




This work is published in IEEE Transactions on Medical Imaging.

P. Thiagarajan, P. Khairnar and S. Ghosh, "Explanation and Use of Uncertainty Obtained by Bayesian Neural Network Classifiers for Breast Histopathology Images," in IEEE Transactions on Medical Imaging, doi: 10.1109/TMI.2021.3123300.

Link to the paper can be found [here](#).

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

# Explanation and Use of Uncertainty Quantified by Bayesian Neural Network Classifiers for Breast Histopathology Images

Ponkrshnan Thiagarajan <sup>a,1</sup>, Pushkar Khairnar <sup>a,2</sup>, Susanta Ghosh <sup>b,3</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

Despite the promise of Convolutional neural network (CNN) based classification models for histopathological images, it is infeasible to quantify its uncertainties. Moreover, CNNs may suffer from overfitting when the data is biased. We show that Bayesian-CNN can overcome these limitations by regularizing automatically and by quantifying the uncertainty. We have developed a novel technique to utilize the uncertainties provided by the Bayesian-CNN that significantly improves the performance on a large fraction of the test data (about 6% improvement in accuracy on 77% of test data). Further, we provide a novel explanation for the uncertainty by projecting the data into a low dimensional space through a nonlinear dimensionality reduction technique. This dimensionality reduction enables interpretation of the test data through visualization and reveals the structure of the data in a low dimensional feature space. We show that the Bayesian-CNN can perform much better than the state-of-the-art transfer learning CNN (TL-CNN) by reducing the false negative and false positive by 11% and 7.7% respectively for the present data set. It achieves this performance with only 1.86 million parameters as compared to 134.33 million for TL-CNN. Besides, we modify the Bayesian-CNN by introducing a stochastic adaptive activation function. The modified Bayesian-CNN performs slightly better than Bayesian-CNN on all performance metrics and significantly reduces the number of false negatives and false positives (3% reduction for both). We also show that these results are statistically significant by performing McNemar's statistical significance test. This work shows the advantages of Bayesian-CNN against the state-of-the-art, explains and utilizes the uncertainties for histopathological images. It should find applications in various medical image classifications.

**Keywords:** Bayesian Convolutional Neural Networks, Breast Cancer, Histopathological Imaging, Machine Learning, Uncertainty Quantification, t-SNE.

---

## 1. Introduction

### 1.1. Histopathological imaging for Breast cancer

Breast cancer is the most common cause of cancer in women [1–3]. It has the highest incidence (43.3 per 100 000 population) than any other cancer and the highest mortality rate (15%) of all cancer deaths in women in 2012 [1]. Thus detection and diagnosis of breast cancer are vital in reducing the impact of the disease. Histopathological imaging is considered the gold

standard for breast cancer detection and diagnosis [4]. Histopathology is a diagnostic technique that involves microscopic examinations of tissues to study the sign of a disease. This method preserves the underlying architecture of the tissues thereby providing a significant contribution to the diagnosis of diseases. It is the only way to detect some of the diseases like lymphocytic infiltration of cancer. Histopathological images have a lot of information and structure which makes them highly reliable in the diagnosis of diseases especially in almost all types of cancer [5].

### 1.2. state-of-the-art machine learning algorithms for Histopathological imaging

With the advancement in digital imaging, computer-aided diagnosis (CAD) is studied extensively in the

---

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

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

<sup>3</sup>email: susantag@mtu.edu

literature [5–11]. Machine learning algorithms such as conventional neural networks and deep neural networks have shown tremendous potential in CAD applications. Feature extraction in a conventional neural network played a significant role in the predictions of the neural network. Various methods of feature extraction followed by classifiers such as support vector machine, K-nearest neighbor, decision tree, Naive Bayes amongst others were explored in the literature for classification and segmentation of breast cancer images [12–18]. Though these methods performed reasonably well in terms of accuracy, manual feature extraction was a bottleneck in improving the results.

Convolutional neural networks (CNNs) were the most successful in learning advanced features automatically from the input images. CNNs perform very well in classification, both binary [19–21] and multiclass classifications [4, 22], and metastasis detection [23] on breast cancer images. Apart from being able to perform classification, CNNs were also successfully implemented in problems involving segmentation [24–26] and detecting regions of interest that contain in-depth discriminatory information for classification in large whole-slide images [27–29]. CNNs also provided the scope to implement transfer learning. In transfer learning, the parameters learned for one data set are utilized to accelerate learning for another data set to solve a similar problem. Transfer learning CNN (TL-CNN) performs the best and is considered as the state-of-the-art classifier for breast cancer images [30–32]. A grand challenge on breast cancer histology images was conducted to advance the state-of-the-art in classifying these images. It was concluded that CNN was the most successful method to classify these breast cancer images [33]. The top performers [34–36] in this challenge used the architecture of an existing network such as Resnet [37], Densenet [38], Inception [39], VGG16 [40], etc and pre-trained these networks using ImageNet [41].

Uncertainty quantification provides a measure of trust in machine decisions and this metric is crucially important in medical applications [42]. Given the success of automatic classifiers for medical images, their uncertainty quantification is essential. However, there are only a few studies that report uncertainty quantification [43]. Although CNN performs better for image classification than other machine learning techniques in terms of accuracy, their parameters are deterministic and thus can not provide any measure of uncertainty in its predictions. In addition, predictions based on deterministic CNNs might provide incorrect results and without any estimate of confidence, these might lead to undesirable consequences. To evaluate the confidence of predic-

tions, uncertainty quantification is important [42]. The Bayesian CNN is an efficient state-of-the-art machine learning technique to quantify uncertainties [44–46]. It endows the uncertainty quantification capability in the CNN framework with nominal additional computational effort.

In a Bayesian-CNN, the weights and biases are random variables as opposed to deterministic variables used in CNN. Through the stochastic nature of the parameters, Bayesian-CNN captures the variability in the data set and computes the uncertainties in its predictions. Towards this, Blundell et al. [44] introduced an efficient algorithm called “Bayes by backprop”. This algorithm can be used to learn the probability distribution on weights and it was compatible with the conventional backpropagation scheme. They have also shown that it is possible to regularise the weights by minimizing a loss function known as the variational free energy and demonstrated that this method of regularisation showed performance comparable with dropout on MNIST classification. Shridhar et al. [45, 46] extended the “Bayes by backprop” algorithm to convolutional neural networks. The proposed Bayesian-CNN architecture was implemented for image classification, image super-resolution, and generative adversarial networks. This method has shown performance as good as conventional CNN and in addition, provided uncertainty measures and regularisation. Kendall et al. [47] provided a Bayesian deep learning framework combining aleatoric and epistemic uncertainties. The authors improved the model’s performance by 1 to 3% over its deterministic counterpart by reducing the effect of noisy data. The trade-offs between modeling aleatoric and epistemic uncertainties were also studied. Kwon et al. [48] proposed a method of quantifying uncertainties in classification problems using Bayesian neural networks. Bayesian networks are used in the literature to perform active learning. Uncertainties were utilized to add new samples in the training data to improve performance by Gal et al [43]. An accurate, reliable active learning framework was introduced by Raczkowski et al [49] utilizing the variational dropout-based uncertainty measures.

