

Improved Trainable Calibration Method for Neural Networks on Medical Imaging Classification¹

Gongbo Liang¹
tliang130@gmail.com

Yu Zhang¹
y.zhang@uky.edu

Xiaoqin Wang²
xiaoqin.wang@uky.edu

Nathan Jacobs¹
jacobs@cs.uky.edu

¹ Department of Computer Science
University of Kentucky
Lexington, KY, USA

² Department of Radiology
University of Kentucky
Lexington, KY, USA

Abstract

Recent works have shown that deep neural networks can achieve super-human performance in a wide range of image classification tasks in the medical imaging domain. However, these works have primarily focused on classification accuracy, ignoring the important role of uncertainty quantification. Empirically, neural networks are often miscalibrated and overconfident in their predictions. This miscalibration could be problematic in any automatic decision-making system, but we focus on the medical field in which neural network miscalibration has the potential to lead to significant treatment errors. We propose a novel calibration approach that maintains the overall classification accuracy while significantly improving model calibration. The proposed approach is based on expected calibration error, which is a common metric for quantifying miscalibration. Our approach can be easily integrated into any classification task as an auxiliary loss term, thus not requiring an explicit training round for calibration. We show that our approach reduces calibration error significantly across various architectures and datasets.

1 Introduction

Recent advances in deep learning research have dramatically impacted the research field of medical imaging analysis [18, 28, 31]. Many high-performance deep learning models have been developed in the field [2, 21, 37, 40]. Researchers are actively pushing convolutional neural networks (CNNs) to have higher and higher accuracy, while uncertainty quantification is often ignored when evaluating these models [19, 25, 26, 35]. However, uncertainty quantification of neural networks is important, especially in automatic decision-making settings in the medical field. An automated method that achieves high accuracy, but captures uncertainty inaccurately, such as providing inaccurate confidence or probability of a specific prediction, could lead to significant treatment errors [11].

¹This study is sponsored by Grant No. IIS-1553116 from the U.S. National Science Foundation and Grant No. IRG-19-140-31 from the American Cancer Society.

© 2020. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.




Figure 1: Left: The train loss/accuracy and test loss/accuracy of the uncalibrated model. The model is overfitted after the 7th epoch, where the train loss keeps decreasing but the test loss keeps increasing. Right: The train loss/accuracy and test loss/accuracy of our method. The DCA term penalizes the model when the loss reduces but the accuracy is plateaued. Both the train and test losses maintain at the same level after the 7th epoch.

Unfortunately, deep neural networks are poorly calibrated [16, 24], which are overconfident in their predictions [4, 16, 24]. One reason for miscalibration of classification models is that the models can overfit the cross-entropy loss easily without overfitting the 0/1 loss (i.e., accuracy) [4, 38]. We propose to add the difference between predicted confidence and accuracy (DCA) as an auxiliary loss for classification model calibration. The DCA term applies a penalty when the cross-entropy loss reduces but the accuracy is plateaued (Figure 1).

We evaluate the proposed method across four public medical datasets and four widely used CNN architectures. The results show that our approach reduces calibration error significantly by an average of 65.72% compared to uncalibrated methods (from 0.1006 ECE to 0.0345 ECE), while maintaining the overall accuracy across all the experiments—83.08% and 83.58% for the uncalibrated method and our method, respectively.

2 Background

The problem we are addressing is the miscalibration issue of deep neural networks for classification tasks. The confidence associated with a prediction (i.e., probability of being one specific class) should reflect the true correctness likelihood of a model [4]. However, deep neural networks tend to be overconfident in their predictions [16, 24].

