

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

Ponkrshnan Thiagarajan <sup>a,1</sup>, Susanta Ghosh <sup>b,2</sup>

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

<sup>b</sup>*Mechanical Engineering-Engineering Mechanics and the Institute of Computing and Cybersystems, Michigan Technological University, MI, USA.*

---

## Abstract

The Kullback-Leibler (KL) divergence is widely used in state-of-the-art Bayesian Neural Networks (BNNs) to approximate the posterior distribution of weights. However, the KL divergence is unbounded and asymmetric, which may lead to instabilities during optimization or may yield poor generalizations. To overcome these limitations, we examine the Jensen-Shannon (JS) divergence that is bounded, symmetric, and more general. Towards this, we propose two novel loss functions for BNNs. The first loss function uses the geometric JS divergence (JS-G) that is symmetric, unbounded and offers an analytical expression for Gaussian priors. The second loss function uses the generalized JS divergence (JS-A) that is symmetric and bounded. We show that the conventional KL divergence-based loss function is a special case of the two loss functions presented in this work. To evaluate the divergence part of the loss we use analytical expressions for JS-G and use Monte Carlo methods for JS-A. We provide algorithms to optimize the loss function using both these methods. The proposed loss functions offer additional parameters that can be tuned to control the degree of regularisation. The regularization performance of the JS divergences is analyzed to demonstrate their superiority over the state-of-the-art. Further, we derive the conditions for better regularization by the proposed JS-G divergence-based loss function than the KL divergence-based loss function. Bayesian convolutional neural networks (BCNN) based on the proposed JS divergences perform better than the state-of-the-art BCNN, which is shown for the classification of the CIFAR data set having various degrees of noise and a histopathology data set having a high bias.

*Keywords:* Bayesian convolutional neural networks, Variational inference, KL divergence, JS divergence, Uncertainty quantification.

---

## 1. Introduction

### 1.1. Background

Despite the widespread success of deep neural networks (DNNs) and convolutional neural networks (CNNs) in numerous science and engineering applications [1–10], they suffer from overfitting when the data set is small, noisy or biased [3, 11]. Further, due to a very large number of (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 necessitate high accuracies such as autonomous driving and medical diagnosis. Several methods were developed to provide prediction intervals as a measure of uncertainty in neural networks. Some of these include the delta method, bootstrap method, mean-variance estimation method, Bayesian method, etc. [12]. Amongst these, Bayesian methods have gained eminence due to their rigorous mathematical foundation for uncertainty quantification through their stochastic parameters [12, 13]. Key aspects of Bayesian Neural Networks (BNNs) are briefly summarised in the following.

---

<sup>1</sup>email: thiagara@mtu.edu

<sup>2</sup>email: susantag@mtu.edu

This paper is submitted to IEEE for peer review.

The work is partially supported by the grant DE-SC0023432 funded by the U.S. Department of Energy, Office of Science. We acknowledge the HPC facilities from the SUPERIOR at Michigan Tech and XSEDE (Request#: MSS200004).

### 1.1.1. Bayesian Neural Networks (BNNs)

A BNN is a neural network with stochastic parameters. The posterior distribution of these parameters is learned through the Bayes rule. The first known BNN developed by Tishby et al [14] showed that the posterior distributions of weights and biases can be obtained via the Bayes rule from a chosen prior. Though their work provided theoretical insights, a method to perform inference on the posterior distribution of parameters was still unavailable. To perform inference a Laplace approximation was proposed by Denker and LeCun [15]. Detailed reviews on BNNs are provided in [16, 17]. The two most commonly used techniques to approximate the posterior distribution of parameters (i.e. Bayesian inference) of a BNN are the Variational Inference (VI) and the Markov Chain Monte Carlo Methods (MCMC). 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 [18]. An additional limitation of MCMC algorithms is that they do not scale well with the model size. In view of these, the VI-based approaches are focused on this work. A brief review of some of the early works VI techniques is presented in the following.

### 1.1.2. Variational Inference (VI) in Bayesian Neural Networks

The VI is a technique to approximate an intractable posterior distribution by a tractable distribution called the variational distribution. The variational distribution is learned by minimizing its dissimilarity with respect to the true posterior. The first work that resembles VI was proposed by Hinton and Van Camp [19] to learn the posterior distribution of network parameters using an information theory-based approach. Their work factorized the posterior distribution of neural network parameters thereby neglecting the correlations between them. A full correlation between the network parameters was introduced by Barber and Bishop [20]. These two methods ([19, 20]) are focused on analytical representations of the optimization objective which restricts their use to single hidden layer networks and could not be used for larger networks.

### 1.1.3. Recent Advances on Bayesian Neural Networks

High computational cost was the bottleneck for BNNs in the late 90s. With the advancement in GPU computation, research in the field of BNN has resurged. A data sub-sampling technique to scale for large amounts of data in a VI objective has been proposed by Graves [21]. In addition, their method minimized the evidence lower bound (ELBO) to learn the network parameters. While their method enabled the application of VI for large networks, it performs poorly due to high variance in the computed gradients via Monte Carlo approximation and also due to a lack of correlations between weights. To improve upon their approach a novel algorithm for backpropagation has been proposed by Hernandez et al. [22]. Blundell et al. [23] reparametrized the stochastic variables and obtained unbiased gradient estimates of the loss function to facilitate backpropagation.

To summarise, VI methods are efficient and they 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 dissimilarity between the true and the variational posterior. However, the KL divergence has two key limitations, first, it is asymmetric with respect to its arguments, and second, its values are unbounded. These limitations may lead to poor generalization or failure during training as reported in Hensman et al. [24]. Therefore, it is imperative to explore alternative divergences for VI that can alleviate these limitations.

In this regard, Renyi's  $\alpha$ -divergences have been introduced for VI by Li and Turner [25]. They proposed a family of variational methods that unified various existing approaches. A  $\chi$ -divergence-based VI has been proposed by 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 by Wan et al. to use VI for all f-divergences that unified the Reyni divergence [25] and  $\chi$  divergence [26] based VIs. A modification to the skew-geometric Jensen-Shanon (JS) divergence has been proposed by Deasy et al. to introduce a new loss function for Variational Auto Encoders (VAEs) [27] which has shown a better reconstruction and generation as compared to existing VAEs. However, a JS divergence-based loss function is not yet available for BNNs or Bayesian CNNs which are widely used.

## 1.2. Present work

In the present work, we propose the following two novel loss functions for BNNs, which are based on:

1. The skew-geometric JS divergence.
2. A modified generalized JS divergence.

These loss functions are obtained through a constrained optimization approach in contrast to the variational inference used in BNNs. These JS divergence-based loss functions are generalizations of the state-of-the-art KL divergence-based ELBO loss function [21–23]. We explain why these loss functions should perform better. In addition, we derive the conditions under which the proposed skew-geometric JS divergence-based loss function regularises better than the KL divergence-based loss function. Further, we show that the loss functions presented in this work perform better for image classification problems where the data set has noise or is biased towards a particular class.