### *1.3. Present work: Quantification, explanation, and use of uncertainty, and the Modified Bayesian-CNN*

Bayesian-CNNs are gaining popularity due to their ability to provide uncertainties associated with the predictions. Yet, there has been no known work to the authors’ knowledge that demonstrates the advantages of Bayesian-CNN over CNN to classify histopathological

images. In this work, at first, we perform uncertainty quantification for the classification of breast histopathological images and show the advantages of Bayesian–CNN over CNN. We found that the Bayesian–CNN improves accuracy and reduces overfitting in comparison to CNN in addition to quantifying uncertainties. To exploit the advantages of uncertainty quantification, we have developed a novel technique to divide the data into low and high uncertainty subsets. We show significant improvement in accuracy in the majority of data that has low uncertainty. A small fraction of the data that has higher uncertainty can be referred to the experts. Thus, uncertainty provides an avenue to determine which images require human intervention. We explain the reason for this high uncertainty in a subset of data by projecting the data into a low dimensional space (latent space) using t-distributed stochastic neighbor embedding (t-SNE). We choose a three-dimensional latent space to facilitate the visualization of data. To the best of the authors’ knowledge, such an explanation of uncertainty is not reported in the literature.

We propose a new model to extend the Bayesian–CNN to further improve its performance and name it as the modified Bayesian–CNN. Adaptive activation functions have shown improvement in performance for both deep neural networks and convolutional neural networks [50–52]. However such adaptive activation functions are not applied to Bayesian Neural networks. The novelty of the proposed adaptivity lies in the fact that the learnable parameter is random as opposed to deterministic parameters reported in the literature. Thus the activation function is now sampled from an ensemble of such functions based on the probability distribution of parameters learned from the data.

The rest of the paper is organized as follows: in section 2 the data set and the methods used in this work are described; results obtained in the work and their analysis are presented in section 3 followed by conclusions in section 4.

## 2. Data set and methods

### 2.1. Data acquisition

The breast histopathological images used in this work are from a publicly available data set [53]. These are images containing regions of Invasive Ductal Carcinoma which is the most common subtype of all breast cancers [53]. These regions are separated from the whole slide images by pathologists. 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  $50 \times 50$  pixels were extracted (198,738 negative and 78,786 positive) and provided as a data set for classification.

### 2.2. Data Preparation

The data set consists of two classes: positive (1) and negative (0). The classification of the data was carefully done by experienced pathologists providing the ground truth for training. 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 image size is  $3 \times 50 \times 50$  (D×H×W), where D is the depth (color channels), H is the height, and W is the width. The images were shuffled and converted from uint8 to float format for normalizing. For most of the images, the majority of the pixel values are greater than 200. This makes the computation extremely expensive especially with the large size of data. Further, it leads to problems such as singularity/gradient explosion during the evaluation of loss function. In addition, the original images are such that the information that needs to be highlighted has low pixel values. To solve these problems, 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 was carried out as  $p_n = (255 - p)/255$ . Where  $p$  is the original pixel value and  $p_n$  is the pixel value after normalization and complement.

### 2.3. Methods

The mathematical background for Bayesian–CNN and its uncertainty quantification [45], and for t-SNE [54] is revisited in this section. The implementation of stochastic adaptive activation on Bayesian–CNN is also described.

#### 2.3.1. Bayesian Neural Networks

The neural network can be defined as a probabilistic model  $P(y|x, w)$ , where  $x \in \mathbb{R}^p$  is an input to the network,  $y \in \mathcal{Y}$  is the output. Where  $\mathcal{Y}$  is a set of all possible outputs. The network consists of a set of trainable random parameters  $w$ . This set of parameters  $w$  is learned using a complete Bayesian approach. In this approach given training data  $D$ , the posterior distribution of weights  $P(w|D)$  is calculated using Bayesian inference for neural networks which involves marginalization over all possible values of  $w$ . Once the posterior distribution  $P(w|D)$  is obtained, predictions on the unseen data are obtained by taking expectations on the

predictive distributions. The predictive distribution of an unknown label  $\widehat{y}$  of a test data item  $\widehat{x}$  is given by  $P(\widehat{y}|\widehat{x}) = \mathbb{E}_{P(w|D)}[P(\widehat{y}|\widehat{x}, w)] = \int_{\Omega_w} P(\widehat{y}|\widehat{x}, w)P(w|D)dw$ .

To estimate the posterior distribution we use Bayes' Rule which gives:

$$P(w|D) = \frac{P(D|w)P(w)}{P(D)} \quad (1)$$

The term  $P(D|w)$  is the likelihood of the training data ( $D$ ) given a parameter setting ( $w$ ). Assuming each training data to be independent and identically distributed, the above term becomes the product of likelihood,  $P(D|w) = \prod_{n=1}^N P(y^n|w, x^n)$  where,  $(x^n, y^n)$  are the training data item and known label respectively.

The prior  $P(w)$  is our belief about the distribution of weights without seeing the data. The term  $P(D)$  in the equation (1) is intractable which makes the posterior distribution  $P(w|D)$  intractable. The term  $P(D)$  involves marginalization over the weight distribution:  $P(D) = \int_{\Omega_w} P(D|w)P(w)dw$ .

For variational inference [44], the posterior distribution  $P(w|D)$  which is intractable is approximated with a tractable simpler distribution over the model weights  $q(w)$ . If we assume our weight distributions to be Gaussian, this simpler tractable distribution has variational parameters  $\theta$ , where  $\theta = (\mu, \sigma)$ . We fit the variational parameters  $\theta$  such that  $q(w|\theta) \approx P(w|D)$ . The variational posterior  $q(w|\theta)$  is used instead of the intractable posterior  $P(w|D)$  for the inference,  $P(\widehat{y}|\widehat{x}) = \mathbb{E}_{q(w|\theta)}[P(\widehat{y}|\widehat{x}, w)]$ . Estimating  $q(w|\theta) \approx P(w|D)$  we can say that we have learned the distribution of weights given training data.

To make the variational posterior  $q(w|\theta)$  and the true posterior  $P(w|D)$  similar, we minimize the KL divergence between them:

$$\begin{aligned} KL[q(w|\theta)||P(w|D)] &= \mathbb{E}_{q(w|\theta)} \left[ \log \frac{q(w|\theta)}{P(w|D)} \right] \\ &= \int q(w|\theta) \log \frac{q(w|\theta)}{P(w)} dw + \int q(w|\theta) \log(P(D)) dw \\ &\quad - \int q(w|\theta) \log(P(D|w)) dw \end{aligned} \quad (2)$$

The term  $\log(P(D))$  makes the above equation intractable. Although intractable, it is constant. Therefore, minimizing the KL divergence can be defined as:

$$\begin{aligned} \min(KL[q(w|\theta)||P(w|D)]) &= \min(KL[q(w|\theta)||P(w)] \\ &\quad - \mathbb{E}_{q(w|\theta)}[\log(P(D|w))]) \end{aligned} \quad (3)$$

The term  $(KL[q(w|\theta)||P(w)] - \mathbb{E}_{q(w|\theta)}[\log(P(D|w))])$  is called the Variational Free Energy (VFE) which is to be

minimized. Exact minimisation of the cost function is computationally expensive, therefore it is approximated through a Monte Carlo sampling procedure as follows:

$$\mathcal{F}(D, \theta) \approx \sum_{i=1}^n [\log q(w^{(i)}|\theta) - \log P(w^{(i)}) - \log P(D|w^{(i)})] \quad (4)$$

Where,  $w^{(i)}$  denotes the  $i^{th}$  Monte Carlo sample drawn from the variational posterior  $q(w^{(i)}|\theta)$ . Each term in  $\mathcal{F}(D, \theta)$ , the loss function that needs to be minimized, also known as the cost function has weights that are drawn from the variational posterior. For further details on the cost function, readers are referred to [44, 45]. To implement and compute the above cost we need three terms: 1) log of variational posterior, 2) log of prior (Gaussian or scale mixture), 3) log-likelihood of the data.

