

Studying the Impact of Augmentations on Medical Confidence Calibration

Adrit Rao^{1,2} Joon-Young Lee³ Oliver Aalami¹

adritrao@stanford.edu jolee@adobe.com aalami@stanford.edu

¹Stanford University ²Palo Alto High School ³Adobe Research

Abstract

The clinical explainability of convolutional neural networks (CNN) heavily relies on the joint interpretation of a model's predicted diagnostic label and associated confidence. A highly certain or uncertain model can significantly impact clinical decision-making. Thus, ensuring that confidence estimates reflect the true correctness likelihood for a prediction is essential. CNNs are often poorly calibrated and prone to overconfidence leading to improper measures of uncertainty. This creates the need for confidence calibration. However, accuracy and performance-based evaluations of CNNs are commonly used as the sole benchmark for medical tasks. Taking into consideration the risks associated with miscalibration is of high importance. In recent years, modern augmentation techniques, which cut, mix, and combine images, have been introduced. Such augmentations have benefited CNNs through regularization, robustness to adversarial samples, and calibration. Standard augmentations based on image scaling, rotating, and zooming, are widely leveraged in the medical domain to combat the scarcity of data. In this paper, we evaluate the effects of three modern augmentation techniques, CutMix, MixUp, and CutOut on the calibration and performance of CNNs for medical tasks. CutMix improved calibration the most while CutOut often lowered the level of calibration.

1. Introduction

The applications of computer vision to medical image analysis have been widely studied [52, 38, 36, 19]. In medical image analysis, models are trained to assist clinicians in the triaging of image diagnosis through the interpretation of medical images (e.g. CT, MRI, CXR, skin, etc.). The goal of these models is to increase diagnostic accuracy and efficiency by reducing manual time-consuming tasks [9]. In recent years, such models have become increasingly accurate [47, 43, 28, 48, 58] leading to the wider adoption of computer vision-based tools in clinical settings. Currently, the accuracy of medical image analysis models are evaluated primarily based on performance-based statistical met-




Figure 1. **Modern image augmentations applied to various medical modalities to create unseen samples.** In this study, we evaluate the effects of modern augmentations on the confidence calibration and performance of CNNs for medical image analysis. Augmentations studied are: CutMix [60], MixUp [62], and CutOut [17]. CNNs are trained and benchmarked on augmented datasets of the Skin, CT, CXR, and MRI medical modalities.

rics [3]. Upon obtaining high accuracy, these models are typically deployed in the clinical setting for validation or comparison against trained clinicians [46, 7, 4, 44].

In addition to maintaining a high accuracy, it is important to ensure the reliability of these models to support safe clinical decision-making. In most medical image analysis tasks, a predicted diagnostic label and associated confidence probability are presented jointly to the clinician. A clinician must interpret this information taking into consideration not

only the provided diagnosis but the certainty or uncertainty of the model to gain trust [13]. While *accuracy* solely encompasses the ability to correctly predict a certain class, a model must provide a *reliable* confidence estimation to the clinician. When using these models for diagnostic aid, confidence estimates can provide major insight to clinicians and significantly influence diagnosis. For example, a model which outputs a confidence of 98% should provide the clinician with a higher certainty for that diagnostic label. On the other hand, a low confidence of 60% should help the clinician consider the uncertainty and potentially revise the diagnosis. Ensuring that confidence estimates are reflective of the *true correctness likelihood* of diagnostic labels is essential, creating the need for **confidence calibration** [22].

It has been shown that state-of-the-art convolutional neural networks (CNN) are often poorly calibrated and tend to be overconfident or *overly certain* in predictions [22]. This poses a significant problem and potentially a high risk in the medical domain. Calibrated confidence estimates are not captured in performance-based metrics commonly used to evaluate medical image analysis models [3]. A low level of calibration can easily go unnoticed leading to improper measures of uncertainty and potentially false interpretations when deployed in the clinical setting. In the general computer vision community, various techniques have been proposed to improve the calibration, reliability, and transparency of neural networks [16, 39, 20, 35]. Modifying certain features (e.g. network width, height, depth) of a CNN can also significantly effect confidence calibration [22]. However, such calibration techniques are not commonly leveraged for medical image analysis as many applications perform transfer learning of pretrained CNNs [33, 49, 45]. These applications may yield high performance but could potentially lack high levels of calibration.