The rest of the manuscript is organized as follows: Sec. 2 provides the necessary mathematical backgrounds. Sec. 3 describes the proposed modification to the JS divergence and the proposed loss functions. The details of the numerical experiments are provided in Sec. 4. The results obtained by the present method are described and discussed in Sec. 5. Finally, conclusions are provided in Sec. 6.

## 2. Mathematical Background

In this section, we provide some of the mathematical backgrounds for KL and JS divergences and the variational inference.

### 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

$$\text{KL}[p \parallel q] = \int_{\Omega} p(x) \log \left[ \frac{p(x)}{q(x)} \right] dx \quad (1)$$

where  $p(x)$  and  $q(x)$  are the probability densities of  $\mathcal{P}$  and  $\mathcal{Q}$  respectively. For two  $n$ -dimensional multivariate Gaussian distributions<sup>3</sup>  $\mathcal{N}_1(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$  and  $\mathcal{N}_2(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)$ , the KL divergence is given by

$$\text{KL}(\mathcal{N}_1 \parallel \mathcal{N}_2) = \frac{1}{2} \left[ \text{tr}(\boldsymbol{\Sigma}_2^{-1} \boldsymbol{\Sigma}_1) + \ln \frac{|\boldsymbol{\Sigma}_2|}{|\boldsymbol{\Sigma}_1|} + (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1)^T \boldsymbol{\Sigma}_2^{-1} (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1) - n \right] \quad (2)$$

The KL divergence is widely used in literature to represent the dissimilarity between two probability distributions for various applications such as VI. However, it has limitations such as the asymmetric property i.e  $\text{KL}[p \parallel q] \neq \text{KL}[q \parallel p]$ , and unboundedness. These limitations may lead to difficulty in approximating light-tailed posteriors as reported in [24].

To overcome these limitations a symmetric JS divergence can be employed. JS divergence is defined as

$$\text{JS}[p \parallel q] = \frac{1}{2} \text{KL} \left[ p \parallel \frac{p+q}{2} \right] + \frac{1}{2} \text{KL} \left[ q \parallel \frac{p+q}{2} \right] \quad (3)$$

The above divergence can be further generalized with the weighted arithmetic mean of  $p$  and  $q$  defined as  $A_{\alpha} = (1 - \alpha)p + \alpha q$ ,

$$\text{JS}^{A_{\alpha}}[p \parallel q] = (1 - \alpha) \text{KL}(p \parallel A_{\alpha}) + \alpha \text{KL}(q \parallel A_{\alpha}) \quad (4)$$

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 abstract means, specifically the geometric mean was proposed by Nielsen [28]. 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

$$\text{JS}^{G_{\alpha}}[p \parallel q] = (1 - \alpha) \text{KL}(p \parallel G_{\alpha}(p, q)) + \alpha \text{KL}(q \parallel G_{\alpha}(p, q)) \quad (5)$$

---

<sup>3</sup>  $\mathcal{N}$  denotes a Gaussian distribution throughout this work.

The parameter  $\alpha$ , called a skew parameter, controls the divergence skew between  $p$  and  $q$ . However, the skew-geometric divergence in Eq.(5) fails to capture the divergence between  $p$  and  $q$  when  $\alpha = 0$  or  $\alpha = 1$  since it yields a value of zero for the divergence. To resolve this issue, Deasy et al. [27] used the following reverse form of geometric mean  $G'_\alpha(x, y) = x^\alpha y^{1-\alpha}$ , with  $\alpha \in [0, 1]$  to implement JS divergence 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

$$\text{JS-G}[p \parallel q] = (1 - \alpha)\text{KL}(p \parallel G'_\alpha(p, q)) + \alpha\text{KL}(q \parallel G'_\alpha(p, q)) \quad (6)$$

This yields KL divergences in the limiting values of the skew parameter.

In the case of two multivariate Gaussian distributions, the JS-G can be written as

$$\begin{aligned} \text{JS-G}(\mathcal{N}_1 \parallel \mathcal{N}_2) &= (1 - \alpha)\text{KL}(\mathcal{N}_1 \parallel \mathcal{N}'_\alpha) + \alpha\text{KL}(\mathcal{N}_2 \parallel \mathcal{N}'_\alpha) \\ &= \frac{1}{2} \left[ \text{tr} \left( \frac{(1 - \alpha)\mathbf{\Sigma}_1 + \alpha\mathbf{\Sigma}_2}{\mathbf{\Sigma}'_\alpha} \right) + \log \left( \frac{|\mathbf{\Sigma}'_\alpha|}{|\mathbf{\Sigma}_1|^{1-\alpha} |\mathbf{\Sigma}_2|^\alpha} \right) + (1 - \alpha)(\boldsymbol{\mu}'_\alpha - \boldsymbol{\mu}_1)^T \mathbf{\Sigma}_\alpha^{-1} (\boldsymbol{\mu}'_\alpha - \boldsymbol{\mu}_1) + \alpha(\boldsymbol{\mu}'_\alpha - \boldsymbol{\mu}_2)^T \mathbf{\Sigma}_\alpha^{-1} (\boldsymbol{\mu}'_\alpha - \boldsymbol{\mu}_2) - n \right] \end{aligned} \quad (7)$$

where,  $\mathcal{N}'_\alpha = \mathcal{N}(\boldsymbol{\mu}'_\alpha, \mathbf{\Sigma}'_\alpha)$  is the geometric mean distribution of  $\mathcal{N}_1$  and  $\mathcal{N}_2$  which has the following the parameters

$$\mathbf{\Sigma}'_\alpha = \left[ \alpha\mathbf{\Sigma}_1^{-1} + (1 - \alpha)\mathbf{\Sigma}_2^{-1} \right]^{-1}; \quad \boldsymbol{\mu}'_\alpha = \mathbf{\Sigma}_\alpha \left[ \alpha\mathbf{\Sigma}_1^{-1}\boldsymbol{\mu}_1 + (1 - \alpha)\mathbf{\Sigma}_2^{-1}\boldsymbol{\mu}_2 \right] \quad (8)$$

Note that for  $\alpha \in [0, 1]$ ,  $\text{JS-G}(\mathcal{N}_1 \parallel \mathcal{N}_2)|_\alpha = \text{JS-G}(\mathcal{N}_2 \parallel \mathcal{N}_1)|_{1-\alpha}$  which is not symmetric. However, for  $\alpha = 0.5$ , the JS-G is symmetric with  $\text{JS-G}(\mathcal{N}_1 \parallel \mathcal{N}_2)|_{\alpha=0.5} = \text{JS-G}(\mathcal{N}_2 \parallel \mathcal{N}_1)|_{\alpha=0.5}$ .