2.1 Problem Definition

Mathematically, the problem can be defined in the following way. The input $X \in x$ and label $Y \in y = \{1, \dots, k\}$ are random variables that follow a joint distribution $\pi(X, Y) = \pi(Y|X)\pi(X)$. Let h be a deep neural network with $h(X) = (\hat{Y}, \hat{P})$, where \hat{Y} is the predicted class label and \hat{P} is the associated confidence. We would like the confidence estimate \hat{P} to be calibrated, which intuitively means that \hat{P} represents a true probability. For instance, given 100 predictions with the average confidence of 0.95, we expect that 95 predictions should be correct. In reality, the average confidence of a deep neural network is often higher than its accuracy [4, 16, 24]. The perfect calibration can be defined as:

$$\mathbb{P}(\hat{Y} = Y | \hat{P} = p) = p, \forall p \in [0, 1]. \quad (1)$$

Difference in expectation between confidence and accuracy (i.e., the calibration error) can be defined as:

$$\mathbb{E}_{\hat{p}} [|(\hat{Y} = Y | \hat{P} = p) - p|]. \quad (2)$$

2.2 Measurements

Expected Calibration Error (ECE) is a commonly used criterion for measuring neural network calibration error. ECE [23] approximates Equation (2) by partitioning predictions into M bins and taking a weighted average of the accuracy/confidence difference for each bin. All the samples need to be grouped into M interval bins according to the predicted probability. Let B_m be the set of indices of samples whose predicted confidence falls into the interval $I_m = (\frac{m-1}{M}, \frac{m}{M}]$, $m \in M$. The accuracy of B_m is

$$\text{acc}(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} 1(\hat{y}_i = y_i), \quad (3)$$

where \hat{y}_i and y_i are the predicted and ground-truth label for sample i . The average predicted confidence of bin B_m can be defined as

$$\text{conf}(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} \hat{p}_i, \quad (4)$$

where \hat{p}_i is the confidence of sample i . ECE can be defined with $\text{acc}(B_m)$ and $\text{conf}(B_m)$

$$\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{n} |\text{acc}(B_m) - \text{conf}(B_m)|, \quad (5)$$

where n is the number of samples.

Maximum Calibration Error (MCE) [23] is another common criterion for measuring neural network calibration error that partitions predictions into M equally-spaced bins and estimates the worst-case scenario. MCE can be computed as:

$$\text{MCE} = \max_{m \in \{1, \dots, m\}} |\text{acc}(B_m) - \text{conf}(B_m)|. \quad (6)$$

3 Existing Calibration Methods

In this section, we introduce some existing calibration methods, including temperature scaling [4, 7], entropy regularization [24], MMCE regularization [16], label smoothing [22, 32], and Mixup training [33, 39]. Temperature scaling is a widely used calibration method, which treats model calibration as a post-processing task. All the other methods fix neural network calibration during the classification training stage.

3.1 Temperature Scaling

Temperature scaling [4, 7] is a widely-used approach for deep learning model calibration. It fixes the miscalibration issue by dividing the logits by a temperature parameter of T ($T > 0$). The method involves two steps, in general. The first step is to train a classification model. Once the model is trained, the temperature parameter is added to the model and needs to

be trained on the validation set while all the other parameters are frozen [4]. After that, the temperature parameter will be used for calibration at the testing time. The calibrated confidence, \hat{q}_i , using temperature scaling is

$$\hat{q}_i = \max_k \theta_{SM}\left(\frac{z_i}{T}\right)^{(k)}, \quad (7)$$

where k is the class label ($k = 1, \dots, K$), $\theta_{SM}(z_i)$ is the predicted confidence. As $T \rightarrow \infty$, the confidence \hat{q}_i approaches the minimum, which indicates maximum uncertainty.

Temperature scaling is easy to use and performs well. The optimization process of the temperature parameter is not expensive and only needs to be done once. However, as a post-processing approach, temperature scaling does not help with feature learning. In addition, a neural network model should be able to calibrate itself without any post-processing [34].

3.2 Trainable Calibration Methods

Trainable calibration methods are proposed to integrate model calibration into classification training. No explicit training round for calibration is needed in such a fashion. One of the earliest trainable approaches is the entropy regularization [24]. The method proposes to use entropy as a regularization term in loss functions for model calibration. The final classification loss can be written as:

$$\text{Loss} = \text{CrossEntropy} + \beta \text{Entropy}, \quad (8)$$

where β is a weight scalar. One disadvantage of entropy regularization is that the method is very sensitive to the value of β [16]. Kumar et al. propose to use MMCE replacing entropy for model calibration [16]. MMCE is computed in a reproducing kernel Hilbert space (RKHS) [3]. The completely loss function can be written as:

