAR-10 data set for training, validation, and testing. We also used a histopathology data set which is highly biased towards one class. Further details on these data sets and the pre-processing steps used here are provided in Appendix G.

## 6.2. Hyperparameter optimisation and network architecture

Regression experiments on multiple datasets are conducted following the framework of [29, 28]. The network architecture is a single layer with 50 hidden units and a ReLU activation function. We assume the priors to be Gaussians  $p \sim \mathcal{N}(0, I)$  and 100 MC approximations are used to calculate the NLL part of the loss function for the KL and JS-G divergence-based losses, and 10 MC samples are used to evaluate the JS-A divergence based loss. The likelihood function is Gaussian with noise parameter  $\sigma$  which is also learned along with the parameters of the network. Six datasets (Airfoil, Aquatic, Building, Concrete, Real Estate, and Wine) are evaluated and they are split into 90% and 10% for training and testing respectively.

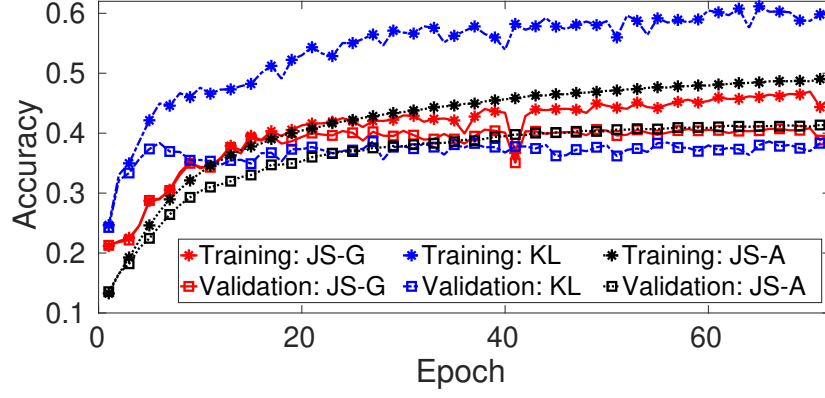
Hyperparameters of the networks used for classification are chosen through hyperparameter optimization. It is performed using a Tree-structured Parzen Estimator (TPE) algorithm [37] to maximize the validation accuracy for different hyperparameter settings of the network. Further details and the results of the hyperparameter optimization are given in Appendix H. The architecture of the networks used for classification follows the ResNet-18 V1 model [38] without the batch normalization layers. The network parameters are initialized with the weights of ResNet-18 trained on the Imagenet data set [39].

## 7. Classification results

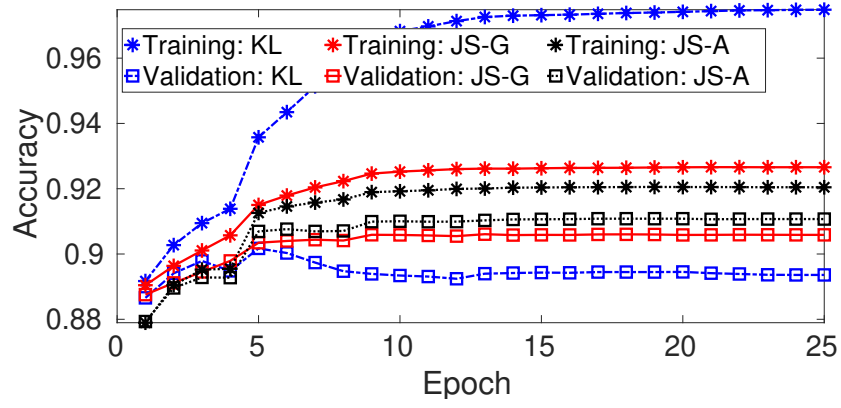
This section presents the results of the classification experiments and the performance comparison between the KL loss and the proposed JS losses.

### 7.1. Training and validation

Three Bayesian CNNs were trained by minimizing the KL and proposed JS losses. Training of the networks is done until the loss converges or the validation accuracy starts to decrease. Training of the CIFAR-10 data set is performed with varying levels of noise intensity. Accuracy of training and validation sets for noise  $\mathcal{N}(\mu = 0, \sigma = 0.9)$  is presented for both KL loss and the proposed JS losses in Fig. 3(a). For the histopathology data set, a learning rate scheduler is used during training in which the learning rate is multiplied by a factor of 0.1 in the 4th, 8th, 12th, and 20th epochs. Fig. 3(b) shows the accuracy of training and validation of the histopathology set for the KL loss and the proposed JS losses. It is evident that the KL loss learns the training data too well and fails to generalize for the unseen validation



(a)



(b)

Figure 3: Training and validation of (a) CIFAR-10 with added Gaussian noise (b) histopathology data set with bias.

set on both data sets. Whereas, the proposed JS losses regularise better and provide more accurate results for the validation set.

### 7.2. Testing

Results obtained on the test sets of the CIFAR-10 data set and the histopathology data set are presented in this section. The test results correspond to the epoch in which the validation accuracy was maximum. Five runs were performed with different mutually exclusive training and validation tests to compare the results of the KL loss and the proposed JS losses. The accuracy of the noisy CIFAR-10 test data set at varying noise levels is presented in Fig. 4(a) and Fig. 4(b). It is evident that the accuracy of both the proposed JS losses is better than KL for all the noise level cases. Further, the difference in accuracy between KL loss and the JS losses shows an increasing trend with increasing noise levels. This demonstrates the regularising capability of the proposed JS losses. The results of the five runs of the KL loss and the proposed JS losses on the biased histopathology data set are compared in Fig. 4(c). It is evident that both the proposed JS losses perform better than the KL loss in all five runs with different training and validation sets. Since this data set is biased toward the negative class, the performance improvement shown by the proposed JS losses is attributed to better regularisation and generalization capabilities of the loss functions.

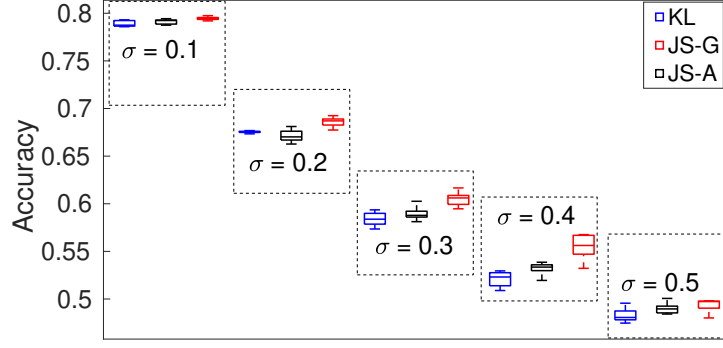
The receiver operating characteristic (ROC) curve is plotted in Fig. 4(d) for the classification of the histopathology data set. The proposed JS losses perform better than the KL loss in terms of the area under the curve. The confusion matrices in Fig. 5(a)-5(c) show that in addition to improving the accuracy of predictions, the proposed JS-G and the JS-A losses reduce the number of false negative predictions by 11.7% and 12.8% respectively, as compared to the KL loss. Given that the data set is biased towards the negative class, this is a significant achievement.

### 7.3. Comparison against baselines

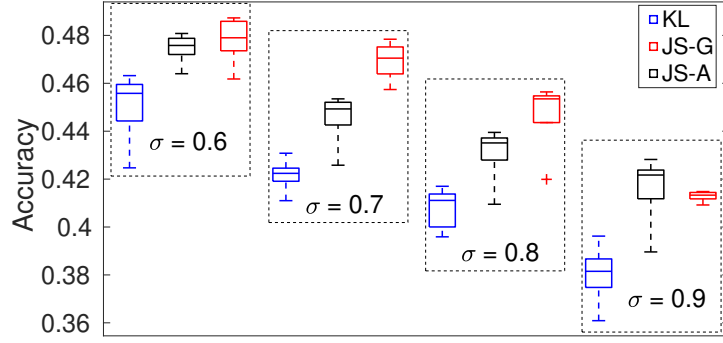
We performed additional experiments to evaluate the performance of the proposed loss functions against various baselines. The results of these experiments are presented here.