The geometric JS divergences,  $\text{JS}^{\text{G}_\alpha}$  and JS-G, given in Eq.(5) and Eq.(6) have analytical expressions when  $p$  and  $q$  are Gaussians. However, they are unbounded like the KL divergence. Whereas, the generalized JS divergence  $\text{JS}^{A_\alpha}$  in Eq.(4) 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 are interested in learning a data-driven model (for instance a BNN) to predict the probability  $P(\mathbf{y}|\mathbf{x}, \mathbb{D})$  of output  $\mathbf{y} \in \mathcal{Y}$ , where  $\mathcal{Y}$  is the output space. Towards this, we train a neural network with random parameters  $\mathbf{w}$  and learn the posterior probability distribution of these parameters  $P(\mathbf{w}|\mathbb{D})$  from the training data  $\mathbb{D}$  using the Bayes' rule as follows:

$$P(\mathbf{w}|\mathbb{D}) = \frac{P(\mathbb{D}|\mathbf{w})P(\mathbf{w})}{P(\mathbb{D})} \quad (9)$$

Where,  $P(\mathbb{D}|\mathbf{w})$  is the likelihood of the training data  $\mathbb{D}$  given parameters  $\mathbf{w}$ .  $P(\mathbf{w})$  is the prior distribution of weights. 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

$$P(\mathbf{y}|\mathbf{x}, \mathbb{D}) = \int_{\Omega_{\mathbf{w}}} P(\mathbf{y}|\mathbf{x}, \mathbf{w})P(\mathbf{w}|\mathbb{D})d\mathbf{w} \quad (10)$$

The term  $P(\mathbb{D})$  in Eq. (9) 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  $q(\mathbf{w}|\boldsymbol{\theta})$  over the model weights.  $q(\mathbf{w}|\boldsymbol{\theta})$  is referred to as the variational posterior having parameters  $\boldsymbol{\theta}$ . For example, if  $q$  is assumed to be a Gaussian distribution then the parameters  $\boldsymbol{\theta}$  are the mean  $\boldsymbol{\mu}$  and the standard deviation  $\boldsymbol{\sigma}$ . In variational inference, 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 [23].

$$\begin{aligned} \boldsymbol{\theta}^* &= \arg \min_{\boldsymbol{\theta}} \text{KL} [q(\mathbf{w}|\boldsymbol{\theta}) \parallel P(\mathbf{w}|\mathbb{D})] \\ &= \arg \min_{\boldsymbol{\theta}} \int q(\mathbf{w}|\boldsymbol{\theta}) \log \left[ \frac{q(\mathbf{w}|\boldsymbol{\theta})}{P(\mathbf{w})P(\mathbb{D}|\mathbf{w})} P(\mathbb{D}) \right] d\mathbf{w} \end{aligned} \quad (11)$$

Note that the term  $P(\mathbb{D})$  can be eliminated from Eq. (11) since it is a constant and does not affect the optimization. The resulting loss function  $\mathcal{F}(\mathbb{D}, \theta)$ , which is to be minimised to learn the optimal parameters  $\theta^*$  is expressed as:

$$\mathcal{F}_{KL}(\mathbb{D}, \theta) = \text{KL}[q(\mathbf{w}|\theta) \| P(\mathbf{w})] - \mathbb{E}_{q(\mathbf{w}|\theta)}[\log P(\mathbb{D}|\mathbf{w})] \quad (12)$$

This loss function is known as the variational free energy or the evidence lower bound (ELBO) [21, 23].

### 3. Methods

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

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

The Generalised JS divergence given in Eq. (4) fails to capture the divergence between  $p$  and  $q$  in the limiting cases of  $\alpha$  since,

$$\left[ \text{JS}^{\alpha}(p \| q) \right]_{\alpha=0} = 0; \quad \left[ \text{JS}^{\alpha}(p \| q) \right]_{\alpha=1} = 0 \quad (13)$$

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

$$\text{JS-A}[p \| q] = (1 - \alpha)\text{KL}(p \| A'_\alpha) + \alpha\text{KL}(q \| A'_\alpha) \quad (14)$$

This modification yields KL divergences in the limiting cases of  $\alpha$  as,

$$\left[ \text{JS-A}(p \| q) \right]_{\alpha=0} = \text{KL}(p \| q); \quad \left[ \text{JS-A}(p \| q) \right]_{\alpha=1} = \text{KL}(q \| p) \quad (15)$$

#### **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,

$$\text{JS-A} \leq -(1 - \alpha) \log \alpha - \alpha \log(1 - \alpha)$$

#### **Proof:**

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

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,  $\mathcal{H} \geq 0$ ;  $\forall t \in \Omega$

We provide a comparison of symmetry and boundedness for various divergences used in this work in Table 1.

Table 1: Properties of various divergences

Divergence	Bounded	Symmetric	Analytical expression
KL	×	×	✓
JS-A	✓	✓	×
JS-G	×	✓	✓

### 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 divergence is used instead of the KL divergence in the VI setting (see Eq. (11)), the optimization problem becomes,

$$\theta^* = \arg \min_{\theta} \text{JS-G} [q(\mathbf{w}|\theta) \parallel P(\mathbf{w}|\mathbb{D})] \quad (16)$$

The loss function can then be written as,

$$\begin{aligned} \mathcal{F}_{\text{JS-G}}(\mathbb{D}, \theta) &= \text{JS-G} [q(\mathbf{w}|\theta) \parallel P(\mathbf{w}|\mathbb{D})] \\ &= (1 - \alpha)\text{KL}(q \parallel G'_\alpha(q, P)) + \alpha\text{KL}(P \parallel G'_\alpha(q, P)) \end{aligned} \quad (17)$$

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

Rewriting the first term in Eq. (17) as,

$$\begin{aligned} T_1 &= (1 - \alpha)\text{KL}(q \parallel G'_\alpha(q, P)) \\ &= (1 - \alpha)\text{KL}[q(\mathbf{w}|\theta) \parallel q(\mathbf{w}|\theta)^\alpha P(\mathbf{w}|\mathbb{D})^{(1-\alpha)}] \\ &= (1 - \alpha) \int q(\mathbf{w}|\theta) \log \left[ \frac{q(\mathbf{w}|\theta)}{q(\mathbf{w}|\theta)^\alpha P(\mathbf{w}|\mathbb{D})^{(1-\alpha)}} \right] d\mathbf{w} \\ &= (1 - \alpha)^2 \int q(\mathbf{w}|\theta) \log \left[ \frac{q(\mathbf{w}|\theta)}{P(\mathbf{w}|\mathbb{D})} \right] d\mathbf{w} \end{aligned} \quad (18)$$

which is equivalent to loss function in Eq. (12) multiplied by a constant  $(1 - \alpha)^2$ . Similarly rewriting the second term in Eq. (17) as,

$$\begin{aligned} T_2 &= \alpha\text{KL}(P \parallel G'_\alpha(q, P)) \\ &= \alpha\text{KL}[P(\mathbf{w}|\mathbb{D}) \parallel q(\mathbf{w}|\theta)^\alpha P(\mathbf{w}|\mathbb{D})^{(1-\alpha)}] \\ &= \alpha \int P(\mathbf{w}|\mathbb{D}) \log \left[ \frac{P(\mathbf{w}|\mathbb{D})}{q(\mathbf{w}|\theta)^\alpha P(\mathbf{w}|\mathbb{D})^{(1-\alpha)}} \right] d\mathbf{w} \\ &= \alpha^2 \int P(\mathbf{w}|\mathbb{D}) \log \left[ \frac{P(\mathbf{w}|\mathbb{D})}{q(\mathbf{w}|\theta)} \right] d\mathbf{w} \end{aligned} \quad (19)$$

The term  $P(\mathbf{w}|\mathbb{D})$  in Eq. (19) is intractable as explained in section 2.2, which makes  $T_2$  intractable.<sup>4</sup> Therefore the JS-G divergence-based loss function given in Eq. (17) cannot be used to find the optimum parameter  $\theta^*$  which contrasts the KL divergence based loss function in Eq. (12). 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 [27, 29] 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.