In recent years, various *modern image augmentation* techniques have been introduced. Three notable techniques which are studied in this paper are MixUp [62], CutMix [60], and CutOut [17]. As implied in the names, these augmentations perform unique operations such as cutouts, blends, and mixes on images to generate unseen samples. Each of these techniques have not only benefited CNNs in terms of performance but also have significant benefits for confidence calibration. Various studies have shown the advantages of these augmentations for calibration of CNNs using standard calibration-based metrics on general computer vision benchmarks [56, 63, 8, 12, 14]. Additionally, similar to standard image augmentations (e.g. crop, rotate, flip) these techniques improve the regularization and robustness of CNNs to out-of-distribution (OOD) samples.

Standard image augmentations [34, 54, 27] such as random scaling, rotation, and zooming of image samples are widely leveraged for medical image analysis [11]. The reason augmentations are used in the medical domain is to:

increase the *size* and the *diversity* of datasets to improve robustness and reduce overfitting (i.e improve regularization). From a clinical perspective, augmentations are also beneficial in combatting the scarcity of large clinically-acquired and annotation intensive datasets.

Due to the various benefits of modern augmentations, most importantly on the calibration of CNNs, the adoption of these techniques can be very beneficial for medical image analysis to improve the reliability and uncertainty measurements of models. Additionally, as augmentations are already used for medical image analysis and present various other advantages, modern augmentation techniques present a low barrier of entry. In comparison to other calibration techniques, modern augmentations do not effect the structure of CNN architectures and can effect calibration solely based on modifications made to datasets. Additionally, augmentations have a very low computational cost. By using modern image augmentations, major modifications do not have to be made to medical image analysis pipelines. In this study, we evaluate the performance of modern image augmentations for medical confidence calibration using various open-source medical image datasets.

Our main **contributions** are summarized below:

1. We evaluate the effects of modern augmentations on the *performance* of CNNs for medical image analysis
2. We understand the effects of modern augmentations on *medical confidence calibration*
3. We conduct experiments across various *medical modalities* to more deeply understand the effects of modern augmentations across an array of diseases and image types

2. Related Work

As follows is a review of prior work regarding standard calibration methods in the *general* computer vision domain, the use of *modern image augmentations* to improve CNN calibration, and the *current applications* of confidence calibration methods to medical image analysis tasks.

Confidence Calibration: In a study by Gou *et al.* [22], various observations are made on the calibration of CNNs and factors which influence this. Based on in-depth empirical experimentation, the following objective observations have been made relating to CNNs: 1. Increasing the network depth and width of CNNs typically increases accuracy [61] however this has negative effects on calibration. 2. Batch normalization, used for neural network optimization, often leads to miscalibration. 3. Weight decay, a regularization mechanism [57] commonly replaced for batch normalization, has positive effects on calibration. Popular

methodologies which provide benefits in confidence calibration and quantifying predictive uncertainty include temperature scaling [22], bayesian neural networks [16, 39], dropout as bayesian approximation [20, 53], and ensembles of networks [35]. Calibration methods are evaluated based on *calibration metrics* [41, 22] and *reliability plotting* [15, 42, 22]. These two techniques provide both a quantitative and qualitative assessment of confidence calibration.

Modern Augmentations for Calibration: The effects of modern augmentations on the regularization of CNNs is evident from the original studies of MixUp [62], CutMix [60], and CutOut [17]. MixUp [62], an augmentation built around convex combinations of image and label pairs, was the first modern augmentation to be thoroughly studied for confidence calibration. MixUp, when benchmarked using calibration metrics, presented CNNs with significant calibration benefits according to various studies [56, 63, 8]. These studies showed the benefits of MixUp, which was first proposed for regularization, to confidence calibration of CNNs. Subsequent studies covered the calibration of CNNs for the CutMix and CutOut modern augmentations [12, 14]. These studies also concluded that these modern augmentation-based regularization techniques present CNNs with significant calibration benefits.