#### 7.3.1. Deterministic networks


Bayesian CNNs with the proposed loss functions are compared with a deterministic CNN of the same architecture and the results are presented in




(a)



(b)



(c)



(d)

Figure 4: Accuracy on (a-b) the CIFAR-10 test data at different noise levels (c) histopathology test data. Each box chart displays the median as the center line, the lower and upper quartiles as the box edges, and the minimum and maximum values as whiskers. (d) ROC curves, (e)-(g) Confusion matrices for histopathology data.




Figure 5: Confusion matrices for the histopathology dataset

Fig. 6. The deterministic CNN significantly overfits the training data and hence generalizes poorly for the validation dataset. Whereas, the Bayesian CNNs with their stochastic parameters can regularize much better than their deterministic counterparts. The validation accuracy of the deterministic CNN is about 8% lower than the proposed BNNs and 5% lower than the KL-based BNN. The difference between the training and validation accuracies are 63% and 86% less in KL- and JS-based BNNs respectively compared to the DNN, affirming greater generalization by KL- and JS-based BNNs.

### 7.3.2. $\lambda$ KL baseline

Implementing KL divergence in the constrained optimization framework (Eq. 19) yields a loss function which is henceforth called the  $\lambda$ KL baseline. We performed hyperparameter optimization on this baseline to tune the hyperparameters. The results of the  $\lambda$ KL baseline with the best-performing hyperparameters are presented in Fig. 7. We found that the maximum validation accuracy is 40.12% for this baseline on the CIFAR-10 images with  $\sigma=0.9$ . This validation accuracy is higher than the ELBO loss, which is 39.72 %, which corresponds to  $\lambda=1$ . However, it is still less than the proposed JS loss, which has a maximum validation accuracy of 41.78%. In addition, for the  $\lambda$ KL baseline, the regularization performance improves with  $\lambda=68.9$  although there are broad regions (between 20 to 90 epochs) where the network overfits.

The improved performance of JS divergences is a result of optimal penalization when the posterior is away from the prior. The functional form of the regularization term is now changed since the JS divergence is a weighted average of forward and reverse KL divergences. Through the  $\alpha$  term of JS divergence, we can choose the optimal weights between the forward and reverse KL which is not present in the standard KL divergence. This

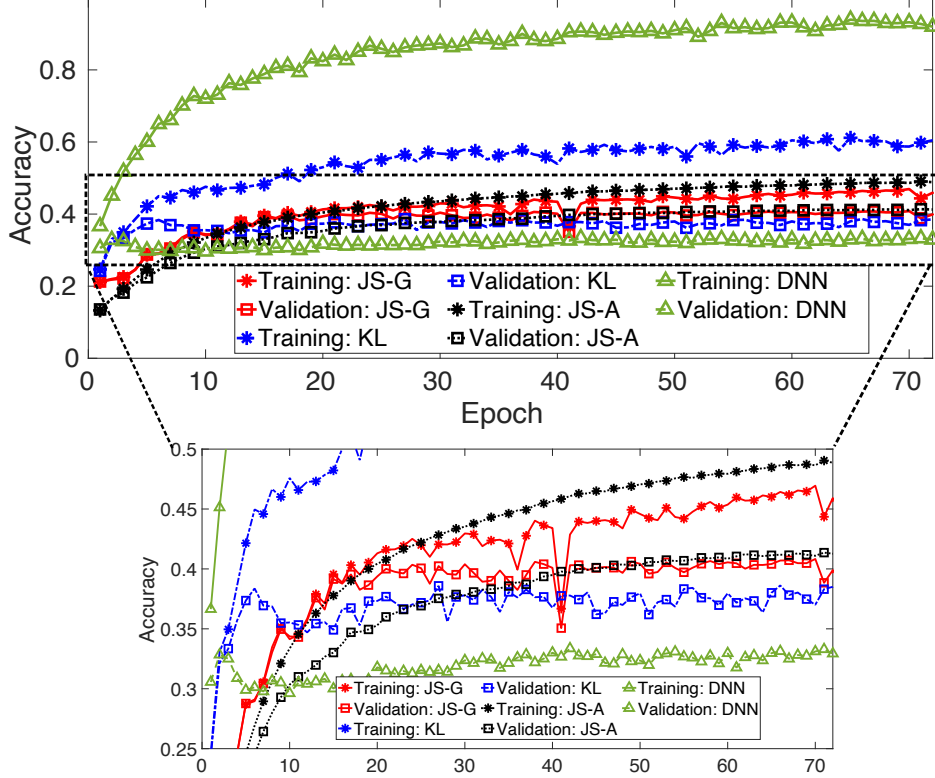


Figure 6: Performance of a deterministic CNN on CIFAR-10 dataset with added Gaussian noise  $\mathcal{N}(\mu = 0, \sigma = 0.9)$  and its magnified view.

combination of forward and reverse KL divergence allows us to alter the shape of the multi-dimensional regularization term by adapting to the data, which is not possible to achieve by scalar multiplication of the regularization term.

### 7.3.3. Unmodified JS divergences

The unmodified JS divergences in Eq. 1 and Eq. 2 fail to capture the dissimilarity between two distributions in the limiting cases of  $\alpha$  as explained in the main text. However, we implemented these unmodified JS divergences in the loss function in Eq. 19 and compared their performance with the modified versions for completeness. From Fig. 8 it is evident that the modified JS divergences-based losses outperform the unmodified JS divergences-based losses on the histopathology dataset. The validation accuracy is less by about 3% and 2% for the unmodified JS divergences in Eq. 1 and Eq. 2 as compared

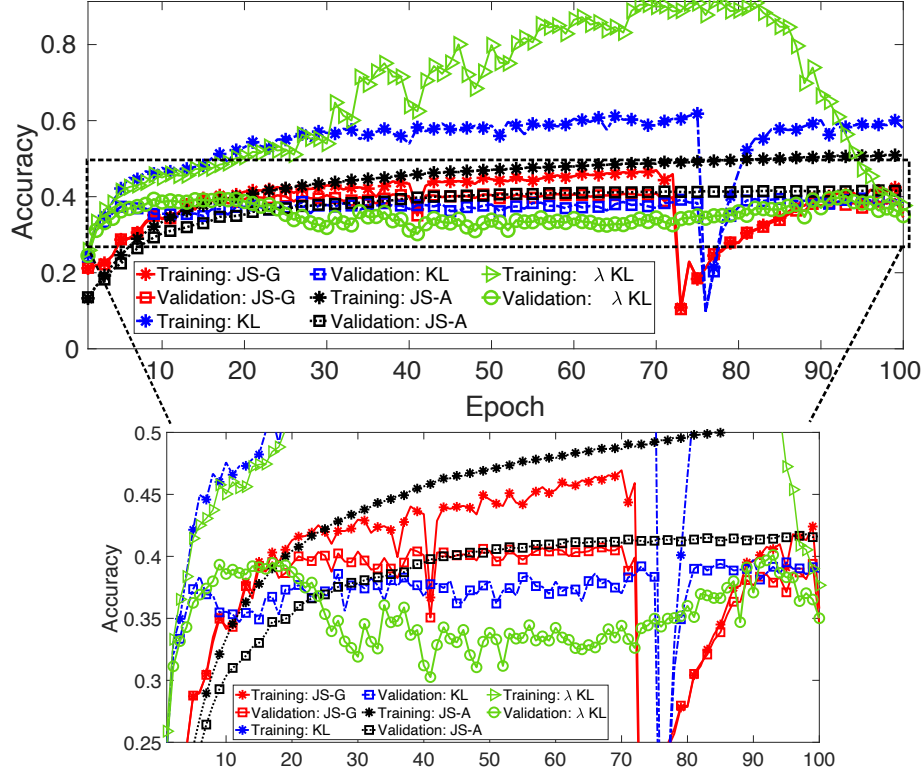


Figure 7: Performance of the  $\lambda$  KL baseline for  $\lambda = 68.9$  with a magnified view.

to the modified ones respectively.