<sup>4</sup> We obtain a similar result when the geometric mean  $G_\alpha$  is used instead of  $G'_\alpha$ . Only the constants outside the integral change in that case.

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

$$\max_{\theta} \mathbb{E}_{q(\mathbf{w}|\theta)} [\log P(\mathbb{D}|\mathbf{w})] \quad (20)$$

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

$$\max_{\theta} \mathbb{E}_{q(\mathbf{w}|\theta)} [\log P(\mathbb{D}|\mathbf{w})] \quad \text{subject to } D(q(\mathbf{w}|\theta) \| P(\mathbf{w})) < \epsilon \quad (21)$$

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

$$\mathcal{L} = \mathbb{E}_{q(\mathbf{w}|\theta)} [\log P(\mathbb{D}|\mathbf{w})] - \lambda(D(q(\mathbf{w}|\theta) \| P(\mathbf{w})) - \epsilon) \quad (22)$$

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>5</sup>.

$$\tilde{\mathcal{F}}_D = \lambda D(q(\mathbf{w}|\theta) \| P(\mathbf{w})) - \mathbb{E}_{q(\mathbf{w}|\theta)} [\log P(\mathbb{D}|\mathbf{w})] \quad (23)$$

This loss function reproduces the ELBO loss [23] when KL divergence is used and  $\lambda$  is taken as 1.

In the following, we obtain loss functions for two JS divergences, namely, the Geometric J-S divergence, and the Modified Generalised J-S divergence.

### 3.3.1. Geometric J-S divergence

Using the modified skew-geometric J-S divergence as the measure of divergence in Eq.(23) leads to the following loss function:

$$\tilde{\mathcal{F}}_{JSG} = \lambda \text{JS-G}(q(\mathbf{w}|\theta) \| P(\mathbf{w})) - \mathbb{E}_{q(\mathbf{w}|\theta)} [\log P(\mathbb{D}|\mathbf{w})] \quad (24a)$$

$$= \lambda(1 - \alpha) \text{KL}(q \| G'_{\alpha}(q, P_w)) + \lambda\alpha \text{KL}(P_w \| G'_{\alpha}(q, P_w)) - \mathbb{E}_{q(\mathbf{w}|\theta)} [\log P(\mathbb{D}|\mathbf{w})] \quad (24b)$$

Note,

$$\begin{aligned} \text{KL}(q \| G'_{\alpha}(q, P_w)) &= \int q(\mathbf{w}|\theta) \log \frac{q(\mathbf{w}|\theta)}{q(\mathbf{w}|\theta)^{\alpha} P(\mathbf{w})^{1-\alpha}} d\mathbf{w} \\ &= (1 - \alpha) \int q(\mathbf{w}|\theta) \log \frac{q(\mathbf{w}|\theta)}{P(\mathbf{w})} d\mathbf{w} \end{aligned} \quad (25)$$

and

$$\begin{aligned} \text{KL}(P_w \| G'_{\alpha}(q, P_w)) &= \int P(\mathbf{w}) \log \frac{P(\mathbf{w})}{q(\mathbf{w}|\theta)^{\alpha} P(\mathbf{w})^{1-\alpha}} d\mathbf{w} \\ &= \alpha \int P(\mathbf{w}) \log \frac{P(\mathbf{w})}{q(\mathbf{w}|\theta)} d\mathbf{w} \end{aligned} \quad (26)$$

Hence, the loss function can be written as,

$$\tilde{\mathcal{F}}_{JSG} = \lambda(1 - \alpha)^2 \int q(\mathbf{w}|\theta) \log \frac{q(\mathbf{w}|\theta)}{P(\mathbf{w})} d\mathbf{w} + \lambda\alpha^2 \int P(\mathbf{w}) \log \frac{P(\mathbf{w})}{q(\mathbf{w}|\theta)} d\mathbf{w} - \mathbb{E}_{q(\mathbf{w}|\theta)} [\log P(\mathbb{D}|\mathbf{w})] \quad (27)$$

Eq. (27) can be rewritten in terms of expectations as,

$$\tilde{\mathcal{F}}_{JSG} = \lambda(1 - \alpha)^2 \mathbb{E}_{q(\mathbf{w}|\theta)} \left[ \log \frac{q(\mathbf{w}|\theta)}{P(\mathbf{w})} \right] + \lambda\alpha^2 \mathbb{E}_{P(\mathbf{w})} \left[ \log \frac{P(\mathbf{w})}{q(\mathbf{w}|\theta)} \right] - \mathbb{E}_{q(\mathbf{w}|\theta)} [\log P(\mathbb{D}|\mathbf{w})] \quad (28)$$

<sup>5</sup> The constrained optimization approach-based loss functions are marked by an overhead tilde.

In Eq.(27), the first term is the the *mode seeking* reverse KL divergence  $\text{KL}(q(\mathbf{w}|\theta)\|P(\mathbf{w}))$  and the second term is the *mean seeking* forward KL divergence  $\text{KL}(P(\mathbf{w})\|q(\mathbf{w}|\theta))$ . 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. 3.5.1. The parameters  $\lambda$  and  $\alpha$  can be used to control the amount of regularisation.

### 3.3.2. Modified Generalised J-S divergence

Using the modified Generalised JS divergence as the measure of divergence in Eq.(23) leads to the following loss function:

$$\begin{aligned}\widetilde{\mathcal{F}}_{JSA} &= \lambda \text{JS-A}(q(\mathbf{w}|\theta) \| P(\mathbf{w})) - \mathbb{E}_{q(\mathbf{w}|\theta)} [\log P(\mathbb{D}|\mathbf{w})] \\ &= \lambda(1 - \alpha) \text{KL}(q \| A'_\alpha(q, P_w)) + \lambda\alpha \text{KL}(P_w \| A'_\alpha(q, P_w)) - \mathbb{E}_{q(\mathbf{w}|\theta)} [\log P(\mathbb{D}|\mathbf{w})]\end{aligned}\quad (29)$$

Where,  $A'_\alpha(q, P_w) = \alpha q + (1 - \alpha)P_w$ . The above equation, Eq.(29), can be expanded as,

$$\widetilde{\mathcal{F}}_{JSA} = \lambda(1 - \alpha)\mathbb{E}_{q(\mathbf{w}|\theta)} \left[ \log \frac{q(\mathbf{w}|\theta)}{A'_\alpha(q, P_w)} \right] + \lambda\alpha\mathbb{E}_{P(\mathbf{w})} \left[ \log \frac{P(\mathbf{w})}{A'_\alpha(q, P_w)} \right] - \mathbb{E}_{q(\mathbf{w}|\theta)} [\log P(\mathbb{D}|\mathbf{w})] \quad (30)$$

It is to be noted that the proposed loss functions in Eq. (28) and Eq. (30) yield the ELBO loss for  $\alpha = 0$  and  $\lambda = 1$ . The minimization of the two loss functions Eq. (28) and Eq. (30) are given in the following subsection.