Medical Image Analysis Calibration: We have not identified prior studies which validate the efficacy of the MixUp, CutMix, and CutOut modern augmentations for the *confidence calibration* of CNNs for medical image analysis tasks. A study by Galdran *et al.* [21] performed experimental validation of MixUp for medical image classification. This study solely used performance-based metrics for evaluation and not calibration-based metrics. MixUp has additionally been used for medical image segmentation however, it has not been benchmarked using calibration metrics [18]. Confidence calibration using methods other than augmentation have been studied in the medical domain [37]. In our study, we focus on understanding the effects of modern augmentations on the calibration of CNNs for medical image analysis to improve to reliability of models. Apart from augmentation-based calibration, other methods have been studied for medical image analysis calibration namely in medical image segmentation [40, 55, 50, 30].

3. Methods

In this study, we perform experiments based on training CNNs across the MixUp, CutMix, and CutOut modern augmentations for various medical image modalities. With this, we evaluate the calibration of each CNN variant against the baseline using conventional calibration metrics and reliability plotting. The goal is to understand the effects of modern

image augmentations on the confidence calibration of these CNNs. We additionally evaluate the accuracy of models using standard performance-based metrics. The formulation of the modern augmentations are briefly reviewed and described in 3.1. The metrics used to evaluate the calibration of the models are documented in Sec 3.2. The various medical image modalities used are reviewed in Sec 3.3. Sec 3.4 and 3.5 review model architectures and our implementation.

3.1. Modern Data Augmentations

MixUp Zhang *et al.* [62] proposed MixUp as a modern augmentation technique for training neural networks on a *blend* between a pair of images and labels based on convex combinations. MixUp has proven various benefits in terms of increasing robustness of neural networks when learning from corrupt labels and adversarial examples. MixUp is based on the Vicinal Risk Minimization (VRM) [10] principle, where the *vicinity* of the training data distribution can be used to draw *virtual* samples and shows improvements over Empirical Risk Minimization (ERM). The original formulation of MixUp from the original paper [62] is:

$$\begin{aligned}\tilde{x} &= \lambda x_i + (1 - \lambda)x_j \\ \tilde{y} &= \lambda y_i + (1 - \lambda)y_j,\end{aligned}\tag{1}$$

where x_i, y_i are raw randomly sampled input vectors and x_j, y_j are the corresponding one-hot label encodings. λ are values in the range [0, 1] which are randomly sampled from the Beta distribution for each augmented example. Samples of the MixUp augmentation technique applied to various medical images are shown in Figure 1.

CutMix Yun *et al.* [60] introduced CutMix, an augmentation built upon the original formulation of MixUp and idea of combining samples. CutMix removes a patch from an image and swaps it for a region of another image generating a locally natural unseen sample. Similar to MixUp, CutMix not only combines two samples but also their corresponding labels. The formulation for CutMix is as follows:

$$\begin{aligned}\tilde{x} &= Mx_i + (1 - M)x_j \\ \tilde{y} &= \mu y_i + (1 - \mu)y_j,\end{aligned}\tag{2}$$

where M indicates the binary mask used to perform the cutout and fill-in operation from two randomly drawn images. μ are values (in [0,1]) randomly drawn from the Beta distribution. Samples of the CutMix technique applied to various medical images are shown in Figure 1.

CutOut This technique was proposed by DeVries *et al.* [17] and is a simple augmentation technique for improving the regularization of CNNs. CutOut was formulated based on the idea of extending dropout [26] to a spatial prior in

the input space. CutOut performs occlusions of an input image similar to the idea proposed in [5]. Rather than partially occluding portions of an image [5], CutOut performs fixed-size zero-masking to fully obstruct a random location of an image. CutOut differentiates from dropout as it is an augmentation technique and visual features are dropped at the input stage of the CNN whereas in dropout, this occurs in intermediate layers. The goal of CutOut is to not only improve regularization of CNNs but improve robustness to occluded samples in real-world applications. Samples of CutOut applied to medical images are shown in Figure 1.

3.2. Calibration and Performance Metrics

As follows are descriptions of the two techniques used to evaluate the effects of the modern augmentations on the calibration of CNNs. The first technique is a quantitative metric based on error and the second technique allows for visualizing calibration through reliability plotting.