#### 7.4. CIFAR-100 dataset

The CIFAR-100 data set [33] consists of 60,000 images of size  $32 \times 32 \times 3$  belonging to 100 mutually exclusive classes. Experiments on this dataset were carried out to evaluate the performance of the loss functions. The images were normalized using the min-max normalization technique and the training data set was split into 80% 20% for training and validation respectively. For this dataset, a ResNet18-V1 type architecture without the batch norm layers was used where the first two layers (convolution and max pooling layer) are replaced with a single convolution layer with  $3 \times 3$  kernel and  $1 \times 1$  stride.

The results of the experiment are presented in Fig. 9. The test accuracies of KL, JS-G, and JS-A divergence-based losses were 22.81 %, 24.51 %, and



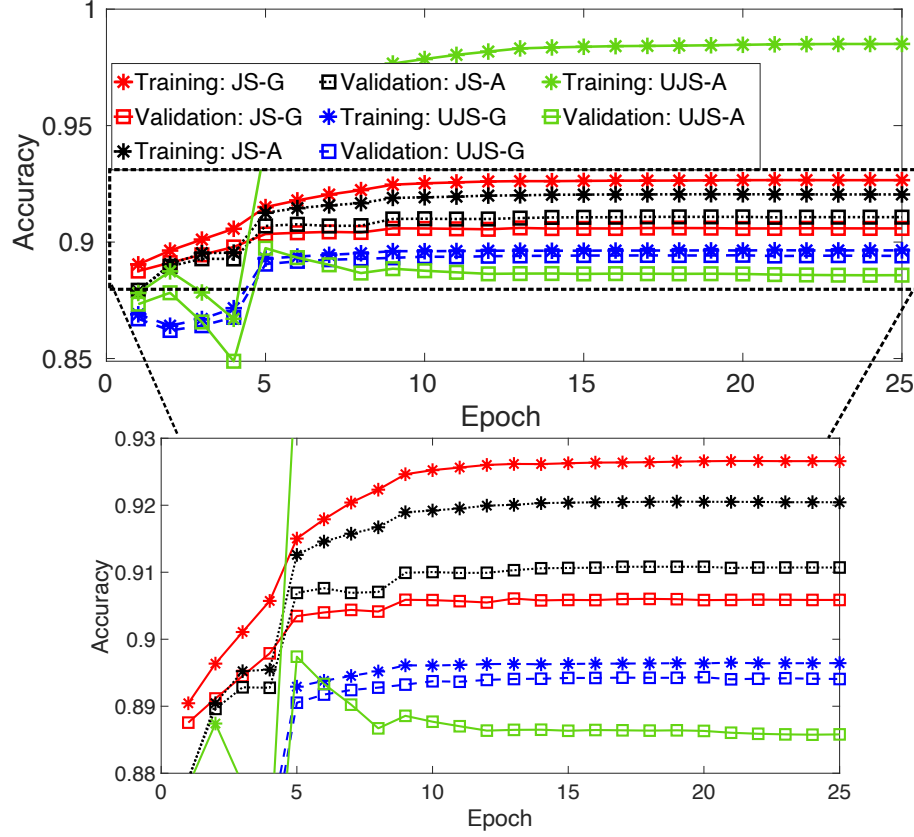


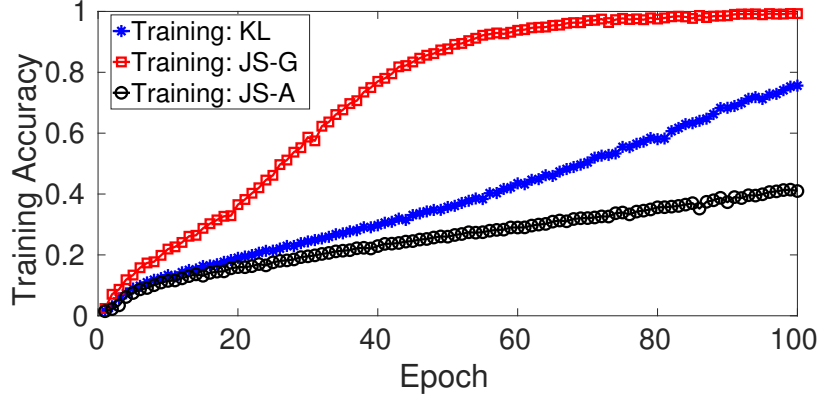
Figure 8: Performance of unmodified divergences (UJS) on histopathology with a magnified view.

24.02 % respectively. Both the proposed JS divergence-based losses perform better than the KL loss in terms of test and validation accuracies. The regularization performance of the JS-A divergence was better than KL for this noise level. However, in terms of regularization the KL divergence performs better than the JS-G divergence for this dataset at the given noise level.

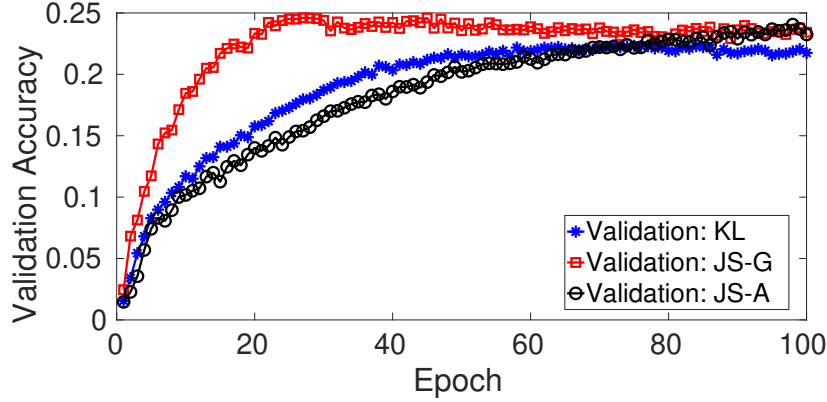
### 7.5. Computational cost

All the convolutional neural networks presented in the paper were built on Python 3 using Apache MXNet with CUDA 10.2. Training, validation, and testing of all the networks were performed using the Nvidia Tesla V100 32GB GPUs.

A comparison of the computational time per epoch during training of the



(a)



(b)

Figure 9: Performance comparison of the KL, JS-G, and JS-A divergence-based loss functions with CIFAR-100 dataset with added Gaussian noise  $\mathcal{N}(\mu = 0, \sigma = 0.5)$ .

histopathology dataset is provided for the three loss functions in Table. 1. The number of MC samples for the KL divergence and JS-A divergence-based loss functions was taken as 1. Thus, the computational effort is almost equal for the MC sampled KL divergence and the closed-form evaluated JS-G divergence. Whereas, to evaluate the JS-A divergence both the prior and the posterior distributions of the parameters need to be sampled. Due to this, there is an increased computational effort to evaluate the JS-A divergence-based loss function which is reflected in the increased computational time in Table. 1

Table 1: A comparison of computational time for the three loss functions.

Divergence	Training Time per epoch (s)
KL	1140
JS-G	1168
JS-A	1856

### 7.6. Comparison of numerical and theoretical results

The evolution of the divergence and negative log-likelihood part of the loss function during optimization is shown in Fig. 10 for the histopathology dataset. It is to be noted that the values are normalized by the number of test samples. The results are aligned with the theoretical results that the JS-G divergence penalizes higher than the KL divergence when the distribution  $q$  is farther from  $p$ . It is also seen that the negative log-likelihood is higher for the JS-G divergence as compared to the KL divergence in training. However, for the test dataset, the negative log-likelihood is lower for the JS divergences than the KL divergence which is desirable.




Figure 10: Comparison of the evolution of training loss for the histopathology dataset

### 7.7. Comparing uncertainty metrics

In this section, we compare the Test NLL, Test loss, and expected calibration error for the three losses for the histopathology dataset. From Table. 2

It is observed that the proposed JS divergences perform better in terms of both ECE and negative log-likelihood than the KL divergence for the histopathology dataset. Note that the test loss and test NLL are normalized by the number of test samples. The JS-A loss is significantly higher because of the value of  $\lambda = 100$ .

Table 2: Performance comparison for the three loss functions.