## 3.4. Minimisation of the proposed loss functions

### 3.4.1. 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. The closed-form expression of the JS-G divergence for two Gaussian distributions is presented earlier in Sec.3.3. 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  $P_N(\mathbf{w}|\theta) = \mathcal{N}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2^2)$  and  $q_N(\mathbf{w}) = \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1^2)$  respectively <sup>6</sup>. The closed-form expression of the JS-G divergence between  $q_N(\mathbf{w})$  and  $P_N(\mathbf{w}|\theta)$  can be written as,

$$\text{JS-G}(q_N\|P_N) = \frac{1}{2} \sum_{i=1}^n \left[ \frac{(1 - \alpha)\sigma_{1i}^2 + \alpha\sigma_{2i}^2}{\sigma_{ai}^2} + \log \frac{(\sigma'_{ai})^2}{\sigma_{1i}^{2(1-\alpha)}\sigma_{2i}^{2\alpha}} + (1 - \alpha) \frac{(\mu'_{ai} - \mu_{1i})^2}{(\sigma'_{ai})^2} + \frac{\alpha(\mu'_{ai} - \mu_{2i})^2}{(\sigma'_{ai})^2} - 1 \right] \quad (31)$$

where,

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

Therefore, the divergence term of the proposed loss function, the first term in Eq. (24a), can be evaluated by this closed-form expression given in Eq. (31). The expectation of the log-likelihood, the second term in Eq. (24a), can be approximated by a Monte-Carlo sampling <sup>7</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.

### 3.4.2. 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 is applicable to both the JS-G and the JS-A divergences with no restrictions on the priors. The loss functions given in Eq. (28) and Eq. (30) can be approximated using Monte Carlo samples from the corresponding distributions as shown in Algorithm 2.

<sup>6</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  $\boldsymbol{\mu}_1 = [\mu_{11}, \mu_{12}, \dots, \mu_{1n}]^T$  and  $\boldsymbol{\Sigma}_1^2 = \text{diag}(\sigma_{11}^2, \sigma_{12}^2, \dots, \sigma_{1n}^2)$  (similarly for  $\boldsymbol{\mu}_2$  and  $\boldsymbol{\Sigma}_2^2$ ).

<sup>7</sup>The approximations to the loss functions are denoted by F



---

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

---

**Evaluate** JS-G term of Eq. 24a analytically using Eq. 31

**Evaluate**  $\mathbb{E}_{q(\mathbf{w}|\theta)}$   $[\log P(\mathbb{D}|\mathbf{w})]$  term of Eq. 24a

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

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

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

**Loss:**

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

**Gradients:**

$$\frac{\partial F}{\partial \boldsymbol{\mu}} \leftarrow \sum_i \frac{\partial F}{\partial \mathbf{w}_i} + \frac{\partial F}{\partial \boldsymbol{\mu}}$$
$$\frac{\partial F}{\partial \boldsymbol{\rho}} \leftarrow \sum_i \frac{\partial F}{\partial \mathbf{w}_i} \frac{\varepsilon_i}{1 + \exp(-\boldsymbol{\rho})} + \frac{\partial F}{\partial \boldsymbol{\rho}}$$

**Update:**

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

---

### 3.5. 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 two aspects of the proposed loss functions:

1. The ability to regularize.
2. The advantage of using the closed-form expression over the Monte Carlo approximation while evaluating the JS-G divergence.

#### 3.5.1. Regularisation performance of JS-G divergence

Let two Gaussian distributions  $q = \mathcal{N}(\mu_1 = \mu, \sigma_1^2 = 0.01)$  and  $P = \mathcal{N}(\mu_2 = 0, \sigma_2^2 = 0.1)$  represent the posterior and the prior distribution of a parameter in a BNN. The closed-form divergences for KL and JS-G are evaluated by varying the mean of the distribution  $q$ . This emulates the learning of the network parameter during training. In Fig. 1, the values of the divergences are plotted against  $\mu$ , which is the difference between the mean of the prior and posterior distributions. Fig. 1 shows that as the mean of the posterior distribution moves away from the mean of the prior distribution, the JS-G divergence increases more rapidly than the KL divergence. This implies that a greater penalization is offered by JS-G divergence than the KL divergence as the mean of the posterior deviates away from the prior. Thus, by assuming small values for the means of prior distributions we can regularize better by the proposed JS-G divergence. This regularization process requires finding an optimal value of  $\alpha$  through hyperparameter optimization. In the following subsection, we provide the condition for the existence of  $\alpha$  for any two arbitrary Gaussian distributions  $P$  and  $q$ .

---

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

---

**Approximate**  $\mathbb{E}_{q(\mathbf{w}|\theta)}$  terms of Eq. (28) or Eq. (30)

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

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

Evaluate first and third terms of Eq. (28):

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

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

(or)

Evaluate first and third terms of Eq. (30):

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

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

**Approximate**  $\mathbb{E}_{P(\mathbf{w})}$  terms of Eq. (28) or Eq. (30)

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

Evaluate second term of Eq. (28):

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

$$\text{where, } c_2 = \lambda\alpha^2$$

(or)

Evaluate second term of Eq. (30):

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

$$\text{where, } c_2 = \lambda\alpha^2$$

**Loss:**

$$F \leftarrow f_1 + f_2$$

**Gradients:**

$$\frac{\partial F}{\partial \boldsymbol{\mu}} \leftarrow \sum_i \frac{\partial F}{\partial \mathbf{w}_i^q} + \frac{\partial F}{\partial \boldsymbol{\mu}}$$

$$\frac{\partial F}{\partial \boldsymbol{\rho}} \leftarrow \sum_i \frac{\partial F}{\partial \mathbf{w}_i^q} \frac{\boldsymbol{\varepsilon}_i}{1 + \exp(-\boldsymbol{\rho})} + \frac{\partial F}{\partial \boldsymbol{\rho}}$$

**Update:**

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




Figure 1: The value of divergence is plotted against the mean  $\mu$  of  $q$ . Since the mean of  $P$  is fixed to zero,  $\mu$  is also the difference in the mean values between the two distributions  $q$  and  $P$

### 3.5.2. Condition for better regularisation of $\widetilde{\mathcal{F}}_{JSG}$

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  $q$ ,  $\exists \alpha \in [0, \infty) : \widetilde{\mathcal{F}}_{JSG} > \mathcal{F}_{KL}$

**Proof:**

In the following, we seek the condition for

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

This leads to the following necessary condition on  $\alpha$

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

since KL divergence is always non-negative for any distributions, we have

$$\frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} \geq 0$$

Which yields,

$$\alpha > \frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} \geq 0 \quad (33)$$

This proves the theorem, that there exists a positive value of  $\alpha$  for any  $P$  and  $q$  such that  $\widetilde{\mathcal{F}}_{JSG}$  is greater than  $\mathcal{F}_{KL}$ .