$$\text{Loss} = \text{CrossEntropy} + \beta (\text{MMCE}_m^2(D))^{\frac{1}{2}}, \quad (9)$$

where D denotes a dataset. The performance of MMCE may be limited by imbalance predictions of a neural network. For instance, the number of correct predictions is usually larger than the number of incorrect predictions. Thus, the MMCE term needs to be re-weighted as:

$$\begin{aligned} \text{MMCE}_w^2 = & \sum_{c_i=c_j=0} \frac{p_i, p_j, k(p_i, p_j)}{(m-n)^2} + \\ & \sum_{c_i=c_j=1} \frac{(1-p_i)(1-p_j)k(p_i, p_j)}{n^2} - \\ & 2 \sum_{c_i=1, c_j=0} \frac{(1-p_i)p_j k(p_i, p_j)}{(m-n)n}, \end{aligned} \quad (10)$$

where c is the predicted label, m is the number of correct predictions, n is the batch size, and k is a universal kernel [29].

Label smoothing [32] was proposed to improve classification performance of the Inception architecture. Müller et al. demonstrate that label smoothing improves classification

performance by calibrating models implicitly [22]. Instead of targeting a hard probability, 1.0, for the correct class, label smoothing tries to predict a softer version of it:

$$y_k^{LS} = y_k(1 - \alpha) + \frac{\alpha}{K}, \quad (11)$$

where y_k is original targeting probability ($y_k = 1.0$ for the correct class and $y_k = 0.0$ for the rest), K is the number of class labels, α is a hyperparameter that determines the amount of smoothing. Mixup [39] is another method that aims to predict a softer target by randomly mixing training samples. During the training, two samples from different classes are randomly mixed together. Instead of predicting one target label, the network needs to predict the two corresponding labels' probability. The target probabilities equal to the portion of the pixels from each image. Thulasidasan et al. [33] demonstrate that Mixup is also useful for neural network calibration.

4 Proposed Method

We propose to add the difference between confidence and accuracy (DCA) as an auxiliary loss term to the cross-entropy loss for classification tasks. DCA is based on expected calibration error by minimizing the difference between the predicted confidence and accuracy directly. Therefore, the proposed method can calibrate neural networks effectively. The proposed method is easy to implement and suitable for any classification tasks. In general, classification loss can be written as follows:

$$\text{Loss} = -\frac{1}{N} \sum_{i=1}^N y_i \cdot \log(p(y_i)) + \beta DCA, \quad (12)$$

where y_i is the true label and $p(y_i)$ is the predicted confidence (i.e., probability) of the true label. The DCA term can be computed for each mini-batch using the following equation:

$$DCA = \left| \frac{1}{N} \sum_{i=1}^N c_i - \frac{1}{N} \sum_{i=1}^N p(\hat{y}_i) \right|, \quad (13)$$

where \hat{y}_i is the predicted label; $c_i = 1$, if $\hat{y}_i = y_i$; otherwise, $c_i = 0$. The final loss function can be written as:

$$\text{Loss} = -\frac{1}{N} \sum_{i=1}^N y_i \cdot \log(p(y_i)) + \beta \left| \frac{1}{N} \sum_{i=1}^N c_i - \frac{1}{N} \sum_{i=1}^N p(\hat{y}_i) \right|. \quad (14)$$

The DCA auxiliary loss fixes the miscalibration issue by penalizing deep learning models when the cross-entropy loss can be reduced, but the accuracy does not change (i.e., when the model is overfitting). The term forces the average predicted confidence to match the accuracy over all training examples without strict constraint on each example, which pushes the network closer to the ideal situation, in which the accuracy reflects the true correctness likelihood of a model. The averaging mechanism of DCA also smooths the predictions that have extremely high or low confidence.

DCA is differentiable in the predicted confidence term but not strictly in the prediction accuracy term due to the argmax step for computing the predicted label. During the training phase, gradients can be backpropagated through the confidence terms but not through the accuracy.

5 Experiments

We compare the proposed method with temperature scaling and uncalibrated models (trained with cross-entropy loss without applying any calibration methods) on four medical imaging datasets across four popular CNN networks. The trainable methods are not compared in this work because the literature shows that they have a worse or similar calibration performance with temperature scaling [16, 22, 33]. Thus, it may not be necessary to be compared in this paper explicitly.