<b>Divergence</b>	<b>Test NLL</b>	<b>Test loss</b>	<b>ECE</b>
KL	0.0810	177.6	0.0323
JS-G	0.0706	201.2	0.0158
JS-A	0.0689	15151.3	0.0091

## 8. Results on regression

The setup for the regression experiments follows [29, 28] in general. 20 trials are performed for each of the six regression datasets and the average test root mean squared error and average negative log-likelihood for the JS divergences are compared in Table 3 and 4 with various other divergence proposed for variational inference of BNNs in the literature [29]. Since these datasets and the architecture do not demand regularization, the proposed JS divergence-based loss functions perform as good as the state-of-the-art methods or in some cases better. Competitive performances on these datasets demonstrate the versatility of the loss proposed loss functions. From these experiments, it is evident that the proposed loss functions can perform better than the state-of-the-start in applications requiring regularization and can perform as good as the state-of-the-art in cases where the need for regularization is absent.

## 9. Limitations and Future work

The proposed losses have two additional hyperparameters ( $\alpha$  and  $\lambda$ ) that need to be optimized to realize their full potential in regularization performance, which increases computational expenses. This might limit the application of the proposed methods to tasks involving excessive computational cost, either due to large datasets or extensive models. However, when

Table 3: Average root mean squared error. Except JS-G and JS-A all other values are taken from [29]

<b>Data</b>	JS-G	JS-A	KL-VI	$\chi$ -VI
Air	2.22 $\pm$ .25	2.32 $\pm$ .19	<b>2.16<math>\pm</math>.07</b>	2.36 $\pm$ .14
Aq	1.63 $\pm$ .19	1.13 $\pm$ .13	<b>1.12<math>\pm</math>.06</b>	1.20 $\pm$ .06
Bos	3.34 $\pm$ .88	2.91 $\pm$ 0.73	<b>2.76<math>\pm</math>.36</b>	2.99 $\pm$ .37
Con	5.10 $\pm$ .67	4.88 $\pm$ .63	5.40 $\pm$ .24	<b>3.32<math>\pm</math>.34</b>
R.E	<b>6.77<math>\pm</math>1.1</b>	7.37 $\pm$ 2.3	7.48 $\pm$ 1.4	7.51 $\pm$ 1.4
Ya	0.82 $\pm$ .35	<b>0.69<math>\pm</math>.25</b>	0.78 $\pm$ .12	1.18 $\pm$ .18
<b>Data</b>	$\alpha$ -Vi	TV-VI	$f_{c1}$ -VI	$f_{c2}$ -VI
Air	2.30 $\pm$ .08	2.47 $\pm$ .15	2.34 $\pm$ .09	2.16 $\pm$ .09
Aq	1.14 $\pm$ .07	1.23 $\pm$ .10	1.14 $\pm$ .06	1.14 $\pm$ .06
Bos	2.86 $\pm$ .36	2.96 $\pm$ .36	2.87 $\pm$ .36	2.89 $\pm$ .38
Con	5.32 $\pm$ .27	5.27 $\pm$ .24	5.26 $\pm$ .21	5.32 $\pm$ .24
R.E	7.46 $\pm$ 1.4	8.02 $\pm$ 1.6	7.52 $\pm$ 1.4	7.99 $\pm$ 1.6
Ya	0.99 $\pm$ .12	1.03 $\pm$ .14	1.00 $\pm$ .18	0.82 $\pm$ .16

Table 4: Average negative log-likelihood. Except JS-G and JS-A all other values are taken from [29]

<b>Data</b>	JS-G	JS-A	KL-VI	$\chi$ -VI
Air	2.22 $\pm$ .09	2.72 $\pm$ .10	<b>2.17<math>\pm</math>.03</b>	2.27 $\pm$ .03
Aq	1.94 $\pm$ .13	1.78 $\pm$ .23	<b>1.54<math>\pm</math>.04</b>	1.60 $\pm$ .08
Bos	2.69 $\pm$ .31	3.3 $\pm$ 1.12	2.49 $\pm$ .08	2.54 $\pm$ .18
Con	3.1 $\pm$ .17	3.11 $\pm$ .20	3.10 $\pm$ .04	<b>2.61<math>\pm</math>.18</b>
R.E	<b>3.49<math>\pm</math>.06</b>	3.56 $\pm$ .20	3.60 $\pm$ .30	3.70 $\pm$ .45
Ya	1.51 $\pm$ .15	<b>1.43<math>\pm</math>0.09</b>	1.70 $\pm$ .02	1.79 $\pm$ .03
<b>Data</b>	$\alpha$ -Vi	TV-VI	$f_{c1}$ -VI	$f_{c2}$ -VI
Air	2.26 $\pm$ .02	2.28 $\pm$ .04	2.29 $\pm$ .02	2.18 $\pm$ .03
Aq	1.54 $\pm$ .07	1.56 $\pm$ .07	1.54 $\pm$ .06	1.55 $\pm$ .04
Bos	<b>2.48<math>\pm</math>.13</b>	2.51 $\pm$ .18	2.49 $\pm$ .13	2.51 $\pm$ .10
Con	3.09 $\pm$ .04	3.10 $\pm$ .05	3.09 $\pm$ .03	3.10 $\pm$ .04
R.E	3.59 $\pm$ .32	3.86 $\pm$ .52	3.62 $\pm$ .33	3.74 $\pm$ .37
Ya	1.82 $\pm$ .01	1.78 $\pm$ .02	2.05 $\pm$ .01	1.86 $\pm$ .02

such expenses can not be afforded, the parameters can be set to the fixed values  $\alpha = 0$  and  $\lambda = 1$  to recover the KL loss. As a future direction, the proposed loss functions can be used in an application involving adversarial training to see the effectiveness of the method where a high degree of noise is involved. Further, the uncertainty quantification aspects of the proposed loss functions can be extensively analyzed to compare them against the existing uncertainty estimates of machine learning models' predictions.

## 10. Conclusions


We summarize the main findings of this work in the following. *Firstly*, the proposed loss function using a novel bounded JS-A divergence resolves the issue of unstable optimization associated with KL divergence-based loss. The proposed loss encompasses the KL divergence-based loss and extends it to a wider class of bounded divergences. *Secondly*, we also introduced a loss function based on the JS-G divergence that can be evaluated analytically and hence is computationally efficient. *Thirdly*, better regularization performance by the two proposed loss functions compared to the state-of-the-art is established analytically and/or numerically. *Fourthly*, empirical experiments on the histopathology data sets having bias show about a 13% reduction in false negative predictions, and the CIFAR-10 dataset with various degrees of added noise shows a 5% improvement in accuracy over the state-of-the-art with the proposed losses. Similarly, experiments on standard regression datasets show an 8% reduction in the root mean square error, in the best-case scenario, when using the proposed losses as compared to the KL loss.

## Acknowledgement

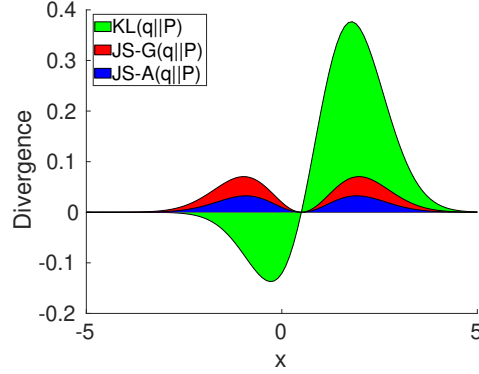
This work was supported by grant DE-SC0023432 funded by the U.S. Department of Energy, Office of Science. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, using NERSC awards BES-ERCAP0025205 and BES-ERCAP0025168. The authors thank the Superior HPC facility at MTU, and the Applied Computing and MRI GPU clusters at MTU for making available some of the computing resources used in this work. PT acknowledges the Doctoral Finishing Fellowship awarded by the Graduate School at Michigan Technological University.

## Appendix A. Divergences

Two distributions,  $P(x) = \mathcal{N}(0, 1)$  and  $q(x) = \mathcal{N}(1, 1)$  are shown in Fig. A.11 along with the divergence as a function of  $x$ . The area under the curve shown as a shaded region is the area to be integrated to obtain the divergence. We provide a comparison of symmetry (at  $\alpha = 0.5$ ) and



(a) Distributions  $q$  and  $P$



(b) Divergence shown as shaded regions.