**Theorem 3.** If  $P$  and  $q$  are Gaussian distributions with  $P = \mathcal{N}(\mu_p, \sigma_p^2)$  and  $q = \mathcal{N}(\mu_q, \sigma_q^2)$  then,  $\exists \alpha \in [0, 1]$  such that  $\widetilde{\mathcal{F}}_{JSG} > \mathcal{F}_{KL}$  if and only if  $\sigma_p^2 > \sigma_q^2$

**Proof:**

From Eq.(33) and  $\alpha \in [0, 1]$ , the necessary condition on  $\alpha$  is

$$1 \geq \alpha > \frac{2 \text{KL}(q||P)}{\text{KL}(q||P) + \text{KL}(P||q)} > 0 \quad (34)$$

Which leads to the following necessary condition

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

or,

$$\text{KL}(P||q) > \text{KL}(q||P) \quad (35)$$

Let,  $P$  and  $q$  be Gaussian distributions with  $P = \mathcal{N}(\mu_p, \sigma_p^2)$  and  $q = \mathcal{N}(\mu_q, \sigma_q^2)$ . Then, Eq.(35) 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,

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

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

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

or,

$$\ln \left[ \frac{1}{\gamma^2} \exp \left( \gamma - \frac{1}{\gamma} \right) \right] + \frac{(\mu_q - \mu_p)^2}{\sigma_q^2} \left(1 - \frac{1}{\gamma}\right) > 0 \quad (36)$$

This condition Eq.(36) is satisfied only when  $\gamma > 1$  or

$$\sigma_p^2 > \sigma_q^2 \quad (37)$$

Thus, when  $P$  and  $q$  are Gaussian distributions with  $\sigma_p^2 > \sigma_q^2$ , there exists an  $\alpha \in [0, 1]$  such that  $\tilde{\mathcal{F}}_{JSG} > \mathcal{F}_{KL}$ .

Therefore, Theorems 1 and 2 provide us with the conditions under which the JS-G divergence-based loss function regularises better than the KL divergence-based loss function.

### 3.5.3. Regularisation performance of JS-A divergence

In absence of the closed-form expressions for the JS-A divergence, it is evaluated using a Monte Carlo method and compared against the KL divergence for the distributions  $q = \mathcal{N}(\mu_1 = \mu, \sigma_1^2 = 0.01)$  and  $P = \mathcal{N}(\mu_2 = 0, \sigma_2^2 = 0.1)$  in Fig.2. The JS-A divergence part of the loss function can be tuned with the parameter  $\lambda$  to achieve a desired level of




Figure 2: The JS-A divergence evaluated using an MC method is compared against the KL divergence for various means of  $q$ .

regularization as shown in Fig. 2. Thus, the JS-A-based loss function  $\tilde{\mathcal{F}}_{JSA}$  should result in better regularisation than the KL divergence-based loss function when an appropriate value for the hyperparameter  $\lambda$  is chosen.

### 3.5.4. 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 Appendix Appendix A. 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.

## 4. Experiments

In order to demonstrate the advantages of the proposed losses in comparison to the KL loss, we performed experiments. We have implemented the JS-G loss and the JS-A loss via a closed-form expression and a Monte-Carlo method respectively, in these experiments <sup>8</sup>.

### 4.1. Data sets

The following experiments were performed on two Data sets: the CIFAR-10 data set and a benchmark histopathology data set.

#### 4.1.1. CIFAR-10 data set

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.

#### 4.1.2. Histopathology data set

We used a benchmark breast histopathology data set [30] which is highly biased towards one class. Further details on these data sets and the pre-processing steps used here are provided in Appendix Appendix B.

### 4.2. Hyperparameter optimisation and network Architecture

Hyperparameters for all the networks considered here are chosen through hyperparameter optimization. A Tree-structured Parzen Estimator (TPE) algorithm [31] is used here 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 [32] is used to implement this optimization algorithm over a given search space. Data is split into a training set and a validation set as explained in Appendix Appendix B. An optimization is performed to maximize the validation accuracy for different hyperparameter settings of the network. The results of the hyperparameter optimization are given in Appendix Appendix C.

The architecture of all the networks used in this work follows the ResNet-18 V1 model [33] without the batch normalization layers. The network parameters are initialized with the weights of ResNet-18 trained on the Imagenet data set[34].

## 5. Results and discussions

This section presents the classification results on the aforementioned two data sets. A comparison of the performance is presented between the KL loss and the JS losses proposed in this work.

### 5.1. Training and Validation

Three Bayesian CNNs were trained in this work by minimizing the KL loss and the proposed JS losses. Training of the networks is done until the loss converges or the validation accuracy starts to decrease.

#### 5.1.1. CIFAR-10

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. It is evident from Fig. 3 that KL loss over-fits the training data by learning the noise whereas the JS losses regularise better implying a greater generalization. Fig. 3 provides the result for a representative case of  $\sigma = 0.9$ . A similar trend was observed for all other noise levels as well.

---

<sup>8</sup> In this section, the state-of-the-art KL divergence-based ELBO loss function, the JS-G-based loss function, and the JS-A-based loss function are referred to as KL loss, JS-G loss, and JS-A loss respectively.

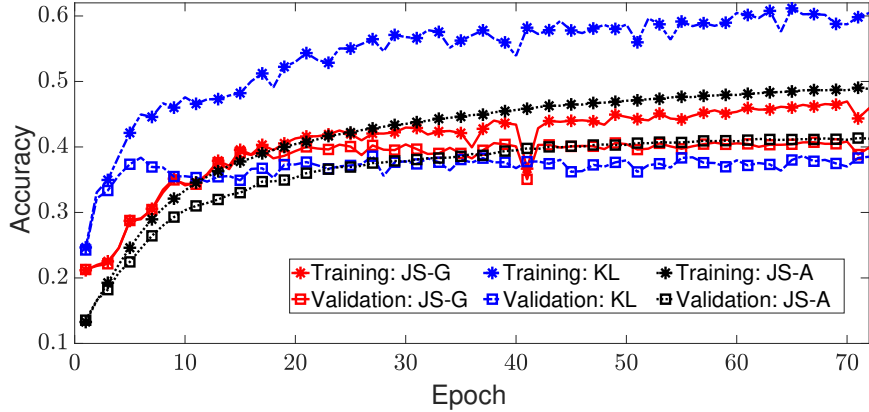


Figure 3: Training and validation of CIFAR-10 data set with added Gaussian noise  $\mathcal{N}(\mu = 0, \sigma = 0.9)$  using the KL-loss function and the proposed JS losses.

### 5.1.2. Histopathology

Training of the histopathology data set is performed using the hyperparameters obtained in Table. C.3. 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. 4 shows the accuracy of training and validation sets for the KL loss and the proposed JS losses.