The variational posterior ( $q$ ) is composed of independent Gaussian distribution for each parameter. The sample of weights are obtained by sampling the unit Gaussian, shifting and scaling by the mean  $\mu$  and a standard deviation  $\sigma$  respectively. To ensure that the standard deviation is always non-negative, it is expressed as  $\sigma = \text{softplus}(\rho) = \log(1 + \exp(\rho))$ , point-wise. Parameters,  $w$ , of the variational posterior,  $q$ , can be re-written in terms of a parameter-free noise as  $[w = \mu + \log(1 + \exp(\rho)) \circ \varepsilon]$  or  $[w = \mu + \sigma \circ \varepsilon]$ . Where,  $\circ$  is point-wise multiplication,  $\varepsilon$  is a sample drawn from unit normal [44]. The steps for optimization are presented in supplementary document.

### 2.3.2. Bayesian Convolution Neural Networks

The parameters of CNNs are filters or kernels which are to be learned during training. In the case of a Bayesian-CNN these kernels are represented by probability distributions as shown in Fig. 1.




Figure 1: Schematics of the convolution of Bayesian-CNN showing probability distributions of random parameters.

During the training of Bayesian-CNN the reparameterization trick is applied on these filters or kernels, which are of shape  $h \times w \times d$ . They are sampled from the

variational posterior  $q(w|\theta)$  using the following equation:

$$w_{h,w,d} = \mu_{h,w,d} + \log(1 + \exp(\rho_{h,w,d})) \circ \varepsilon_{h,w,d} \quad (5)$$

where  $h$  is the height,  $w$  is the width and  $d$  is the depth of the filter and  $\circ$  represents point-wise multiplication. After the sampling from the variational posterior, similar optimization steps as explained in the previous section are followed.

### 2.3.3. Uncertainty Quantification

The uncertainty quantification becomes extremely important when dealing with the applications related to autonomous vehicles, medical imaging, etc. Bayesian deep learning makes it possible to quantify the uncertainties in the prediction as we have probability distribution over weights. Taking an expectation of the predictive posterior probability distribution:  $\mathbb{E}_{q(w|\theta)}[P(\hat{y}|\hat{x}, w)]$  gives us the most probable prediction of the unknown data  $\hat{x}$ . The variance of the predictive posterior probability distribution:  $\text{Var}_{q(w|\theta)}[P(\hat{y}|\hat{x}, w)]$  quantifies the uncertainties. There are two types of uncertainties: the Aleatoric and the Epistemic uncertainty. The variance of the predictive posterior probability distribution is the sum of both these uncertainties:

$$\text{Var}_{q(w|\theta)}[P(\hat{y}|\hat{x}, w)] = \text{aleatoric} + \text{epistemic} \quad (6)$$

The aleatoric uncertainty corresponds to the noise in the data set whereas the epistemic uncertainty corresponds to the variability of the model developed[45]. One of the promising approaches to quantify the uncertainties is explained in [48]. The uncertainties are obtained from the variance of the predictive posterior probability distribution,

$$\begin{aligned} \text{Var}_{q(w|\theta)}[P(\hat{y}|\hat{x}, w)] &= \mathbb{E}_{q(w|\theta)} [(y - \mathbb{E}[y])^2] \\ &= \int_{\Omega_w} \left[ \left[ \text{diag}(\mathbb{E}_{P(\hat{y}|\hat{x}, w)}[\hat{y}]) - \mathbb{E}_{P(\hat{y}|\hat{x}, w)}[\hat{y}] \mathbb{E}_{P(\hat{y}|\hat{x}, w)}[\hat{y}]^T \right] \right. \\ &\quad \left. q(w|\theta) dw \right] + \int_{\Omega_w} \left[ \left[ \mathbb{E}_{P(\hat{y}|\hat{x}, w)}[\hat{y}] - \mathbb{E}_{q_\theta(\mathcal{Y}|\hat{x})}[\hat{y}] \right] \right. \\ &\quad \left. \left[ \mathbb{E}_{P(\hat{y}|\hat{x}, w)}[\hat{y}] - \mathbb{E}_{q_\theta(\mathcal{Y}|\hat{x})}[\hat{y}] \right]^T q(w|\theta) dw \right] \end{aligned} \quad (7)$$

The above expression is the sum of aleatoric and the epistemic uncertainties which is derived from the variant of the law of the total variance [48]

The first term of Eq. 7 is defined as the aleatoric uncertainty and the second term is the epistemic uncertainty. Due to the integral term in Eq. 7, it is intractable and requires approximations as follows. The equation of the aleatoric uncertainty is

defined as:  $\frac{1}{N} \sum_{n=1}^N \text{diag}(\hat{p}_n) - \hat{p}_n \hat{p}_n^T$  where,  $\hat{p}_n = p(\hat{w}_n) = \text{softmax}\{f^{\hat{w}_n}(\hat{x})\}$ . And the equation of the epistemic uncertainty is defined as:  $\frac{1}{N} \sum_{n=1}^N (\hat{p}_n - \bar{p})(\hat{p}_n - \bar{p})^T$  where,  $\hat{p}_n = p(\hat{w}_n) = \text{softmax}\{f^{\hat{w}_n}(\hat{x})\}$  and  $\bar{p} = \sum_{n=1}^N \frac{\hat{p}_n}{N}$ . The overall equation for the variance is:

$$\begin{aligned} \text{Var}_{q(w|\theta)}[P(\hat{y}|\hat{x}, w)] &= \left( \frac{1}{N} \sum_{n=1}^N \text{diag}(\hat{p}_n) - \hat{p}_n \hat{p}_n^T \right) \\ &\quad + \left( \frac{1}{N} \sum_{n=1}^N (\hat{p}_n - \bar{p})(\hat{p}_n - \bar{p})^T \right) \end{aligned} \quad (8)$$

The variance of the predictive distribution can be calculated by Eq. 8 which provide us with the confidence of the network in making predictions for a given image.