The evaluation results show that the proposed method significantly improves model calibration while maintaining the overall classification accuracy. The proposed method reduces calibration error by an average of 65.72% compared to uncalibrated methods (from 0.1006 ECE to 0.0345 ECE) and performs about 20% better than temperature scaling on average.

5.1 Experiment Setup

Four medical imaging datasets (RSNA [27], DDSM [6], Mendeley V2 [13], and Kather 5000 [12]) were used in this study for both binary and multi-class classification tasks. See Table S1 and Section S1 for more details. Four CNN models that trained with transfer learning mechanism were evaluated. More specifically, we use ImageNet [1] pre-trained AlexNet [15], ResNet-50 [5], DenseNet-121 [8], and SqueezeNet 1-1 [9] as fixed feature extractors, a 1×1 convolutional layer and two fully connected layers were added on top to each feature extractor. See Section S2 for more details.

5.2 Calibration Results

Table 1 shows the expected calibration error (ECE) and the accuracy of the uncalibrated models (Unca.), temperature scaling (Temp.), and the proposed method (DCA). Each model was trained for two times. The average value is shown in the table.

The table shows that our method is consistently better than the uncalibrated method on model calibration, which reduces the ECE by 65.71% on average (from 0.1006 to 0.0345). Temperature scaling has the second smallest average ECE (0.0427). However, it is still 23.77% worse than the proposed method. On average, the uncalibrated method and temperature scaling have an 83.08% accuracy, while the proposed method has an 83.51% accuracy. The proposed method increases the accuracy of 11 out of 16 tests.

It is worth noting that temperature scaling increases calibration error of 3 out of 4 models on the Kather 5000 dataset, while the proposed method is still able to reduce the calibration error of most cases on the same dataset. The Kather 5000 dataset is a relatively simple and large dataset for its task. The dataset is considered as the MNIST of histology images. It is speculated to have a sufficient amount of training data to train a model end-to-end, with a smaller overfitting effect (i.e., miscalibration). In such a case, since the temperature parameter (T) of temperature scaling is learned on only the validation set, it may actually hurt the calibration. However, the proposed method jointly optimizes the accuracy and modal calibration simultaneously, and it can still reduce the calibration error.

Table 2 shows the maximum calibration error (MCE) of the compared models trained using the RSNA, DDSM, and Mendeley datasets. Though there is no clear winner on MCE, the proposed model decreases the average MCE by about 3%, while temperature scaling increases the MCE slightly. Table 3 shows the MCE results on the Kather 5000 dataset. Both temperature scaling and the proposed method increase the MCE quite well. According

Table 1: Expected Calibration Error (ECE) for Each Model

Dataset	Model	ECE (smaller is better)			Accuracy ¹ (larger is better)	
		Unca.	Temp.	DCA	Unca.	DCA
		RSNA	AlexNet	0.0113	0.0239	0.0120
ResNet	0.0276		0.0231	0.0122	0.8569	0.8762
DenseNet	0.0102		0.0814	0.0077	0.8502	0.8543
SqueezeNet	0.0253		0.0317	0.0097	0.8671	0.8841
DDSM	AlexNet	0.2164	0.0658	0.0591	0.6766	0.6291
	ResNet	0.1844	0.0307	0.0798	0.7195	0.6987
	DenseNet	0.1798	0.0337	0.0754	0.7076	0.7106
	SqueezeNet	0.2173	0.0458	0.0805	0.6853	0.6771
Mendeley	AlexNet	0.1693	0.0396	0.0273	0.8585	0.8785
	ResNet	0.1475	0.0475	0.0291	0.8520	0.8767
	DenseNet	0.1136	0.0746	0.0285	0.8331	0.8796
	SqueezeNet	0.1871	0.0468	0.0252	0.8742	0.8750
Kather	AlexNet	0.0279	0.0344	0.0243	0.9062	0.9052
	ResNet	0.0248	0.0318	0.0304	0.9355	0.9229
	DenseNet	0.0302	0.0286	0.0237	0.9385	0.9410
	SqueezeNet	0.0372	0.0439	0.0269	0.8932	0.9038
Average		0.1006	0.0427	0.0345	0.8308	0.8351

¹The temperature scaling method has the same accuracy as the uncalibrated models.