Expected Calibration Error (ECE) [41] is a very widely leveraged metric for quantifying the calibration of neural networks. This approach provides a scalar summary statistic of calibration by grouping a models predictions into equally-spaced bins (B). The weighted average of the difference between accuracy and confidence across the bins is outputted. The formulation of ECE from [22] is as follows:

$$ECE = \sum_{b=1}^B \frac{n_b}{N} |\text{acc}(b) - \text{conf}(b)|, \quad (3)$$

where n represents the number of samples. Gaps in calibration or miscalibration is represented by the difference between acc and conf . In terms of the subsequently described reliability plotting, this represents the visual gaps between the identify function and plotted model calibration line.

Reliability Plotting allows for visualizing the calibration of neural networks in a qualitative manner [15, 42]. The plot shows expected accuracy as a function of the confidence. In the case of a perfectly calibrated model, the plotted line will be identical to the identity function. Deviations from the diagonal identity function line represents miscalibrations which have occurred. The reliability diagram implementation and formulation is based on this paper [22].

Accuracy and AUROC are the two statistical metrics used to assess the general performance of the CNNs. Accuracy measures the fraction of predictions from the validation dataset which the model predicted correctly after training is completed. The area under the receiver operating characteristic (AUROC) is a robust measure of the ability for the binary classifier to discriminate between class labels [23].

3.3. Medical Image Datasets

As follows are brief descriptions of the open-source medical image datasets used in our experiments. For training of the CNNs, 80% of the dataset is partitioned and 20% is used to perform the validation respectively.

Skin Cancer Dataset The Skin Cancer Dataset was sourced from the International Skin Imaging Collaboration (ISIC) organization [1]. The dataset is open-source and consists of 3,297 processed dermatological skin images of mole lesions partitioned into malignant (diseased) and benign (normal) classes. Factors which differentiate images are mainly based on the pigmented skin lesions [29].

CXR Pneumonia Dataset The chest radiograph (CXR) dataset was gathered from the open-source "Chest X-Ray Images for Classification" repository from UCSD [31]. The dataset consists of 5,863 x-ray images (both anterior and posterior) from the normal and pneumonia classes. Images between class labels are differentiated based on hazy shadowing and opacity's found in x-rays with pneumonia.

MRI Tumor Dataset The magnetic resonance imaging (MRI) dataset of the human brain was sourced from an open-source repository on Kaggle [51]. The dataset contains 3,264 images split into tumorous and no tumor classes. The main differentiating factor between classes are the circular tumorous lesions which are typically in a difference shade compared to other regions [6].

CT COVID-19 Dataset The CT (computed tomography) dataset of COVID-19 is from the open-source UCSD COVID-CT repository [59]. The dataset consists of 812 CT scans split into the COVID-19 positive and negative classes. COVID-19 is identified in a CT based on ground-glass opacity, vascular enlargements, and white or hazy shadowing of the lung [25].

3.4. Model Architectures and Augmentations

To perform the experiments, the widely leveraged CNN architecture, ResNet is utilized [24]. ResNet is applied to various medical image analysis tasks for transfer learning thus providing a robust ground-truth for experimentation [33]. Taking into account varying CNN sizes, both ResNet-50 and ResNet-101 are benchmarked across the modern augmentation techniques. Implementations of ResNet follow the standard Keras Tensorflow [2] applications plugin¹. The implementation of CutMix [60], CutOut [17], and MixUp [62] augmentations follow open-source developments based on the original formulations².

¹<https://keras.io/api/applications>

²<https://github.com/ayulockin/DataAugmentationTF>

3.5. Training Details

All models across each dataset are trained for 100 epochs with cross entropy loss. Each dataset used contains two distinctive class labels. Thus, models are trained with two output logits for each input. Experiments are carried out using the stochastic gradient descent (SGD) optimizer [32], a batch size of 64, and learning rate 0.001. Input images are scaled to 224x224 pixels. For CutOut, a mask size of 50x50 pixels is used. Other augmentations follow the same parameters from open-source implementations.