### 2.3.4. Adaptive activation

The weights and bias of a perceptron in a neural network perform a linear transformation of the inputs. The output of this linear transformation is passed to an activation function that determines if a particular perceptron is activated for a given input. A non-linear activation function is a key component of a neural network, which enables it to learn complex functions with a small number of perceptrons. However, this nonlinearity of the activation function is known to introduce problems such as exploding or vanishing gradients. Thus, there is a trade-off between the learning capabilities vs training complexities of a perceptron, which can be optimized. To achieve this, we modify the Bayesian-CNN by introducing a learnable stochastic activation function that adapts to the training data. Such adaptivity in the context of other machine learning models has been reported in the literature [50, 52], in which the learnable parameter is deterministic. In this work, we introduce a probabilistic learning parameter for the adaptivity of Bayesian-CNN. This adaptive activation function contains a probabilistic parameter that is learned during the training of the neural network. To this end, a trainable probabilistic hyperparameter ( $\alpha$ ) is introduced in the activation functions of a Bayesian-CNN and the resulting network is the modified Bayesian-CNN. The details of the proposed adaptive activation function are given below.

$$\sigma(\alpha f_k(x^{k-1})) \quad (9)$$

where  
 $f_k(x^{k-1}) = \bar{w}^k x^{k-1} + b^k$   
 $\sigma$  is the activation function  
 $\alpha$  is the probabilistic hyperparameter that can be trained  
 $\bar{w}$  and  $b$  are the weights and bias of the  $k^{\text{th}}$  layer and

$x^{k-1}$  is the output from the previous layer of the neural network.

To incorporate the stochastic adaptive activation, the loss function  $\mathcal{F}$  (Eq. 4) is modified to include the additional stochastic parameter  $\alpha$ . Specifically, the set of trainable network parameters  $w = \{\bar{w}, b\}$  in  $\mathcal{F}$  is extended to  $W = \{\bar{w}, b, \alpha\}$ . Thus the modified loss function  $\tilde{\mathcal{F}}(D, \theta)$  is given as

$$\tilde{\mathcal{F}}(D, \theta) \approx \sum_{i=1}^n [\log q(W^{(i)}|\theta) - \log P(W^{(i)}) - \log P(D|W^{(i)})]$$

The parameter  $\alpha$  is learned in the same way as the parameters  $w$  as described in Sec. 2.3.1. A prior distribution (a mixture of Gaussian) for the parameter  $\alpha$  is assumed and the posterior distribution is obtained via Bayes' rule. This posterior distribution is approximated by a Gaussian variational posterior  $q_\alpha$  with mean  $\mu_\alpha$  and standard deviation  $\sigma_\alpha$ . The parameters of this variational posterior are updated through gradient descent as,

$$\mu_\alpha^{m+1} = \mu_\alpha^m - \eta \nabla_{\mu_\alpha} \tilde{\mathcal{F}}^m \text{ and}$$

$$\sigma_\alpha^{m+1} = \sigma_\alpha^m - \eta \nabla_{\sigma_\alpha} \tilde{\mathcal{F}}^m$$

Where,  $\eta$  is the learning rate. Refer to section II of the supplementary document for further details.

In this work, an adaptive ReLU given by  $\sigma(\gamma) = \max(0, \gamma x)$  is used for the fully connected layers.

The networks were trained, validated, and tested on GPUs using Apache MXNet on Python 3. Details of the computational platform are given in the supplementary document.

### 2.3.5. *t*-distributed stochastic neighbor embedding (*t*-SNE)

t-SNE [54] is an unsupervised non-linear dimensionality reduction technique that helps in visualizing high dimensional data by mapping it into a low dimensional space. In this method, as a first step, a conditional probability ( $p_{j|i}$ ) is assigned between data points  $x_i$  and  $x_j$  in a high dimensional space. This conditional probability  $p_{j|i}$  represents the similarity of data points such that similar points in the high dimensional space have a higher probability than dissimilar points. One such measure of similarity is the Euclidean distance between a pair of data points. In the second step, a similar conditional probability ( $q_{j|i}$ ) is defined for the low dimensional counterpart  $y_i$  and  $y_j$ . If the low dimensional representation of the data correctly models similarity

between the points in the high dimensional space, the two conditional probabilities  $p_{j|i}$  and  $q_{j|i}$  will be equal. Thus, in the third step, by minimizing the KL divergence between the two conditional probabilities a low dimensional map of the data points is obtained via t-SNE. In this work, we explain the high uncertainty of images by projecting them into a latent space via t-SNE.

## 3. Results and analysis

This section presents classification results for histopathological data, its uncertainty quantification obtained through the proposed modified Bayesian-CNN (Section 2.3). It also provides an explanation of the predicted uncertainty via t-SNE. Further, the proposed model is analyzed and juxtaposed with Bayesian-CNN and transfer learning CNN (TL-CNN).

### 3.1. Hyperparameter optimization

Hyperparameters and architecture for all the networks considered here are chosen through hyperparameter optimization. A Tree-structured Parzen Estimator (TPE) algorithm is used which is a sequential model-based optimization approach [55]. 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 [56] 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 the Data preparation section 2.2. An optimization is performed to maximize the validation accuracy for different configuration and hyperparameter settings of the network.

The architecture of TL-CNN follows VGG-16 [40]. Architectures for Bayesian-CNN and modified Bayesian-CNN are shown in Fig. 2. In the case of modified Bayesian-CNN, the ReLU function after the FC 1 layer (see Fig.2) is changed into a stochastic-adaptive ReLU. Using stochastic-adaptive Relu in convolution layers as well increased the computational cost without any significant improvement in performance. More details are given in the supplementary document. The number of filters for convolution layers and the number of neurons in the fully connected layers are obtained through hyper-parameter optimization. The results of the optimisation are presented in Tables. 1 and 2.




Figure 2: Architecture for Bayesian-CNN and modified Bayesian-CNN. Conv represents convolution layers and FC represent fully connected layers. Number of neurons in FC 2 is fixed as 2 for binary classification. Optimized kernel size of all convolution layers is fixed as  $3 \times 3$  with stride 1. Optimized kernel size of max pool layers is  $2 \times 2$  with stride 2.

Table 1: Selected hyperparameters

Network	Learning rate	Batch Size	Layers learnt
TL-CNN	0.001	64	2
Bayesian-CNN	0.001	128	All
Modified Bayesian-CNN	0.0001	64	All

Table 2: Selected hyper parameters for network architecture

Network	con1	con2	con3	con4	con5	con6	FC1
Bayesian-CNN	16	32	32	64	128	256	512
modified Bayesian-CNN	32	64	64	128	128	128	256

### 3.2. Training, validation, and testing of networks

The networks are trained using the images as inputs and the corresponding known classes as outputs. The cross-entropy loss and the variational free energy loss are minimized for TL-CNN and Bayesian networks respectively. The training is done until the loss converges or the validation accuracy starts to decrease. In the following the details of training, validation, and testing for TL-CNN, Bayesian-CNN, and the modified Bayesian-CNN are described.