Table 2: Maximum Calibration Error (MCE) for Binary Classification Tasks

Dataset	Model	MCE (smaller is better)		
		Unca.	Temp.	DCA
		RSNA	AlexNet	0.0291
ResNet	0.0484		0.0325	0.0399
DenseNet	0.0335		0.2233	0.0142
SqueezeNet	0.0430		0.0663	0.0270
DDSM	AlexNet	0.2527	0.1545	0.1800
	ResNet	0.2897	0.1171	0.1078
	DenseNet	0.2403	0.0941	0.0959
	SqueezeNet	0.332	0.1631	0.1586
Mendeley	AlexNet	0.2454	0.4305	0.1225
	ResNet	0.2321	0.1769	0.297
	DenseNet	0.2653	0.2477	0.4898
	SqueezeNet	0.2521	0.2451	0.2507
Average		0.2812	0.2817	0.2737

to Guo et al., MCE is very sensitive to the number of bins since it measures the worst cases across all bins, making it an improper metric for small test sets [4]. On average, temperature scaling has an MCE of 0.3401 across all the evaluated models, the proposed method has an MCE of 0.3398, and the uncalibrated method has an MCE of 0.2766.

Table 3: Maximum Calibration Error for Multi-Class Classification Tasks

Dataset	Model	MCE (smaller is better)		
		Unca.	Temp.	DCA
Kather 5000	AlexNet	0.2570	0.2974	0.7371
	ResNet	0.3072	0.6379	0.2513
	DenseNet	0.2577	0.8015	0.4268
	SqueezeNet	0.2566	0.3237	0.7371
Average		0.2696	0.5151	0.5381




Figure 2: The t-SNE plots of the representations learned using temperature scaling and the proposed method on the Mendeleev V2 (left) and Kather 5000 (right) datasets. The samples of temperature scaling are spreading in the feature space (left in each subplot). The samples of the proposed method are densely packed for the same class (right in each subplot).

5.3 Model Representation Learning

As a post-processing method, temperature scaling fixes model miscalibration at the output-level, which does not change the learned representation. However, the proposed method integrates calibration into the network training phase, which may help models learn a better representation. In this section, we firstly use t-SNE [20] plot to visualize features extracted by temperature scaling and DCA on two datasets. Then, we compare the recovered probability distribution of the two methods on a toy dataset.

Figure 2 shows the t-SNE plots of features extracted by temperature scaling and the proposed method on the Mendeleev V2 and Kather 5000 datasets. Each dot represents one data sample in a 2D feature space. To generate these plots, We first extract the high-dimensional feature maps using an AlexNet model trained with either temperature scaling or DCA. We then feed the feature maps to t-SNE that projects the high-dimensional feature maps to 2D feature space. Ideally, the samples from the same class should be close to each other; the samples from different classes should be far from each other.

The plots reveal that the samples of temperature scaling (left in each subplot) are spreading in the feature space regardless of class labels, while the samples of the same class are densely packed in our method (right in each subplot). Especially for the Kather 5000 dataset, DCA (the rightmost figure) successfully separated the samples of TE (blue) and GB (gray) classes from the rest. But, the same classes of temperature scaling are mixed with others.

Figure 3 shows the probability distribution recovered by the uncalibrated method (left), temperature scaling (middle), and the proposed method (right) on a toy dataset. For this experiment, we train three simple networks using the uncalibrated method, temperature scaling, and the proposed method, separately. The networks share the same two-layer architecture that takes a one-dimensional input for a binary classification task. We randomly sample a