4. Results

As follows is a summary of our systematic experimentation of modern augmentations on medical image analysis using performance-based and calibration-based evaluations across each medical image modality. Performance metrics are reported in Table 1 and calibration error is reported in Table 2 with reliability plots displayed in Figure 2.

4.1. Skin Cancer Dataset

4.1.1 Performance

The performance metrics (accuracy and AUROC) for both ResNet-50 [24] and ResNet-101 [24] for each augmentation technique on the skin cancer modality [1] are shown in 1a. For the ResNet-50 baseline on the skin cancer dataset, an accuracy of 79.2% and AUROC of 88.9% was achieved (Row 1). All augmentations presented minor increases in accuracy and AUROC, the most significant being MixUp [62] at an accuracy of 80.3% (+1.1% over ResNet-50). In terms of AUROC, CutMix [60] achieved an AUROC of 89.8% (+0.9% over ResNet-50). In summary, for ResNet-50, no highly significant benefits in terms of performance-based metrics were observed using modern augmentations.

For ResNet-101 on the skin cancer modality, the baseline achieved an accuracy of 79.8% and AUROC of 88.5%. For this model, the CutMix [60] and CutOut [17] augmentations performed slightly worse than the baseline in terms of accuracy. CutOut [17] also performed worse than the baseline for AUROC. MixUp [62] performed the best for both performance-based metrics at an accuracy of 82.5% (+2.7% over ResNet-101) and AUROC of 89.7% (+1.2% over ResNet-101). In summary, for ResNet-101, MixUp [62] presented fairly significant benefits in terms of performance-based metrics.

4.1.2 Confidence Calibration

The ECE calibration [22] results for the modern augmentations on the skin cancer modality are shown in Table 2 (Row 1 and Row 2). The lower the ECE, the higher level of calibration for the model. For ResNet-50, the baseline ECE was 0.1812. All modern augmentation techniques

lowered the ECE, the highest decrease was observed in CutMix [60] at 0.1286 (-0.0526). For ResNet-101, the baseline ECE was 0.1676. MixUp [62] and CutMix [60] both lowered the ECE however, CutOut [17] increased the ECE to 0.1967 (+0.0291). The most significant decrease in ECE for ResNet-101 was observed in CutMix [60] at 0.0973 (-0.0703). For both models, across the augmentations, CutMix [60] presented the most significant decreases in ECE providing higher calibration. The reliability plots [22] for the skin cancer modality are shown in Figure 2a.

4.2. CXR Pneumonia Dataset

4.2.1 Performance

The performance-based metrics for the CXR modality [31] are reported in Table 1b. The ResNet-50 baseline performed with an accuracy of 92.7% and AUROC of 94.4%. Both MixUp [62] and CutMix [60] presented significant increases in both accuracy and AUROC however, CutOut [17] decreased both accuracy and AUROC. The most significant increase in the performance-based metrics was observed in MixUp [62] with an accuracy of 94.4% (+1.7%) and AUROC of 98% (+3.6%). In summary, MixUp [62] presented the highest benefits in terms of performance.

For ResNet-101, the baseline accuracy was 87.2% and AUROC was 90.2%. All augmentations presented increases in both performance-based metrics. The highest increase in accuracy was observed in MixUp [62] at 93.9% (+6.7%). The highest increase in AUROC was observed in CutMix [60] at 97.7% (+7.5%). In summary, both MixUp [62] and CutMix [60] presented increases in performance.