#### 3.2.1. Convolutional neural networks

Transfer learning is the state-of-the-art for CNNs and has consistently performed the best for medical image classification [33, 57]. In transfer learning, the learned parameters for one data set are utilized for other data sets to perform a similar task. Here we use the weights of VGG16 [40] architecture trained on Imagenet data set [41] for classification and perform transfer learning to classify the histopathological data. The accuracy for training and validation for each epoch is shown in Fig. 3 for TL-CNN. The accuracy increases with the epoch for the training data whereas the accuracy remains constant

for the validation data after the first few epochs. This problem is known as overfitting. Thus, TL-CNN fails to generalize its predictions to unknown validation data and overfits the training data. Beyond 40 epochs, the validation accuracy remains almost a constant while the training accuracy is still increasing, which is long before the training loss has converged.




Figure 3: Training and validation accuracy of TL-CNN. We can clearly see the problem of overfitting as the training accuracy increases while the validation accuracy does not improve. The results presented for the TL-CNN henceforth are for the maximum validation accuracy.

To alleviate the overfitting of the networks, regularization strategies are used, such as  $L_1$  or  $L_2$  regularization, cross-validation, early stoppage, and dropout [58–60]. However, these may require human intervention during training (such as parameter tuning, determination of early stoppage) or may lead to loss of information due to drop out [59].

In Bayesian networks the prior distribution on weights  $p(w)$  introduces regularisation of weights automatically. A Gaussian prior yields  $L_2$  regularisation on the weights and a Laplace prior yields  $L_1$  regularisation on the weights. In addition, due to the stochastic nature of the parameters, an average across multiple models is computed during training which introduces a regularisation effect on the network. However, similar to the selection of a regularisation parameter in deterministic networks, the selection of prior parameters in Bayesian networks is challenging. Herein we show that through a Bayesian framework the limitations of overfitting is overcome automatically.

#### 3.2.2. Bayesian-CNN and the modified Bayesian-CNN

The accuracy of predictions for both training and validation data of the Bayesian and modified Bayesian-CNN is shown in Fig. 4 and Fig. 5 respectively. It is evident that the problem of overfitting is eliminated by using a Bayesian network as the accuracy of predictions



for both training and validation data increases with an increase in the number of epochs.




Figure 4: Training and validation accuracy of Bayesian-CNN. We can see that both training and validation accuracy increases with epoch which implies automatic regularisation.




Figure 5: Training and validation accuracy of modified Bayesian-CNN.

The values of the accuracy of training, validation, and test data upon the convergence of the loss function are provided in Table 3. It is seen that the problem of overfitting drastically reduces the accuracy for validation and test data in a TL-CNN whereas this problem is not present for Bayesian-CNN and the modified Bayesian-CNN. The training, validation, and test accuracies of the modified Bayesian-CNN is slightly higher than the Bayesian-CNN when the loss function is converged. We found that the modified Bayesian-CNN requires about the same time as that of Bayesian-CNN for the same batch size but twice the computational cost per epoch during the training as compared to TL-CNN. More details on the training times are presented in the supplementary document. We have performed a comparative study of the Bayesian-CNN and the modified Bayesian-CNN on the MNIST data [61]. We found that the modified Bayesian-CNN shows a similar improvement in performance on the accuracy for the MNIST data set. Further details are presented in the supplementary document.

Table 3: A Comparison of TL-CNN, Bayesian-CNN and modified Bayesian-CNN

Network	Training Accuracy	Validation Accuracy	Testing Accuracy
TL-CNN	0.9078	0.8691	0.8676
Bayesian-CNN	0.8848	0.8817	0.8801
modified Bayesian-CNN	0.8898	0.8856	0.8837

### 3.3. Quality assessment metrics

The performance of the classifier is presented using the confusion matrix in Fig. 6. Details of the confusion matrix is provided in the supplementary document.




Figure 6: Comparison of confusion matrices for a) TL-CNN, b) Bayesian-CNN and modified Bayesian-CNN

In addition to the improvement of accuracy, the modified Bayesian-CNN remarkably reduces the number of false-negative predictions as seen in Fig. 6, which is a significant achievement. CNN is known to perform poorly for biased data set [62]. Given the data set used in this work is highly biased to the negative class, the TL-CNN model fails to overcome this bias yielding more negative predictions. We demonstrate that the modified Bayesian-CNN, overcomes this limitation of the CNN and provides an unbiased result. The modified Bayesian-CNN reduces the false negative and false positive predictions by about 13.7% and 10.4% respectively as compared to TL-CNN. The modified Bayesian-CNN reduces both the false negative and false positive by 3% as compared to Bayesian-CNN as seen in Fig. 6, demonstrating unbiased prediction. This is a remarkable advancement given that medical imaging data sets are usually biased towards negative labels. This significant improvement in performance is achieved despite having only 1.11 million parameters as compared to 134.33 million for TL-CNN.

The improvement in the performance of modified Bayesian-CNN is due to the adaptive nature of the activation function. The adaptive activation function with a probabilistic learning parameter works better because

through it the network trains an ensemble of activation functions where each activation function has its weights drawn from a probability distribution which is learned from the data. Learning the probability distribution for the activation parameter from the training data reduces the bias in the model by eliminating the assumption in the functional form of the activation.

To quantitatively assess the performance of the proposed binary classification the following metrics are used in addition to the confusion matrix: Accuracy, Precision, Recall, F1-score, Cohen’s-kappa (CK) coefficient [63]. Cohen’s Kappa coefficient varies between 0-1 where: 0 – agreement is equivalent to chance and 1.0 – perfect agreement. Improvement in Cohen’s Kappa coefficient is achieved by using the modified Bayesian-CNN as seen from table 4 which implies that it works better than both Bayesian-CNN and TL-CNN for an unbalanced data set.

The performance of the neural network classifiers is evaluated based on all of the above metrics and the results are provided in table 4. The modified Bayesian-CNN performs slightly better than Bayesian-CNN and much better than TL-CNN on all the performance metrics for this data set.

Table 4: Performance metrics of TL-CNN and BCNN

Network	Recall	Precision	F1-score	CK
TL-CNN	0.7552	0.7722	0.7636	0.6717
Bayesian-CNN	0.7823	0.7917	0.7870	0.7036
modified Bayesian-CNN	0.7888	0.7981	0.7934	0.7125

### 3.3.1. Statistical significance

In order to ensure that the difference in the results obtained by the Bayesian-CNN and the modified Bayesian-CNN is not by chance, we performed a statistical significance test. Importance of statistical significance tests is highlighted in [64, 65]. The McNemar’s test [66] is performed and the results show that the difference in performance is statistically significant ( $p < 0.001$ ). Therefore, the Bayesian-CNN and the adaptive Bayesian CNN are statistically different models. The details and the results of the statistical significance test are provided in Sec. VIII of the supplementary document.

### 3.4. Uncertainty quantification

Uncertainty quantification informs how confident the network is in its predictions, which is crucial for medical applications. Thus in addition to the accuracy met-

rics, uncertainty is an important measure that can be estimated naturally through the Bayesian approaches. On the contrary, the parameters of CNN are deterministic and large in number (134 million) that makes its uncertainty quantification infeasible. The parameters of the (modified) Bayesian-CNN are probability distributions whose variance provides an estimate of the uncertainties associated with the predictions. The larger the variances the more uncertain the model is on its predictions. The change of the probability distribution of a parameter for the modified Bayesian-CNN is shown for different epochs in Fig. 7. As the training continues, the standard deviation decreases increasing the confidence in predictions (see the supplementary document).