Figure A.11: Depiction of KL, JS-G, and JS-A divergences for two Gaussian distributions. The area under the curves is the value of the divergences.

boundedness for divergences used in this work in Table A.5

Table A.5: Properties of various divergences

Div.	Bounded	Symmetric	Analytical exp.
KL	×	×	✓
JS-A	✓	✓	×
JS-G	×	✓	✓

## Appendix B. Proof of Theorem 1

**Theorem 1:** Boundedness of the modified generalized JS divergence

For any two distributions  $P_1(t)$  and  $P_2(t)$ ,  $t \in \Omega$ , the value of the divergence JS-A is bounded such that,

$$\text{JS-A}(P_1(t)||P_2(t)) \leq -(1 - \alpha) \log \alpha - \alpha \log(1 - \alpha), \quad \text{for } \alpha \in (0, 1)$$

**Proof:**

$$\begin{aligned}
\text{JS-A}(P_1||P_2) &= (1 - \alpha) \int_{\Omega} P_1 \log \frac{P_1}{A'_\alpha} dt + \alpha \int_{\Omega} P_2 \log \frac{P_2}{A'_\alpha} dt \\
&= \int_{\Omega} A'_\alpha \left[ (1 - \alpha) \frac{P_1}{A'_\alpha} \log \frac{P_1}{A'_\alpha} + \alpha \frac{P_2}{A'_\alpha} \log \frac{P_2}{A'_\alpha} \right] dt \\
&= \int_{\Omega} A'_\alpha \left[ \frac{(1 - \alpha) \alpha P_1}{\alpha A'_\alpha} \left( \log \frac{\alpha P_1}{A'_\alpha} - \log \alpha \right) \right. \\
&\quad \left. + \frac{\alpha}{(1 - \alpha)} \frac{(1 - \alpha) P_2}{A'_\alpha} \left( \log \frac{(1 - \alpha) P_2}{A'_\alpha} - \log(1 - \alpha) \right) \right] dt \\
&= \int_{\Omega} -(1 - \alpha) P_1 \log \alpha dt - \int_{\Omega} \alpha P_2 \log(1 - \alpha) dt - \\
&\quad \int_{\Omega} A'_\alpha \left[ \frac{(1 - \alpha)}{\alpha} H \left( \frac{\alpha P_1}{A'_\alpha} \right) + \frac{\alpha}{(1 - \alpha)} H \left( \frac{(1 - \alpha) P_2}{A'_\alpha} \right) \right] dt \\
&= -(1 - \alpha) \log \alpha - \alpha \log(1 - \alpha) - \mathcal{H} \\
&\leq -(1 - \alpha) \log \alpha - \alpha \log(1 - \alpha)
\end{aligned}$$

Where,  $H(f(t)) = -f(t) \log f(t)$  and

$$\mathcal{H} = \int_{\Omega} A'_\alpha \left[ \frac{(1 - \alpha)}{\alpha} H \left( \frac{\alpha P_1}{A'_\alpha} \right) + \frac{\alpha}{(1 - \alpha)} H \left( \frac{(1 - \alpha) P_2}{A'_\alpha} \right) \right] dt$$

Note,  $H(f(t)) \geq 0$ ;  $\forall f(t) \in [0, 1]$ . Therefore,  $\mathcal{H} \geq 0$ ;  $\forall t \in \Omega$



The unboundedness of the KL and the JS-G divergences is depicted through a contrived example in Fig. B.12. The distribution  $q$ , where  $q = \mathcal{N}(0, \sigma)$ , is assumed to be a Gaussian with varying  $\sigma$ . The distribution  $P$ , where  $P = \mathcal{U}(-5, 5)$ , is assumed to be a uniform distribution with support  $(-5, 5)$ . The distributions  $q$  and  $p$  are shown in Fig. 12(a). The KL, the JS-A, and the JS-G divergence of the two distributions  $q$  and  $P$  are evaluated using 50,000 Monte Carlo samples each. The value of the KL and the JS-G divergence explodes to infinity when the distribution  $p$  is zero for a non-zero  $q$  due to the effect of  $\log(q/P)$ . In contrast, the JS-G divergence is always bounded. This can be seen in Fig. 12(b).

Numerical examples depicting Theorem-1 are shown in Fig. B.13. Two distribution  $q = \mathcal{N}(\mu, 1)$  and  $P = \mathcal{N}(0, 1)$  are considered. The JS-A divergence is evaluated for varying  $\mu$ . The value of the JS-A divergence increases with increasing  $\mu$  until the upper bound is reached and it remains constant henceforth as seen in Fig. B.13.

### Appendix C. Intractability of JS-based VI


The JS-G loss is given as

$$\begin{aligned}\mathcal{F}_{JS-G}(\mathbb{D}, \boldsymbol{\theta}) &= \text{JS-G} [q(\mathbf{w}|\boldsymbol{\theta}) \parallel P(\mathbf{w}|\mathbb{D})] \\ &= (1 - \alpha)\text{KL} (q \parallel G'_\alpha(q, P)) + \alpha\text{KL} (P \parallel G'_\alpha(q, P))\end{aligned}\tag{C.1}$$


Where,  $G'_\alpha(q, P) = q(\mathbf{w}|\boldsymbol{\theta})^\alpha P(\mathbf{w}|\mathbb{D})^{(1-\alpha)}$ .

Rewriting the first term in Eq. C.1 as,

$$\begin{aligned}T_1 &= (1 - \alpha)\text{KL} (q \parallel G'_\alpha(q, P)) \\ &= (1 - \alpha)KL[q(\mathbf{w}|\boldsymbol{\theta}) \parallel q(\mathbf{w}|\boldsymbol{\theta})^\alpha P(\mathbf{w}|\mathbb{D})^{(1-\alpha)}] \\ &= (1 - \alpha) \int q(\mathbf{w}|\boldsymbol{\theta}) \log \left[ \frac{q(\mathbf{w}|\boldsymbol{\theta})}{q(\mathbf{w}|\boldsymbol{\theta})^\alpha P(\mathbf{w}|\mathbb{D})^{1-\alpha}} \right] d\mathbf{w} \\ &= (1 - \alpha)^2 \int q(\mathbf{w}|\boldsymbol{\theta}) \log \left[ \frac{q(\mathbf{w}|\boldsymbol{\theta})}{P(\mathbf{w}|\mathbb{D})} \right] d\mathbf{w}\end{aligned}\tag{C.2}$$



(a) Distributions  $q$  and  $P$



(b) Divergence values for various  $\sigma$

Figure B.12: Depiction of the unboundedness of the KL and JS-G divergence and the boundedness of the JS-A divergence.




Figure B.13: Upper bound of the JS-A divergence for various values of  $\alpha$ .

Similarly rewriting the second term in Eq. C.1 as,

$$\begin{aligned}
T_2 &= \alpha \text{KL}(P \parallel G'_\alpha(q, P)) \\
&= \alpha \text{KL}[P(\mathbf{w}|\mathbb{D}) \parallel q(\mathbf{w}|\boldsymbol{\theta})^\alpha P(\mathbf{w}|\mathbb{D})^{(1-\alpha)}] \\
&= \alpha \int P(\mathbf{w}|\mathbb{D}) \log \left[ \frac{P(\mathbf{w}|\mathbb{D})}{q(\mathbf{w}|\boldsymbol{\theta})^\alpha P(\mathbf{w}|\mathbb{D})^{(1-\alpha)}} \right] d\mathbf{w} \\
&= \alpha^2 \int P(\mathbf{w}|\mathbb{D}) \log \left[ \frac{P(\mathbf{w}|\mathbb{D})}{q(\mathbf{w}|\boldsymbol{\theta})} \right] d\mathbf{w} \tag{C.3}
\end{aligned}$$

The term  $T_2$  is intractable due to the posterior distribution  $P(\mathbf{w}|\mathbb{D})$ .

#### Appendix D. Insights into Monte Carlo estimates

A closed-form solution does not exist for KL and JS divergences for most distributions. In cases where such a closed form for the divergence is not available for a given distribution, we resort to Monte Carlo (MC) methods. However, the estimation of the loss function using MC methods is computationally more expensive than the closed-form evaluation as shown in D.14. In addition, for networks with a large number of parameters, the memory requirement increases significantly with the number of MC samples. Therefore, utilizing the closed-form solution when available can save huge computational efforts and memory.