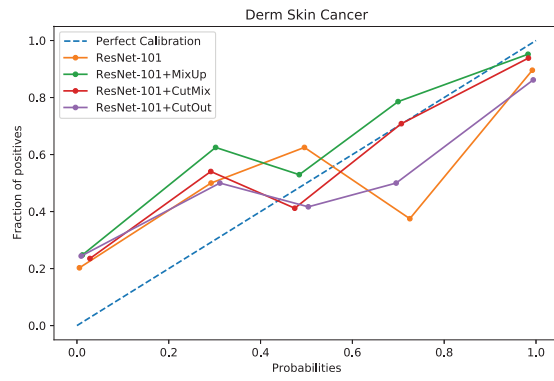

4.2.2 Confidence Calibration

The ECE calibration results for the modern augmentations on the CXR pneumonia modality are shown in Table 2 (Row 2 and Row 3). The ResNet-50 baseline had an ECE of 0.0675. Both MixUp [62] and CutMix [60] decreased the ECE. However, CutOut [17] increased the ECE slightly. The lowest ECE was observed for CutMix at 0.0351 (-0.0324). ResNet-101 had a baseline ECE of 0.1150. All augmentations reduced the ECE, MixUp [62] had the most significant decrease at 0.0340 (-0.081). In summary, both MixUp [62] and CutMix [60] presented the most benefits for calibration. The reliability plots displaying the level of calibration for the CXR pneumonia modality are shown in Figure 2b for both the pneumonia and normal class labels.

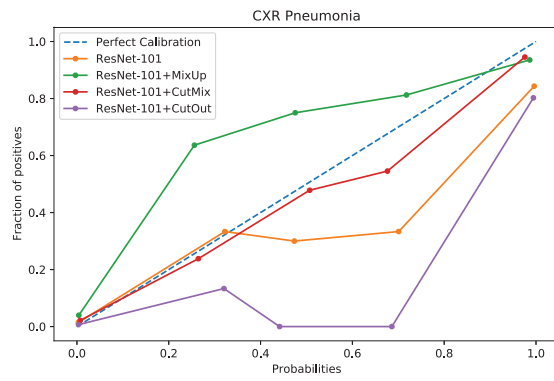
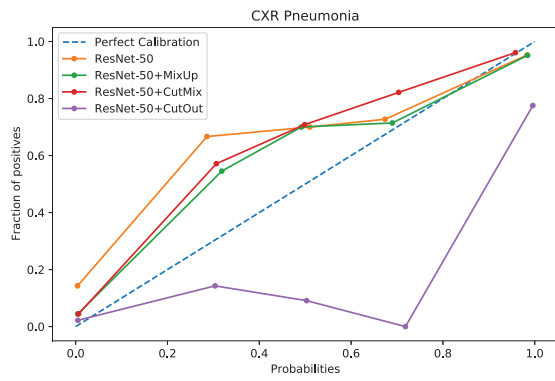
4.3. MRI Tumor Dataset

4.3.1 Performance



The performance-based metrics for the MRI modality are reported in Table 1c. For ResNet-50, the baseline performed with an accuracy of 64.7% and AUROC of 68.4%. In this





(a) Skin Cancer



(b) CXR Pneumonia



(c) MRI Tumor



(d) CT COVID-19

Figure 2. **Calibration Reliability Plots.** Shown for both models across each augmentation technique for all medical modalities. Left: ResNet-50 plot, Right: ResNet-101 plot.

Model	Augmentation	Accuracy	AUROC	Model	Augmentation	Accuracy	AUROC
ResNet-50 [24]	None	0.792	0.889	ResNet-50 [24]	None	0.927	0.944
ResNet-50 [24]	MixUp [62]	0.803	0.890	ResNet-50 [24]	MixUp [62]	0.944	0.980
ResNet-50 [24]	CutMix [60]	0.797	0.898	ResNet-50 [24]	CutMix [60]	0.941	0.977
ResNet-50 [24]	CutOut [17]	0.801	0.894	ResNet-50 [24]	CutOut [17]	0.917	0.941
ResNet-101 [24]	None	0.798	0.885	ResNet-101 [24]	None	0.872	0.902
ResNet-101 [24]	MixUp [62]	0.825	0.897	ResNet-101 [24]	MixUp [62]	0.939	0.976
ResNet-101 [24]	CutMix [60]	0.796	0.889	ResNet-101 [24]	CutMix [60]	0.933	0.977
ResNet-101 [24]	CutOut [17]	0.760	0.879	ResNet-101 [24]	CutOut [17]	0.886	0.915