Figure 7: Probability density of a parameter for different epochs of the modified Bayesian-CNN.




Figure 8: Sample images having negative labels and their uncertainties: a) TN [ $A = 9.25 \times 10^{-3}$ ,  $E = 6.50 \times 10^{-6}$ ], b) TN [ $A = 1.8 \times 10^{-2}$ ,  $E = 5.65 \times 10^{-3}$ ], c) TN [ $A = 4.44 \times 10^{-3}$ ,  $E = 2.53 \times 10^{-6}$ ], d) TN [ $A = 2.77 \times 10^{-1}$ ,  $E = 5.42 \times 10^{-3}$ ], e) FP [ $A = 3.19 \times 10^{-1}$ ,  $E = 4.59 \times 10^{-3}$ ], f) FP [ $A = 3.74 \times 10^{-1}$ ,  $E = 5.57 \times 10^{-3}$ ]. Aleatoric and epistemic uncertainties are given by A and E respectively.





Figure 9: Sample images having positive labels and their uncertainties: a) TP [ $A = 7.0 \times 10^{-2}$ ,  $E = 2.60 \times 10^{-4}$ ], b) TP [ $A = 6.60 \times 10^{-2}$ ,  $E = 1.74 \times 10^{-4}$ ], c) TP [ $A = 1.14 \times 10^{-1}$ ,  $E = 4.91 \times 10^{-4}$ ], d) TP [ $A = 4.72 \times 10^{-1}$ ,  $E = 2.03 \times 10^{-2}$ ], e) FN [ $A = 4.64 \times 10^{-1}$ ,  $E = 9.37 \times 10^{-3}$ ] f) FN [ $A = 4.74 \times 10^{-1}$ ,  $E = 6.26 \times 10^{-3}$ ].

The uncertainty values associated with an individual image denote the confidence with which the network predicts the class label for that image. The aleatoric uncertainty corresponds to the noise in the data which cannot be reduced by improving the model. The epistemic uncertainty represents the variability due to the model, which can be improved. A set of sample images from the test data set, their predicted class and their aleatoric and epistemic uncertainty in prediction are shown in Fig. 8 and 9. The images with negative and positive ground truth labels are shown in Fig. 8 and Fig. 9 respectively. The important features apparent from these images are the texture density and color. It appears that the positive class has darker and denser images. It seems that the images with high epistemic uncertainty, Fig. 8(d,e,f) and 9(d,e,f) has features from both the classes. Therefore, the model gives high epistemic uncertainty to those images indicating mixed features or mislabeling. Thus, the Bayesian approach provides an avenue to identify which images should be referred back to a human expert.


### 3.5. Use of uncertainty to improve performance

We utilize the uncertainty to further improve the performance on a subset of data. The aleatoric uncertainty of the test set is normalised by the maximum and minimum values of aleatoric uncertainty of the training set. For every threshold of aleatoric uncertainty, we divide the test data into two subsets having uncertainty lower and higher than the threshold. The ratio of the number of images in the low uncertainty subset is plotted against the uncertainty-threshold used to create the subset in Fig. 10(a). The accuracy corresponding to each of

this subset are also plotted against the threshold. The number of false-negative and false-positive predictions are plotted for different thresholds of aleatoric uncertainty in Fig. 10(b). For instance, if the threshold value is 0.6, the low uncertainty subset contains 77% of the test data and its accuracy is 94.6% which is about 6% improvement in accuracy over the entire test data set. The remaining 23% of the test data, which has uncertainty higher than this threshold (0.6) may be referred to a human expert for a more accurate prediction. Furthermore, for this subset, the percentage of false negatives reduces from 6.2% for the entire data set to 2.7% and the percentage of false positives reduces from 5.8% for the entire data set to 1.5%, which are significant improvements. We found that normalized uncertainty values lesser than 0.8 provided significant improvement in performance for this data set. For other data sets this threshold on uncertainty can be set by looking at change in slope of the accuracy vs uncertainty plot. A similar study is done using the epistemic uncertainty and the results are presented in the supplementary document. Therefore, we show that uncertainty quantification can be used to improve the performance significantly for a subset of test data.



(a) Fraction of test data having uncertainty below the threshold and their accuracy plotted against the uncertainty-threshold.




(b) Number of false predictions in the low uncertainty subset plotted against the uncertainty-threshold.


Figure 10: Performance improvements for a subset of test data that has uncertainty less than the uncertainty-threshold. Dotted Lines represents two thresholds of aleatoric uncertainty.

### 3.6. Explanation of uncertainty by visualization in low dimension


In the following, we provide an explanation of uncertainty by visualizing the test data in a low dimensional space (latent space). A dimensionality reduction is performed on the data using t-SNE to project the 7500-dimensional feature into 3 dimensions (3D).



(a) Low uncertainty



(b) Medium uncertainty



(c) High uncertainty

Figure 11: Low dimensional representation of the test data for three subsets using t-SNE. blue: negative class, red: positive class

The epistemic uncertainty ( $E$ ) evaluated through the Bayesian approach is normalized to range between  $[0 - 1]$ . The test data set is divided into three subsets: low ( $E \leq 0.01$ ), medium ( $0.01 < E \leq 0.1$ ) and high uncertainty ( $E > 0.1$ ). t-SNE is performed for these three subsets. Low dimensional representations obtained via t-SNE are presented in Fig. 11 where each data point represents a test image. Fig. 11 explains the uncertainty associated with the individual images. The low uncertainty images have a clear separation between the positive and negative classes in the latent space. Thus the model classifies these images with high confidence. The images with medium uncertainty are clustered at different locations in the latent space however have significant overlap. Therefore the model classifies these images with lesser confidence. For the high uncertain images, the negative and the positive classes are not at all separable in the latent space. Thus the model cannot distinguish between the two classes for this subset, yielding poor accuracy and low confidence.

## 4. Conclusions

In this study, we have quantified the uncertainty in the classification of histopathological images through Bayesian-CNN and compared its performance against the state-of-the-art. We have explained the uncertainty and utilized the uncertainty to improve the performance. To conclude this study, we summarize its main findings.

Firstly, we have shown that Bayesian-CNN performs better than the state-of-the-art transfer learning CNN on all performance metrics and provides uncertainties in classifying breast histopathological images. Bayesian-CNN improves the test accuracy by 1.2% and reduces the false negative and false positive by 11% and 7.7% respectively as compare to TL-CNN, despite having only 1.86 million parameters as compared to 134.33 million for TL-CNN.

Secondly, we explained the uncertainties in the test data by projecting them into a low-dimensional space. This low-dimensional projection enables visual interpretation of the data. The data structure in the low dimensional feature space reveals that images with high uncertainty are not separable whereas the images with low uncertainty are clearly separable.

Thirdly, using different uncertainty thresholds, we demonstrate that the accuracy of Bayesian networks can be significantly improved on a large fraction of test data (6% improvement on 77% test data).