To estimate the number of MC samples required to achieve a similar level of accuracy of the closed-form expression, JS-G divergence of two Gaussian distributions  $\mathcal{N}(5, 1)$  and  $\mathcal{N}(0, 1)$  are evaluated and compared with its closed form counterpart. Fig. D.14 shows the results of the comparison. It is seen that at least 600 samples are required to estimate the JS-G divergence within 5% error. This implies evaluating the loss function 600 times for a given input and back-propagating the error which requires huge computational efforts.




Figure D.14: Comparison of MC estimates and the closed form solution of JS-G divergence demonstrating the benefit of closed form solution.

## Appendix E. Detailed proof of Theorem 2

**Theorem 2.** *For any two arbitrary distributions  $P$  and  $q$  such that  $P \neq q$ ,  $\tilde{\mathcal{F}}_{JSG} > \mathcal{F}_{KL}$  if and only if  $\alpha > \frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} \in (0, \infty)$*

**Proof:**

Assuming,

$$\tilde{\mathcal{F}}_{JSG} - \mathcal{F}_{KL} > 0$$

From Eq.(5) and Eq.(18) from the main text we have,

$$\begin{aligned} (1 - \alpha)^2 \text{KL}(q||P) + \alpha^2 \text{KL}(P||q) - \text{KL}(q||P) &> 0 \\ (\alpha^2 - 2\alpha) \text{KL}(q||P) + \alpha^2 \text{KL}(P||q) &> 0 \end{aligned}$$

This leads to,

$$\alpha > \frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)}$$

This proves that if  $\tilde{\mathcal{F}}_{JSG} > \mathcal{F}_{KL}$  then  $\alpha > \frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)}$ .

Conversely, assuming  $\alpha > r$ , where  $r = \frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)}$ , we get

$$\begin{aligned} \alpha [\text{KL}(q||P) + \text{KL}(P||q)] - 2 \text{KL}(q||P) &> \\ r [\text{KL}(q||P) + \text{KL}(P||q)] - 2 \text{KL}(q||P) &\quad (\text{since } \alpha > r) \end{aligned}$$

Substituting  $r$  on the right hand side,

$$\begin{aligned} \alpha [\text{KL}(q||P) + \text{KL}(P||q)] - 2 \text{KL}(q||P) &> 0 \\ \alpha^2 [\text{KL}(q||P) + \text{KL}(P||q)] - 2\alpha \text{KL}(q||P) &> 0 \quad (\text{since } \alpha > 0) \\ (\alpha^2 - 2\alpha)\text{KL}(q||P) + \alpha^2\text{KL}(P||q) &> 0 \\ (1 - \alpha)^2\text{KL}(q||P) + \alpha^2\text{KL}(P||q) - \text{KL}(q||P) &> 0 \\ \tilde{\mathcal{F}}_{JSG} - \mathcal{F}_{KL} &> 0 \end{aligned}$$

This proves that if  $\alpha > \frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)}$  then  $\tilde{\mathcal{F}}_{JSG} > \mathcal{F}_{KL}$ . Hence the theorem is proved.

Note that, since KL divergence is always non-negative for any distributions and  $P \neq q$ , we have  $\frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} > 0$

## Appendix F. Detailed proof of Theorem 3

### Theorem 3.

If  $P = \mathcal{N}(\mu_p, \sigma_p^2)$  and  $q = \mathcal{N}(\mu_q, \sigma_q^2)$  are Gaussian distributions and  $P \neq q$ , then  $\frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} < 1$  if and only if  $\sigma_p^2 > \sigma_q^2$ .

**Proof:**

Assuming,

$$\frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} < 1$$

We have,

$$\text{KL}(P||q) > \text{KL}(q||P) \quad (\text{F.1})$$

Since  $P$  and  $q$  be Gaussian distributions with  $P = \mathcal{N}(\mu_p, \sigma_p^2)$  and  $q = \mathcal{N}(\mu_q, \sigma_q^2)$ , Eq. F.1 can be written as,

$$\begin{aligned} \ln \frac{\sigma_q^2}{\sigma_p^2} + \frac{\sigma_p^2 + (\mu_q - \mu_p)^2}{\sigma_q^2} - 1 &> \ln \frac{\sigma_p^2}{\sigma_q^2} + \frac{\sigma_q^2 + (\mu_p - \mu_q)^2}{\sigma_p^2} - 1 \\ \frac{\sigma_p^2}{\sigma_q^2} + \ln \frac{\sigma_q^2}{\sigma_p^2} + \frac{(\mu_q - \mu_p)^2}{\sigma_q^2} - \frac{\sigma_q^2}{\sigma_p^2} - \ln \frac{\sigma_p^2}{\sigma_q^2} - \frac{(\mu_p - \mu_q)^2}{\sigma_p^2} &> 0 \end{aligned}$$

Denoting  $\gamma = \frac{\sigma_p^2}{\sigma_q^2}$ , we get,

$$\begin{aligned} \gamma - \frac{1}{\gamma} + \ln \frac{1}{\gamma} - \ln \gamma + \frac{(\mu_q - \mu_p)^2}{\sigma_q^2} - \frac{(\mu_p - \mu_q)^2}{\gamma \sigma_q^2} &> 0 \\ \gamma - \frac{1}{\gamma} + \ln \frac{1}{\gamma^2} + \frac{(\mu_q - \mu_p)^2}{\sigma_q^2} \left(1 - \frac{1}{\gamma}\right) &> 0 \\ \ln \left[ \exp \left( \gamma - \frac{1}{\gamma} \right) \right] + \ln \frac{1}{\gamma^2} + \frac{(\mu_q - \mu_p)^2}{\sigma_q^2} \left(1 - \frac{1}{\gamma}\right) &> 0 \end{aligned}$$

or,

$$\ln \left[ \frac{1}{\gamma^2} \exp \left( \gamma - \frac{1}{\gamma} \right) \right] + \frac{(\mu_q - \mu_p)^2}{\sigma_q^2} \left(1 - \frac{1}{\gamma}\right) > 0 \quad (\text{F.2})$$

The second term of Eq. F.2 is greater than 0 only when  $\gamma > 1$ . Consider the first term,

$$\ln \left[ \frac{1}{\gamma^2} \exp \left( \gamma - \frac{1}{\gamma} \right) \right] > 0$$

or,

$$\frac{1}{\gamma^2} \exp \left( \gamma - \frac{1}{\gamma} \right) > 1$$

$\frac{1}{\gamma^2} \exp\left(\gamma - \frac{1}{\gamma}\right) = 1$  for  $\gamma = 1$  and it is a monotonically increasing function for  $\gamma > 1$ . This can be seen from its positive slope

$$\begin{aligned} \frac{d}{d\gamma} \left[ \frac{1}{\gamma^2} \exp\left(\gamma - \frac{1}{\gamma}\right) \right] &= \frac{1}{\gamma^2} \exp\left(\gamma - \frac{1}{\gamma}\right) \times \left(1 + \frac{1}{\gamma^2}\right) - \frac{2}{\gamma^3} \exp\left(\gamma - \frac{1}{\gamma}\right) \\ &= \frac{1}{\gamma^4} \exp\left(\gamma - \frac{1}{\gamma}\right) [\gamma^2 + 1 - 2\gamma] \\ &= \frac{1}{\gamma^4} \exp\left(\gamma - \frac{1}{\gamma}\right) (\gamma - 1)^2 \\ &> 0 \quad \text{for } \gamma \neq 1 \end{aligned}$$

Therefore, the first term of Eq. F.2 is greater than 0 only when  $\gamma > 1$ . Thus, the condition in Eq. F.2 is satisfied only when  $\gamma > 1$ , which implies

$$\sigma_p^2 > \sigma_q^2 \quad (\text{F.3})$$

Thus if  $\frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} < 1$  then  $\sigma_p^2 > \sigma_q^2$ .