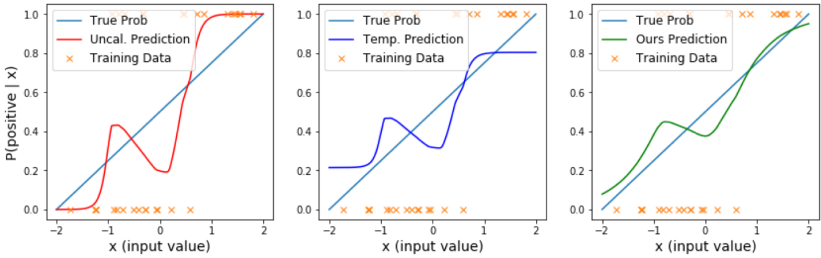


Figure 3: The figure shows the probability distribution that was recovered by the uncalibrated method (left), temperature scaling (middle), and the proposed method (right). The recovered distribution of the uncalibrated model (left) is far from the ground-truth with many overconfident predictions. Temperature scaling (middle) reduces the predicted confidence of the uncalibrated model, but the recovered distribution is still far from the ground-truth. Our method (right) can better recover the true probability.




Figure 4: Train/Test accuracy and loss of Mendeley V2 dataset using AlexNet architecture with four different β values.

dataset between -2 and 2 , and randomly label each sample with either 0 or 1 . The curved line in each figure shows the recovered probability distribution, while the light blue line in each figure shows the ground-truth distribution.

The figures reveal that the uncalibrated model (left) has many high-probability predictions. For instance, when $-2 < x < -1$, the model made many negative predictions with high-probability; when $1 < x < 2$, the model made many positive predictions with high-probability. The majority of these predicted probabilities are far from the true probabilities, which indicates the model is overconfident in its predictions and does not capture the true probability distribution well. The temperature scaling method (middle) can relax those extreme predictions by pushing the predicted probabilities close to 0.5 . However, the recovered probability distribution still quite far from the diagonal line. The proposed method (right) can recover the trend of the ground-truth distribution, and most of the predictions are close to the diagonal line. This experiment shows the models trained with DCA may have a strong ability to recover the true probability distribution more accurately.

5.4 Hyperparameter Effects

One drawback of the proposed method is that the weight scalar β needs to be selected for each model. In this section, we show the testing result of the proposed method with different weights ($\beta = [1, 5, 10, 15, 20, 25]$). Figure 4 shows the train/test accuracy and loss on the




Figure 5: Expected calibration error (ECE) of each dataset with different β values.

Mendeley V2 dataset using AlexNet architecture with four different β values. From the results, we can see that a smaller value such as 1 or 5 usually will not be a good choice since they put a smaller penalty to the model when the cross-entropy loss is overfitted. Among our experiments, most of the best results appeared using a weight between 10 and 25. Figure 5 shows the ECE results of all of the evaluated models with different β values. The figure reveals that the ECE result is not very sensitive to the β value when $\beta \geq 10$, except for the Kather 5000 dataset. In our experiences, we use β values between 10 and 15 for most of the tasks.

6 Conclusion

We proposed a novel approach to neural network calibration that maintains classification accuracy while significantly reducing model calibration error. We evaluated our approach across various architectures and datasets. Our approach reduces calibration error significantly and comes closer to recovering the true probability than other approaches. The proposed method can be easily integrated into any classification tasks as an auxiliary loss term, thus not requiring an explicit training round for calibration. We believe this simple, fast, and straightforward method can serve as a strong baseline for future researchers.

References

- [1] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proc. CVPR*, pages 248–255, 2009.
- [2] Andre Esteva et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639):115, 2017.
- [3] Arthur Gretton. Introduction to rkhs, and some simple kernel algorithms. *Adv. Top. Mach. Learn. Lecture Conducted from University College London*, 2013.
- [4] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *Proc. ICML*, pages 1321–1330, 2017.
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. CVPR*, pages 770–778, 2016.
- [6] Michael Heath, Kevin Bowyer, Daniel Kopans, Richard Moore, and W Philip Kegelmeyer. The digital database for screening mammography. In *Proceedings of the 5th international workshop on digital mammography*, pages 212–218, 2000.