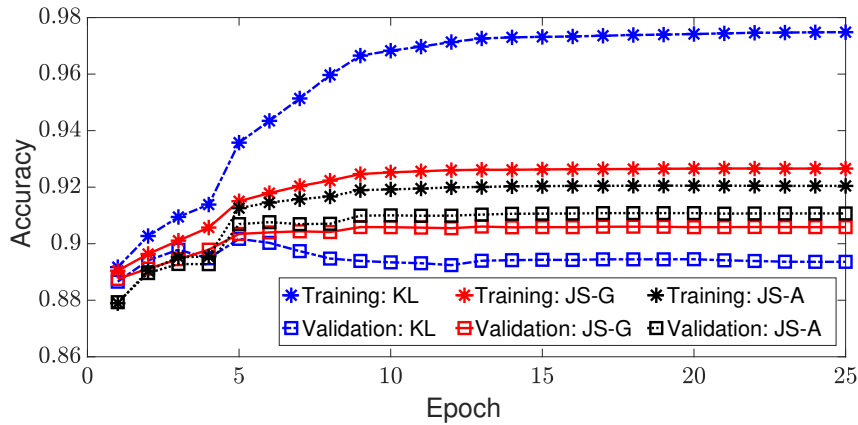


Figure 4: Training and validation of histopathology data set using the KL loss and the proposed JS losses.

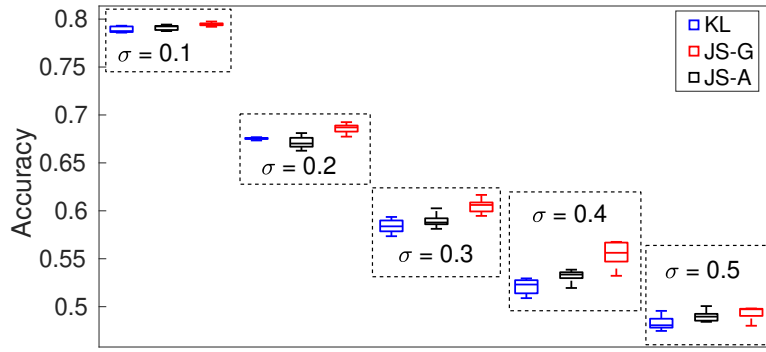
Similar to the CIFAR-10 data set, the KL loss learns the training data too well and fails to generalize for the unseen validation set. Whereas, the proposed JS losses regularise better and provide more accurate results for the validation set.

## 5.2. Testing

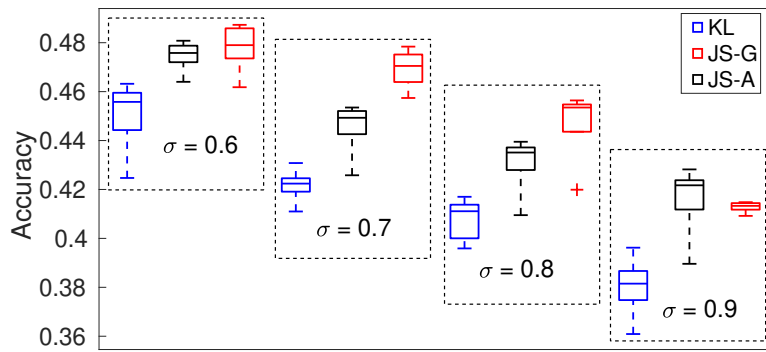
Results obtained on the test sets of the CIFAR-10 data set and the histopathology data set are presented in this section. The test results correspond to the epoch in which the validation accuracy was maximum. Five runs were performed with different mutually exclusive training and validation tests to compare the results of the KL loss and the proposed JS losses.

### 5.2.1. CIFAR-10

The accuracy of the noisy CIFAR-10 test data set at varying noise levels is presented in Fig. 5a and Fig. 5b. It is evident that the accuracy of both the proposed JS losses is better than KL for all the noise level cases. Further, the



(a)



(b)

Figure 5: Accuracy on the test data set at different noise levels. 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.

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.

### 5.2.2. Histopathology

The results of the five runs of the KL loss and the proposed JS losses on the biased histopathology data set are compared in this section. It is evident from Fig. 6 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 improvement in performance 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 a performance evaluation metric that visually plots the classification ability of a binary classifier when the discrimination threshold is varied. For the present binary classification problem on histopathology data set the ROC curve is plotted in Fig. 7. The proposed JS losses perform better than the KL loss in terms of the area under the curve (AUC).




Figure 6: Accuracy on the test data set for five runs. 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.




Figure 7: Receiver operating characteristic curve depicting the classification performance of the three binary classifiers studied.

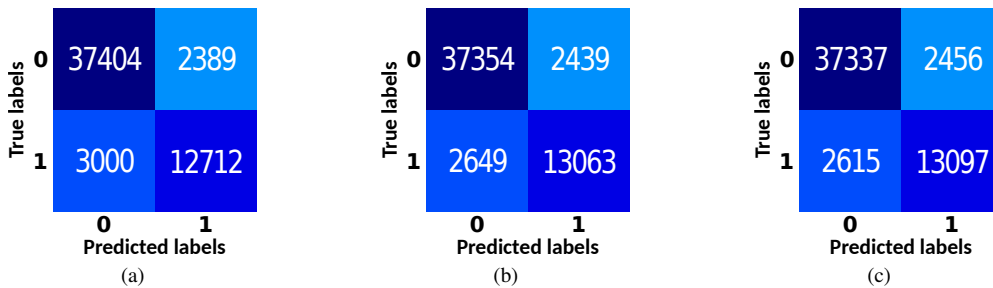


Figure 8: Confusion matrices for networks trained with (a) KL loss, (b) JS-G loss, and (c) JS-A loss on the histopathology data set.

The confusion matrices for various divergences are compared in Fig.8. The confusion matrices 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.



## 6. Conclusions

In this study, we introduced two novel JS divergence-based loss functions for Bayesian neural networks that extend the state-of-the-art KL divergence-based loss function. The proposed loss functions generalize well for data sets with noise and bias. To conclude this study, we summarise its main findings as follows.

Firstly, we have utilized JS Divergence to formulate two novel loss functions for Bayesian neural networks and have shown that the state-of-the-art KL divergence-based loss function is their special case.

Secondly, we have identified the reason for the better performance of the proposed losses compared to the state-of-the-art. Further, for the case of JS-G divergence we have derived the conditions to obtain better regularisation.

Thirdly, we have provided algorithms for the minimization of the proposed loss functions. In the algorithms, the JS divergences are evaluated by either closed-form expression or a Monte Carlo method. The relative merits of these two approaches are discussed.

Fourthly, we have shown the improvement in performance achieved by the proposed loss functions in terms of various performance metrics on data sets having bias and various degrees of noise.

This work would serve as a foundation for exploring other divergences for variational inferences in machine learning models.

## Appendix A. Monte Carlo estimates

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. A.9 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 A.9: Comparison of MC estimates and the closed form solution of JS-G divergence demonstrating the benefit of closed form solution.