Conversely, assuming  $\sigma_p^2 > \sigma_q^2$ , i.e.  $\gamma > 1$  consider,

$$\ln \left[ \frac{1}{\gamma^2} \exp\left(\gamma - \frac{1}{\gamma}\right) \right] + \frac{(\mu_q - \mu_p)^2}{\sigma_q^2} \left(1 - \frac{1}{\gamma}\right) > 0$$

Which leads to,

$$\text{KL}(P||q) > \text{KL}(q||P) \quad (\text{following the steps as above})$$

or,

$$\begin{aligned} \text{KL}(P||q) + \text{KL}(q||P) &> \text{KL}(q||P) + \text{KL}(q||P) \\ \frac{2\text{KL}(q||P)}{\text{KL}(P||q) + \text{KL}(q||P)} &< 1 \end{aligned}$$

Thus if  $\sigma_p^2 > \sigma_q^2$  then  $\frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} < 1$ . Hence the theorem is proved.

**Corollary:** From Theorem 2 and 3:  $\tilde{\mathcal{F}}_{JSG} > \mathcal{F}_{KL}$  if  $\sigma_p^2 > \sigma_q^2$  and  $\forall \alpha \in (0, 1]$  such that  $\alpha > \frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)}$ . Where,  $P$  and  $q$  are Gaussians and  $P \neq q$ .

Fig. F.15 shows the sign of  $\tilde{\mathcal{F}}_{JSG} - \mathcal{F}_{KL}$  for various values of  $\alpha$  and  $\sigma_p - \sigma_q$ . Two Gaussian distributions  $P = \mathcal{N}(0, 1)$  and  $q = \mathcal{N}(0, \sigma_q^2)$  are considered for this purpose. It is evident that at least one value of  $\alpha$  exists such that  $\tilde{\mathcal{F}}_{JSG} - \mathcal{F}_{KL} > 0$  when  $\sigma_p > \sigma_q$ . In addition, for  $\alpha = 0$  or  $\sigma_p = \sigma_q$ ,  $\tilde{\mathcal{F}}_{JSG} = \mathcal{F}_{KL}$ .

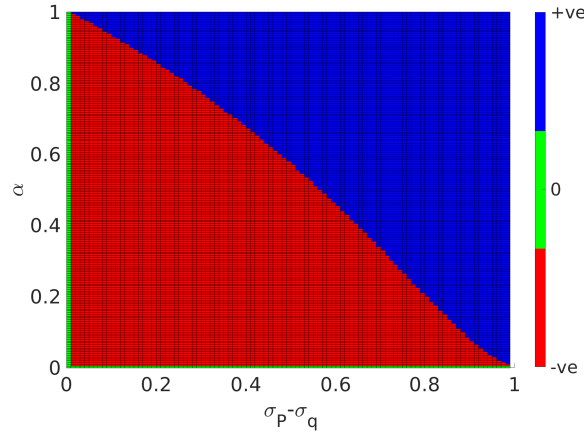


Figure F.15: The sign of  $\tilde{\mathcal{F}}_{JSG} - \mathcal{F}_{KL}$  is plotted for various values of  $\alpha$  and  $\sigma_p - \sigma_q$ .

## Appendix G. Data sets

### Appendix G.1. Regression datasets

**Airfoil:** This NASA dataset includes NACA 0012 airfoils of various sizes tested at different wind tunnel speeds and angles of attack. The goal is to predict the scaled sound pressure.

**Aquatic:** This dataset is used to predict quantitative acute aquatic toxicity towards *Daphnia Magna* from 8 molecular descriptors of 546 chemicals.

**Building:** This dataset includes construction cost, sale prices, project variables, and economic variables corresponding to real estate single-family residential apartments in Tehran, Iran.

**Concrete:** The dataset contains attributes to predict concrete compressive strength.

**Real Estate:** This dataset contains real estate valuation collected from Sindian Dist., New Taipei City, Taiwan.

**Wine:** This dataset is used to predict model wine quality based on physico-chemical tests.



### *Appendix G.2. CIFAR-10*

The CIFAR-10 data set [33] consists of 60,000 images of size  $32 \times 32 \times 3$  belonging to 10 mutually exclusive classes. This data set is unbiased, with each of the 10 classes having 6,000 images.

Images were resized to  $3 \times 224 \times 224$  pixels and normalized using the min-max normalization technique. The training data set was split into 80% 20% for training and validation respectively.

### *Appendix G.3. Histopathology*

The histopathology data set [34, 35] is publicly available under a CC0 license at [36]. The data set consists of images containing regions of Invasive Ductal Carcinoma. The original data set consisted of 162 whole-mount slide images of Breast Cancer specimens scanned at 40x. From the original whole slide images, 277,524 patches of size  $50 \times 50 \times 3$  pixels were extracted (198,738 negatives and 78,786 positives), labeled by pathologists, and provided as a data set for classification.

The data set consists of a positive (1) and a negative (0) class. 20% of the entire data set was used as the testing set for our study. The remaining 80% of the entire data was further split into a training set and a validation set (80%-20% split) to perform hyperparameter optimization. The images were shuffled and converted from uint8 to float format for normalizing. As a post-processing step, we computed the complement of all the images (training and testing) and then used them as inputs to the neural network. The images were resized to  $3 \times 224 \times 224$  pixel-wise normalization and complement were carried out as  $p_n = (255 - p)/255$ .  $p$  is the original pixel value and  $p_n$  is the pixel value after normalization and complement.

## **Appendix H. Results of hyperparameter optimization**

A Tree-structured Parzen Estimator (TPE) algorithm [37] is used for hyperparameter optimization which is a sequential model-based optimization approach. In this approach, models are constructed to approximate the performance of hyperparameters based on historical measurements. New hyperparameters are chosen based on this model to evaluate performance. A python library Hyperopt [40] is used to implement this optimization algorithm over a given search space.

Results of the hyperparameter optimization for the two classification data sets are presented in Tables. H.6 and H.7. In Table. H.6 Div denotes the

Table H.6: CIFAR-10 data set

Div	Parameter	Noise level ( $\sigma$ )				
		0.1	0.2	0.3	0.4	0.5
KL	LR	1e−4	1e−4	1e−4	1e−3	1e−3
JS-G	$\alpha$	0.004	0.1313	0.2855	0.3052	0.2637
	LR	1e−4	1e−4	1e−4	1e−4	1e−5
JS-A	$\lambda$	1000	1000	1000	1000	1000
	$\alpha$	0.7584	0.6324	0.1381	0.6286	0.1588
	LR	1e−4	1e−4	1e−4	1e−3	1e−4

Div	Parameter	Noise level ( $\sigma$ )			
		0.6	0.7	0.8	0.9
KL	LR	1e−3	1e−3	1e−3	1e−3
JS-G	$\alpha$	0.2249	0.3704	0.3893	0.7584
	LR	1e−4	1e−4	1e−5	1e−3
JS-A	$\lambda$	100	1000	1e4	1e5
	$\alpha$	0.4630	0.1220	0.2282	0.5792
	LR	1e−3	1e−4	1e−5	1e−5

divergence measure and LR is the learning rate.  $\lambda$  is taken as 1 for the JS-G divergence-based loss function throughout this work unless specified otherwise.

Table H.7: Histopathology data set

Divergence	$\alpha$	$\lambda$	Learning rate
KL	-	-	1e-4
JS-G	0.0838	1	1e-4
JS-A	0.0729	100	1e-4

## Author Biographies



**Ponkrshnan Thiagarajan** is a postdoctoral fellow at Hopkins Extreme Materials Institute, Johns Hopkins University. He received his Ph.D. in the Department of Mechanical Engineering-Engineering Mechanics at Michigan Technological University, Houghton, MI, USA. His research interests include Bayesian neural networks, Bayesian calibration, and Uncertainty quantification of computational models. He received his master's degree in aerospace engineering from the Indian Institute of Technology - Madras, and a bachelor's degree in aeronautical engineering from Anna University, Chennai.