- [7] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv:1503.02531*, 2015.
- [8] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proc. CVPR*, pages 4700–4708, 2017.
- [9] Forrest N Iandola et al. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and < 0.5 mb model size. *arXiv:1602.07360*, 2016.
- [10] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv:1502.03167*, 2015.
- [11] Xiaoqian Jiang, Melanie Osl, Jihoon Kim, and Lucila Ohno-Machado. Calibrating predictive model estimates to support personalized medicine. *Journal of the American Medical Informatics Association*, 19(2):263–274, 2011.
- [12] Jakob Nikolas Kather et al. Multi-class texture analysis in colorectal cancer histology. *Scientific reports*, 6:27988, 2016.
- [13] DK Kermany and M Goldbaum. Labeled optical coherence tomography (oct) and chest x-ray images for classification. *Mendeley Data*, 2, 2018.
- [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014.
- [15] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *Proc. NIPS*, pages 1097–1105, 2012.
- [16] Aviral Kumar, Sunita Sarawagi, and Ujjwal Jain. Trainable calibration measures for neural networks from kernel mean embeddings. In *Proc. ICML*, pages 2810–2819, 2018.
- [17] R Sawyer Lee, Francisco Gimenez, Assaf Hoogi, and Daniel Rubin. Curated breast imaging subset of ddsM. *The Cancer Imaging Archive*, 2016.
- [18] Gongbo Liang, Sajjad Fouladvand, Jie Zhang, Michael A Brooks, Nathan Jacobs, and Jin Chen. Ganai: standardizing ct images using generative adversarial network with alternative improvement. In *Proc. IEEE ICHI*, pages 1–11, 2019.
- [19] Gongbo Liang, Xiaoqin Wang, Yu Zhang, Xin Xing, Hunter Blanton, Tawfiq Salem, and Nathan Jacobs. Joint 2d-3d breast cancer classification. In *Proc. IEEE BIBM*, pages 692–696, 2019.
- [20] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [21] R. P. Mihail, G. Liang, and N. Jacobs. Automatic hand skeletal shape estimation from radiographs. *IEEE Tran. on NanoBioscience*, 18(3):296–305, 2019.
- [22] Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? In *Proc. NeurIPS*, pages 4694–4703, 2019.
- [23] Mahdi Pakdaman Naeni, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *Proc. AAAI*, 2015.

- [24] Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions. *arXiv:1701.06548*, 2017.
- [25] Dezső Ribli, Anna Horváth, Zsuzsa Unger, Péter Pollner, and István Csabai. Detecting and classifying lesions in mammograms with deep learning. *Scientific reports*, 8(1): 4165, 2018.
- [26] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241. Springer, 2015.
- [27] RSNA. RSNA intracranial hemorrhage detection, 2019. URL www.kaggle.com/c/rsna-intracranial-hemorrhage-detection/overview.
- [28] Berkman Sahiner et al. Deep learning in medical imaging and radiation therapy. *Medical physics*, 46(1):e1–e36, 2019.
- [29] Le Song. Learning via hilbert space embedding of distributions. 2008.
- [30] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [31] Kenji Suzuki. Overview of deep learning in medical imaging. *Radiological physics and technology*, 10(3):257–273, 2017.
- [32] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proc. CVPR*, June 2016.
- [33] Sunil Thulasidasan, Gopinath Chennupati, Jeff A Bilmes, Tanmoy Bhattacharya, and Sarah Michalak. On mixup training: Improved calibration and predictive uncertainty for deep neural networks. In *Proc. NeurIPS*, pages 13888–13899, 2019.
- [34] David Widmann, Fredrik Lindsten, and Dave Zachariah. Calibration tests in multi-class classification: A unifying framework. In *Proc. NeurIPS*, pages 12236–12246, 2019.
- [35] Xin Xing et al. Dynamic image for 3d mri image alzheimer’s disease classification. In *Proc. ECCV Workshops*, 2020.
- [36] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations in convolutional network. *arXiv:1505.00853*, 2015.
- [37] Qingsong Yang et al. Low-dose ct image denoising using a generative adversarial network with wasserstein distance and perceptual loss. *IEEE tran. on medical imaging*, 37(6):1348–1357, 2018.
- [38] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv:1611.03530*, 2016.
- [39] Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *Proc. ICLR*, 2018.
- [40] Yu Zhang, Xiaoqin Wang, Hunter Blanton, Gongbo Liang, Xin Xing, and Nathan Jacobs. 2d convolutional neural networks for 3d digital breast tomosynthesis classification. In *Proc. IEEE BIBM*, pages 1013–1017, 2019.