## Appendix B. Data sets

### Appendix B.1. CIFAR-10

The CIFAR-10 data set consists of 60,000 images of size  $32 \times 32 \times 3$  belonging to 10 mutually exclusive classes. This data set is unbiased data set with each of the 10 classes having 6,000 images. Images were normalized using the min-max normalization technique. The entire data set was split into 60%-20%-20% for training-validation-testing respectively.

## Appendix B.2. Histopathology

The histopathology 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 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 C. Results of hyperparameter optimization

Results of the hyperparameter optimization (described in Section. 4.2) for the two data sets are presented in Tables. C.2 and C.3. In Table. C.2 Div denotes the divergence measure and LR is the learning rate.  $\lambda$  is taken as 1 for the JS-G divergence-based loss function throughout this work.

Table C.2: CIFAR10 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

Table C.3: Histopathology data set

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

## References

- [1] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. Van Der Laak, B. Van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Medical image analysis*, vol. 42, pp. 60–88, 2017.
- [2] D. Shen, G. Wu, and H.-I. Suk, "Deep learning in medical image analysis," *Annual review of biomedical engineering*, vol. 19, p. 221, 2017.

- [3] P. Thiagarajan, P. Khairnar, and S. Ghosh, "Explanation and use of uncertainty quantified by bayesian neural network classifiers for breast histopathology images," *IEEE Transactions on Medical Imaging*, vol. 41, no. 4, pp. 815–825, 2021.
- [4] M. Frank, D. Drikakis, and V. Charissis, "Machine-learning methods for computational science and engineering," *Computation*, vol. 8, no. 1, p. 15, 2020.
- [5] S. Kollmannsberger, D. D'Angella, M. Jokeit, L. Herrmann, *et al.*, *Deep Learning in Computational Mechanics*. Springer, 2021.
- [6] U. Yadav, S. Pathrudkar, and S. Ghosh, "Interpretable machine learning model for the deformation of multiwalled carbon nanotubes," *Physical Review B*, vol. 103, no. 3, p. 035407, 2021.
- [7] K. Choudhary, B. DeCost, C. Chen, A. Jain, F. Tavazza, R. Cohn, C. W. Park, A. Choudhary, A. Agrawal, S. J. Billinge, *et al.*, "Recent advances and applications of deep learning methods in materials science," *npj Computational Materials*, vol. 8, no. 1, pp. 1–26, 2022.
- [8] S. Pathrudkar, H. M. Yu, S. Ghosh, and A. S. Banerjee, "Machine learning based prediction of the electronic structure of quasi-one-dimensional materials under strain," *Physical Review B*, vol. 105, no. 19, p. 195141, 2022.
- [9] C. Beck, M. Hutzenhaler, A. Jentzen, and B. Kuckuck, "An overview on deep learning-based approximation methods for partial differential equations," *arXiv preprint arXiv:2012.12348*, 2020.
- [10] R. Matthey and S. Ghosh, "A novel sequential method to train physics informed neural networks for allen cahn and cahn hilliard equations," *Computer Methods in Applied Mechanics and Engineering*, vol. 390, p. 114474, 2022.
- [11] M. Buda, A. Maki, and M. A. Mazurowski, "A systematic study of the class imbalance problem in convolutional neural networks," *Neural Networks*, vol. 106, pp. 249–259, 2018.
- [12] H. D. Kabir, A. Khosravi, M. A. Hosen, and S. Nahavandi, "Neural network-based uncertainty quantification: A survey of methodologies and applications," *IEEE access*, vol. 6, pp. 36218–36234, 2018.
- [13] L. V. Jospin, H. Laga, F. Boussaid, W. Buntine, and M. Bennamoun, "Hands-on bayesian neural networks—a tutorial for deep learning users," *IEEE Computational Intelligence Magazine*, vol. 17, no. 2, pp. 29–48, 2022.
- [14] N. Tishby, E. Levin, and S. A. Solla, "Consistent inference of probabilities in layered networks: Predictions and generalization," in *International Joint Conference on Neural Networks*, vol. 2, pp. 403–409, IEEE New York, 1989.
- [15] J. Denker and Y. LeCun, "Transforming neural-net output levels to probability distributions," *Advances in neural information processing systems*, vol. 3, 1990.
- [16] E. Goan and C. Fookes, *Bayesian Neural Networks: An Introduction and Survey*, pp. 45–87. Cham: Springer International Publishing, 2020.
- [17] Y. Gal, "Uncertainty in deep learning," *PhD thesis, University of Cambridge*, 2016.
- [18] C. P. Robert, V. Elvira, N. Tawn, and C. Wu, "Accelerating mcmc algorithms," *Wiley Interdisciplinary Reviews: Computational Statistics*, vol. 10, no. 5, p. e1435, 2018.
- [19] G. E. Hinton and 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*, pp. 5–13, 1993.
- [20] D. Barber and C. M. Bishop, "Ensemble learning in bayesian neural networks," *Nato ASI Series F Computer and Systems Sciences*, vol. 168, pp. 215–238, 1998.
- [21] A. Graves, "Practical variational inference for neural networks," *Advances in neural information processing systems*, vol. 24, 2011.
- [22] J. M. Hernández-Lobato and R. Adams, "Probabilistic backpropagation for scalable learning of bayesian neural networks," in *International conference on machine learning*, pp. 1861–1869, Proceedings of Machine Learning Research, 2015.
- [23] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, "Weight uncertainty in neural network," in *International conference on machine learning*, pp. 1613–1622, Proceedings of Machine Learning Research, 2015.
- [24] J. Hensman, M. Zwieße, and N. D. Lawrence, "Tilted variational bayes," in *Artificial Intelligence and Statistics*, pp. 356–364, Proceedings of Machine Learning Research, 2014.
- [25] Y. Li and R. E. Turner, "Rényi divergence variational inference," *Advances in neural information processing systems*, vol. 29, 2016.
- [26] A. B. Dieng, D. Tran, R. Ranganath, J. Paisley, and D. Blei, "Variational inference via  $\chi$  upper bound minimization," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [27] J. Deasy, N. Simidjievski, and P. Liò, "Constraining variational inference with geometric jensen-shannon divergence," *Advances in Neural Information Processing Systems*, vol. 33, pp. 10647–10658, 2020.
- [28] F. Nielsen, "On the jensen–shannon symmetrization of distances relying on abstract means," *Entropy*, vol. 21, no. 5, p. 485, 2019.
- [29] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "beta-VAE: Learning basic visual concepts with a constrained variational framework," in *International Conference on Learning Representations*, 2017.
- [30] Paul Mooney, "Breast Histopathology Images." <https://www.kaggle.com/paultimothymooney/breast-histopathology-images>. 2017.
- [31] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, "Algorithms for hyper-parameter optimization," in *25th annual conference on neural information processing systems (NIPS 2011)*, vol. 24, Neural Information Processing Systems Foundation, 2011.
- [32] J. Bergstra, D. Yamins, and D. Cox, "Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures," in *International conference on machine learning*, pp. 115–123, Proceedings of Machine Learning Research, 2013.
- [33] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- [34] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, pp. 1097–1105, 2012.