**Susanta Ghosh** is an Assistant Professor, Mechanical Engineering-Engineering Mechanics, a Faculty member of Center for Data Sciences at the Institute of Computing and Cybersystems and a Faculty member of the Center for Applied Mathematics and Statistics at Michigan Technological University, Houghton, MI, USA. He has worked as a post-doctoral research scholar at Duke University, the University of Michigan, and The Technical University of Catalunya. His M.S. and Ph.D. degrees are from the Indian Institute of Science, Bangalore.

## References

- [1] S. Samarasinghe, Neural networks for applied sciences and engineering: from fundamentals to complex pattern recognition, Auerbach publications, 2016.
- [2] Z. Li, F. Liu, W. Yang, S. Peng, J. Zhou, A survey of convolutional neural networks: analysis, applications, and prospects, *IEEE transactions on neural networks and learning systems* (2021).
- [3] M. Buda, A. Maki, M. A. Mazurowski, A systematic study of the class imbalance problem in convolutional neural networks, *Neural Networks* 106 (2018) 249–259.
- [4] P. Thiagarajan, P. Khairnar, S. Ghosh, Explanation and use of uncertainty quantified by bayesian neural network classifiers for breast histopathology images, *IEEE Transactions on Medical Imaging* 41 (2021) 815–825.
- [5] H. D. Kabir, A. Khosravi, M. A. Hosen, S. Nahavandi, Neural network-based uncertainty quantification: A survey of methodologies and applications, *IEEE access* 6 (2018) 36218–36234.
- [6] L. V. Jospin, H. Laga, F. Boussaid, W. Buntine, M. Bennamoun, Hands-on bayesian neural networks—a tutorial for deep learning users, *IEEE Computational Intelligence Magazine* 17 (2022) 29–48.
- [7] L. Bortolussi, G. Carbone, L. Laurenti, A. Patane, G. Sanguinetti, M. Wicker, On the robustness of bayesian neural networks to adversarial attacks, *IEEE Transactions on Neural Networks and Learning Systems* (2024) 1–14. doi:10.1109/TNNLS.2024.3386642.
- [8] P. Lemos, M. Cranmer, M. Abidi, C. Hahn, M. Eickenberg, E. Massara, D. Yallup, S. Ho, Robust simulation-based inference in cosmology with bayesian neural networks, *Machine Learning: Science and Technology* 4 (2023) 01LT01.
- [9] X.-X. Dong, R. An, J.-X. Lu, L.-S. Geng, Nuclear charge radii in bayesian neural networks revisited, *Physics Letters B* 838 (2023) 137726.

- [10] S. Pathrudkar, P. Thiagarajan, S. Agarwal, A. S. Banerjee, S. Ghosh, Electronic structure prediction of multi-million atom systems through uncertainty quantification enabled transfer learning, *arXiv preprint arXiv:2308.13096* (2023).
- [11] K. Luo, J. Zhu, Z. Li, H. Zhu, Y. Li, R. Hu, T. Fan, X. Chang, L. Zhuang, Z. Yang, Ultrasonic lamb wave damage detection of cfrp composites using the bayesian neural network, *Journal of Nondestructive Evaluation* 43 (2024) 48.
- [12] N. Tishby, E. Levin, S. A. Solla, Consistent inference of probabilities in layered networks: Predictions and generalization, in: *International Joint Conference on Neural Networks*, volume 2, IEEE New York, 1989, pp. 403–409.
- [13] J. Denker, Y. LeCun, Transforming neural-net output levels to probability distributions, *Advances in neural information processing systems* 3 (1990).
- [14] E. Goan, C. Fookes, *Bayesian Neural Networks: An Introduction and Survey*, Springer International Publishing, Cham, 2020, pp. 45–87.
- [15] Y. Gal, *Uncertainty in deep learning*, PhD thesis, University of Cambridge (2016).
- [16] M. Magris, A. Iosifidis, Bayesian learning for neural networks: an algorithmic survey, *Artificial Intelligence Review* 56 (2023) 11773–11823.
- [17] G. E. Hinton, D. Van Camp, Keeping the neural networks simple by minimizing the description length of the weights, in: *Proceedings of the sixth annual conference on Computational learning theory*, 1993, pp. 5–13.
- [18] D. Barber, C. M. Bishop, Ensemble learning in bayesian neural networks, *Nato ASI Series F Computer and Systems Sciences* 168 (1998) 215–238.
- [19] A. Graves, Practical variational inference for neural networks, *Advances in neural information processing systems* 24 (2011).
- [20] J. M. Hernández-Lobato, R. Adams, Probabilistic backpropagation for scalable learning of bayesian neural networks, in: *International conference*

- on machine learning, *Proceedings of Machine Learning Research*, 2015, pp. 1861–1869.
- [21] R. M. Neal, *Bayesian learning for neural networks*, volume 118, Springer Science & Business Media, 2012.
  - [22] M. Welling, Y. W. Teh, Bayesian learning via stochastic gradient langevin dynamics, in: *Proceedings of the 28th international conference on machine learning (ICML-11)*, Citeseer, 2011, pp. 681–688.
  - [23] C. P. Robert, V. Elvira, N. Tawn, C. Wu, Accelerating mcmc algorithms, *Wiley Interdisciplinary Reviews: Computational Statistics* 10 (2018) e1435.
  - [24] C. Blundell, J. Cornebise, K. Kavukcuoglu, D. Wierstra, Weight uncertainty in neural network, in: *International conference on machine learning*, *Proceedings of Machine Learning Research*, 2015, pp. 1613–1622.
  - [25] J. Hensman, M. Zwiessele, N. D. Lawrence, Tilted variational bayes, in: *Artificial Intelligence and Statistics*, *Proceedings of Machine Learning Research*, 2014, pp. 356–364.
  - [26] A. B. Dieng, D. Tran, R. Ranganath, J. Paisley, D. Blei, Variational inference via  $\chi$  upper bound minimization, *Advances in Neural Information Processing Systems* 30 (2017).
  - [27] J. Deasy, N. Simidjievski, P. Liò, Constraining variational inference with geometric jensen-shannon divergence, *Advances in Neural Information Processing Systems* 33 (2020) 10647–10658.
  - [28] Y. Li, R. E. Turner, Rényi divergence variational inference, *Advances in neural information processing systems* 29 (2016).
  - [29] N. Wan, D. Li, N. Hovakimyan, F-divergence variational inference, *Advances in neural information processing systems* 33 (2020) 17370–17379.
  - [30] F. Nielsen, On the jensen–shannon symmetrization of distances relying on abstract means, *Entropy* 21 (2019) 485.

- [31] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, A. Lerchner, beta-VAE: Learning basic visual concepts with a constrained variational framework, in: International Conference on Learning Representations, 2017. URL: <https://openreview.net/forum?id=Sy2fzU9gl>.
- [32] K. N. Markelle Kelly, Rachel Longjohn, The uci machine learning repository, <https://archive.ics.uci.edu>, Accessed Online; July 14, 2024.
- [33] A. Krizhevsky, G. Hinton, et al., Learning multiple layers of features from tiny images (2009).
- [34] A. Janowczyk, A. Madabhushi, Deep learning for digital pathology image analysis: A comprehensive tutorial with selected use cases, *Journal of pathology informatics* 7 (2016) 29.
- [35] A. Cruz-Roa, A. Basavanahally, F. González, H. Gilmore, M. Feldman, S. Ganesan, N. Shih, J. Tomaszewski, A. Madabhushi, Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks, in: *Medical Imaging 2014: Digital Pathology*, volume 9041, SPIE, 2014, p. 904103.
- [36] Paul Mooney, Breast Histopathology Images, <https://www.kaggle.com/paultimothymooney/breast-histopathology-images>, 2017.
- [37] J. Bergstra, R. Bardenet, Y. Bengio, B. Kégl, Algorithms for hyperparameter optimization, in: 25th annual conference on neural information processing systems (NIPS 2011), volume 24, Neural Information Processing Systems Foundation, 2011.
- [38] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [39] A. Krizhevsky, I. Sutskever, G. E. Hinton, Imagenet classification with deep convolutional neural networks, in: *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [40] J. Bergstra, D. Yamins, D. Cox, Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures, in: *International conference on machine learning*, *Proceedings of Machine Learning Research*, 2013, pp. 115–123.