Fourthly, we have proposed a novel stochastic-adaptive activation enabled Bayesian-CNN and call it

modified Bayesian–CNN. In the proposed model we introduce a probabilistic learnable activation function that adapts to the training data to improve predictive capabilities. We have shown that the modified Bayesian–CNN performs slightly better than the Bayesian–CNN in all performance metrics. Especially it can reduce the false negative and false positive predictions by 3% as compared to Bayesian–CNN.

Fifthly, we have demonstrated that using the proposed modified Bayesian–CNN and the Bayesian–CNN, the problem of overfitting can be nearly eliminated without any recourse to regularization. That is the Bayesian approach work with almost equal accuracy for unknown and known data sets.

This work improves the understanding of the uncertainties in Bayesian–CNN-based classification and helps to leverage these uncertainties to further improve its performance. These findings require further investigation on larger data sets, however, they show potential routes to improve upon the state-of-the-art classifier for a broad range of biomedical applications.

## References

- [1] C. P. Wild, B. W. Stewart, and C. Wild, *World cancer report 2014*. World Health Organization Geneva, Switzerland, 2014.
- [2] C. Fitzmaurice *et al.*, “Global, regional, and national cancer incidence, mortality, years of life lost, years lived with disability, and disability-adjusted life-years for 29 cancer groups, 1990 to 2016: a systematic analysis for the global burden of disease study,” *JAMA oncology*, vol. 4, no. 11, pp. 1553–1568, 2018.
- [3] F. Bray, P. McCarron, and D. M. Parkin, “The changing global patterns of female breast cancer incidence and mortality,” *Breast cancer research*, vol. 6, no. 6, p. 229, 2004.
- [4] Z. Han, B. Wei, Y. Zheng, Y. Yin, K. Li, and S. Li, “Breast cancer multi-classification from histopathological images with structured deep learning model,” *Scientific reports*, vol. 7, no. 1, pp. 1–10, 2017.
- [5] M. N. Gurcan, L. E. Boucheron, A. Can, A. Madabhushi, N. M. Rajpoot, and B. Yener, “Histopathological image analysis: A review,” *IEEE reviews in biomedical engineering*, vol. 2, pp. 147–171, 2009.
- [6] F. Ghaznavi, A. Evans, A. Madabhushi, and M. Feldman, “Digital imaging in pathology: whole-slide imaging and beyond,” *Annual Review of Pathology: Mechanisms of Disease*, vol. 8, pp. 331–359, 2013.
- [7] S. Al-Janabi, A. Huisman, and P. J. Van Diest, “Digital pathology: current status and future perspectives,” *Histopathology*, vol. 61, no. 1, pp. 1–9, 2012.
- [8] T. J. Fuchs and J. M. Buhmann, “Computational pathology: Challenges and promises for tissue analysis,” *Computerized Medical Imaging and Graphics*, vol. 35, no. 7-8, pp. 515–530, 2011.
- [9] M. L. Giger, H.-P. Chan, and J. Boone, “Anniversary paper: history and status of cad and quantitative image analysis: the role of medical physics and aapm,” *Medical physics*, vol. 35, no. 12, pp. 5799–5820, 2008.
- [10] K. Doi, “Computer-aided diagnosis in medical imaging: historical review, current status and future potential,” *Computerized medical imaging and graphics*, vol. 31, no. 4-5, pp. 198–211, 2007.
- [11] M. Aswathy and M. Jagannath, “Detection of breast cancer on digital histopathology images: Present status and future possibilities,” *Informatics in Medicine Unlocked*, vol. 8, pp. 74–79, 2017.
- [12] P. Filipczuk, T. Fevens, A. Krzyżak, and R. Monczak, “Computer-aided breast cancer diagnosis based on the analysis of cytological images of fine needle biopsies,” *IEEE transactions on medical imaging*, vol. 32, no. 12, pp. 2169–2178, 2013.
- [13] Y. M. George, H. H. Zayed, M. I. Roushdy, and B. M. Elbagoury, “Remote computer-aided breast cancer detection and diagnosis system based on cytological images,” *IEEE Systems Journal*, vol. 8, no. 3, pp. 949–964, 2013.
- [14] M. Kowal, P. Filipczuk, A. Obuchowicz, J. Korbicz, and R. Monczak, “Computer-aided diagnosis of breast cancer based on fine needle biopsy microscopic images,” *Computers in biology and medicine*, vol. 43, no. 10, pp. 1563–1572, 2013.
- [15] H. Asri, H. Mousannif, H. Al Moatassime, and T. Noel, “Using machine learning algorithms for breast cancer risk prediction and diagnosis,” *Procedia Computer Science*, vol. 83, pp. 1064–1069, 2016.
- [16] Y. Zhang, B. Zhang, F. Coenen, and W. Lu, “Breast cancer diagnosis from biopsy images with highly reliable random subspace classifier ensembles,” *Machine vision and applications*, vol. 24, no. 7, pp. 1405–1420, 2013.
- [17] F. A. Spanhol, L. S. Oliveira, C. Petitjean, and L. Heutte, “A dataset for breast cancer histopathological image classification,” *IEEE Transactions on Biomedical Engineering*, vol. 63, no. 7, pp. 1455–1462, 2015.
- [18] C. Mercan, S. Aksoy, E. Mercan, L. G. Shapiro, D. L. Weaver, and J. G. Elmore, “Multi-instance multi-label learning for multi-class classification of whole slide breast histopathology images,” *IEEE transactions on medical imaging*, vol. 37, no. 1, pp. 316–325, 2017.
- [19] N. Bayramoglu, J. Kannala, and J. Heikkilä, “Deep learning for magnification independent breast cancer histopathology image classification,” in *2016 23rd International conference on pattern recognition (ICPR)*, pp. 2440–2445, IEEE, 2016.
- [20] A. Cruz-Roa *et al.*, “Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks,” in *Medical Imaging 2014: Digital Pathology*, vol. 9041, p. 904103, International Society for Optics and Photonics, 2014.
- [21] F. A. Spanhol, L. S. Oliveira, C. Petitjean, and L. Heutte, “Breast cancer histopathological image classification using convolutional neural networks,” in *2016 international joint conference on neural networks (IJCNN)*, pp. 2560–2567, IEEE, 2016.
- [22] T. Araújo *et al.*, “Classification of breast cancer histology images using convolutional neural networks,” *PloS one*, vol. 12, no. 6, 2017.
- [23] H. Lin, H. Chen, S. Graham, Q. Dou, N. Rajpoot, and P.-A. Heng, “Fast scannet: fast and dense analysis of multi-gigapixel whole-slide images for cancer metastasis detection,” *IEEE transactions on medical imaging*, vol. 38, no. 8, pp. 1948–1958, 2019.
- [24] K. Sirinukunwattana, S. E. A. Raza, Y.-W. Tsang, D. R. Snead, I. A. Cree, and N. M. Rajpoot, “Locality sensitive deep learning for detection and classification of nuclei in routine colon cancer histology images,” *IEEE transactions on medical imaging*, vol. 35, no. 5, pp. 1196–1206, 2016.
- [25] V. Badrinarayanan, A. Kendall, and R. Cipolla, “Segnet: A deep convolutional encoder-decoder architecture for image segmentation,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 12, pp. 2481–2495, 2017.
- [26] Z. Jia, X. Huang, I. Eric, C. Chang, and Y. Xu, “Constrained

- deep weak supervision for histopathology image segmentation,” *IEEE transactions on medical imaging*, vol. 36, no. 11, pp. 2376–2388, 2017.
- [27] Y. Li, J. Wu, and Q. Wu, “Classification of breast cancer histology images using multi-size and discriminative patches based on deep learning,” *IEEE Access*, vol. 7, pp. 21400–21408, 2019.
- [28] H. Yang, J.-Y. Kim, H. Kim, and S. P. Adhikari, “Guided soft attention network for classification of breast cancer histopathology images,” *IEEE transactions on medical imaging*, 2019.
- [29] B. Xu *et al.*, “Attention by selection: A deep selective attention approach to breast cancer classification,” *IEEE Transactions on Medical Imaging*, 2019.
- [30] M. Nawaz, A. A. Sewissy, and T. H. A. Soliman, “Multi-class breast cancer classification using deep learning convolutional neural network,” *Int. J. Adv. Comput. Sci. Appl.*, vol. 9, no. 6, pp. 316–332, 2018.
- [31] N. H. Motlagh, M. Jannesary, H. Aboulkheyr, P. Khosravi, O. Elemento, M. Totonchi, and I. Hajirasouliha, “Breast cancer histopathological image classification: A deep learning approach,” *bioRxiv*, p. 242818, 2018.
- [32] J. Xie, R. Liu, J. Luttrell IV, and C. Zhang, “Deep learning based analysis of histopathological images of breast cancer,” *Frontiers in genetics*, vol. 10, p. 80, 2019.
- [33] G. Aresta *et al.*, “Bach: Grand challenge on breast cancer histology images,” *Medical image analysis*, vol. 56, pp. 122–139, 2019.
- [34] S. S. Chennamsetty, M. Safwan, and V. Alex, “Classification of breast cancer histology image using ensemble of pre-trained neural networks,” in *International conference image analysis and recognition*, pp. 804–811, Springer, 2018.
- [35] S. Kwok, “Multiclass classification of breast cancer in whole-slide images,” in *International conference image analysis and recognition*, pp. 931–940, Springer, 2018.
- [36] N. Brancati, M. Frucci, and D. Riccio, “Multi-classification of breast cancer histology images by using a fine-tuning strategy,” in *International conference image analysis and recognition*, pp. 771–778, Springer, 2018.
- [37] 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.
- [38] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.
- [39] C. Szegedy *et al.*, “Going deeper with convolutions,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1–9, 2015.
- [40] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [41] 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.
- [42] E. Begoli, T. Bhattacharya, and D. Kusnezov, “The need for uncertainty quantification in machine-assisted medical decision making,” *Nature Machine Intelligence*, vol. 1, no. 1, pp. 20–23, 2019.
- [43] Y. Gal, R. Islam, and Z. Ghahramani, “Deep bayesian active learning with image data,” in *International Conference on Machine Learning*, pp. 1183–1192, PMLR, 2017.
- [44] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, “Weight uncertainty in neural networks,” *arXiv preprint arXiv:1505.05424*, 2015.
- [45] K. Shridhar, F. Laumann, and M. Liwicki, “A comprehensive guide to bayesian convolutional neural network with variational inference,” *arXiv preprint arXiv:1901.02731*, 2019.
- [46] K. Shridhar, F. Laumann, and M. Liwicki, “Uncertainty estimations by softplus normalization in bayesian convolutional neural networks with variational inference,” *arXiv preprint arXiv:1806.05978*, 2018.
- [47] A. Kendall and Y. Gal, “What uncertainties do we need in bayesian deep learning for computer vision?,” in *Advances in neural information processing systems*, pp. 5574–5584, 2017.
- [48] Y. Kwon, J.-H. Won, B. J. Kim, and M. C. Paik, “Uncertainty quantification using bayesian neural networks in classification: Application to ischemic stroke lesion segmentation,” 2018.
- [49] L. Rkaczkowski, M. Możejko, J. Zambonelli, and E. Szczurek, “Ara: accurate, reliable and active histopathological image classification framework with bayesian deep learning,” *Scientific reports*, vol. 9, no. 1, pp. 1–12, 2019.
- [50] S. Qian, H. Liu, C. Liu, S. Wu, and H. San Wong, “Adaptive activation functions in convolutional neural networks,” *Neuro-computing*, vol. 272, pp. 204–212, 2018.
- [51] A. D. Jagtap, K. Kawaguchi, and G. E. Karniadakis, “Adaptive activation functions accelerate convergence in deep and physics-informed neural networks,” *Journal of Computational Physics*, vol. 404, p. 109136, 2020.
- [52] M. Dushkoff and R. Ptucha, “Adaptive activation functions for deep networks,” *Electronic Imaging*, vol. 2016, no. 19, pp. 1–5, 2016.
- [53] Paul Mooney, “Breast Histopathology Images.” <https://www.kaggle.com/paultimothymooney/breast-histopathology-images>. 2017.
- [54] L. Van der Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [55] 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.
- [56] 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, PMLR, 2013.
- [57] H.-C. Shin, H. R. Roth, M. Gao, L. Lu, Z. Xu, I. Noguees, J. Yao, D. Mollura, and R. M. Summers, “Deep convolutional neural networks for computer-aided detection: Cnn architectures, dataset characteristics and transfer learning,” *IEEE transactions on medical imaging*, vol. 35, no. 5, pp. 1285–1298, 2016.
- [58] C. M. Bishop, *Pattern recognition and machine learning*. springer, 2006.
- [59] G. B. Orr and K.-R. Müller, *Neural networks: tricks of the trade*. Springer, 2003.
- [60] J. Friedman, T. Hastie, and R. Tibshirani, *The elements of statistical learning*, vol. 1. Springer series in statistics New York, 2001.
- [61] Y. LeCun, “The mnist database of handwritten digits,” <http://yann.lecun.com/exdb/mnist/>, 1998.
- [62] 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.
- [63] J. Cohen, “A coefficient of agreement for nominal scales,” *Educational and psychological measurement*, vol. 20, no. 1, pp. 37–46, 1960.
- [64] J. Diong, A. A. Butler, S. C. Gandevia, and M. E. Héroux, “Poor statistical reporting, inadequate data presentation and spin persist despite editorial advice,” *PLoS one*, vol. 13, no. 8, p. e0202121, 2018.
- [65] S. Rajaraman, S. Sornapudi, P. O. Alderson, L. R. Folio, and S. K. Antani, “Analyzing inter-reader variability affecting deep ensemble learning for covid-19 detection in chest radiographs,”

*PloS one*, vol. 15, no. 11, p. e0242301, 2020.

- [66] T. G. Dietterich, “Approximate statistical tests for comparing supervised classification learning algorithms,” *Neural computation*, vol. 10, no. 7, pp. 1895–1923, 1998.