(a) Skin Cancer

(b) CXR Pneumonia

Model	Augmentation	Accuracy	AUROC	Model	Augmentation	Accuracy	AUROC
ResNet-50 [24]	None	0.647	0.684	ResNet-50 [24]	None	0.700	0.724
ResNet-50 [24]	MixUp [62]	0.607	0.671	ResNet-50 [24]	MixUp [62]	0.680	0.757
ResNet-50 [24]	CutMix [60]	0.725	0.825	ResNet-50 [24]	CutMix [60]	0.653	0.754
ResNet-50 [24]	CutOut [17]	0.705	0.748	ResNet-50 [24]	CutOut [17]	0.633	0.656
ResNet-101 [24]	None	0.705	0.791	ResNet-101 [24]	None	0.653	0.706
ResNet-101 [24]	MixUp [62]	0.627	0.738	ResNet-101 [24]	MixUp [62]	0.706	0.765
ResNet-101 [24]	CutMix [60]	0.607	0.644	ResNet-101 [24]	CutMix [60]	0.613	0.708
ResNet-101 [24]	CutOut [17]	0.568	0.633	ResNet-101 [24]	CutOut [17]	0.673	0.746

(c) MRI Tumor

(d) CT COVID-19

Table 1. **Performance-based metrics** of various state-of-the-art CNN models across each medical modality.

Dataset	Model	Baseline	MixUp [62]	CutMix [60]	CutOut [17]
Derm	ResNet-50 [24]	0.1812	0.1424 (-0.0388)	0.1286 (-0.0526)	0.1726 (-0.0086)
Derm	ResNet-101 [24]	0.1676	0.1020 (-0.0656)	0.0973 (-0.0703)	0.1967 (+0.0291)
CXR	ResNet-50 [24]	0.0675	0.0409 (-0.0266)	0.0351 (-0.0324)	0.0750 (+0.0075)
CXR	ResNet-101 [24]	0.1150	0.0340 (-0.081)	0.0448 (-0.0702)	0.1024 (-0.0126)
MRI	ResNet-50 [24]	0.3419	0.3675 (+0.0256)	0.1259 (-0.2416)	0.2874 (-0.0801)
MRI	ResNet-101 [24]	0.2665	0.3675 (+0.101)	0.3487 (+0.0822)	0.3770 (+0.1105)
CT	ResNet-50 [24]	0.2866	0.2361 (-0.0505)	0.1909 (-0.0957)	0.3367 (+0.0501)
CT	ResNet-101 [24]	0.3237	0.1975 (-0.1262)	0.2382 (-0.0855)	0.2464 (-0.0773)

Table 2. **Expected Calibration Error (ECE)** ($M = 15$ bins) across various medical imaging modalities and CNN architectures.

scenario, MixUp [62] reduced the performance for both accuracy and AUROC while CutMix [60] and CutOut [17] increased the performance. The highest performing augmentation was observed in CutMix [60] at an accuracy of 72.5% (+7.8%) and AUROC of 82.5% (+14.1%).

For ResNet-101, the baseline performed with an accuracy of 70.5% and AUROC of 79.1%. All calibration metrics performed with lower performance-based metrics leaving the baseline as the highest performing model. The highest decrease in performance was observed in CutOut [17] at an accuracy of 58.6% (-11.9%) and AUROC of 63.3% (-15.8%).

4.3.2 Confidence Calibration

The ECE calibration results for the modern augmentations on the MRI tumor modality are shown in Table 2 (Row 5 and Row 6). For ResNet-50, the baseline ECE was

0.3419. In this scenario, CutMix [60] and CutOut [17] lowered the ECE and MixUp [62] increased the ECE to 0.3675 (+0.0256). The highest decrease in ECE was observed for CutMix [60] at an ECE of 0.1259 (-0.2416). ResNet-101 has a baseline of 0.2665. Interestingly, all augmentations increased the ECE the most significant increase was observed in CutOut [17] at 0.3770 (+0.1105). The baseline for ResNet-101 has the lowest ECE. The reliability plots for the MRI tumor modality are shown in Figure 2c.

4.4. CT COVID-19 Dataset

4.4.1 Performance

The performance-based metrics for the CT modality [59] are reported in Table 1d. For ResNet-50, the baseline performed with an accuracy of 70.0% and AUROC of 72.4%. In terms of accuracy, all augmentations presented lower accuracy compared to the baseline the most signif-

ificant observed in CutOut [17] with an accuracy of 63.3% (-6.7%). For AUROC, both MixUp [62] and CutMix [60] increased performance and CutOut [17] reduced the AUROC to 65.6% (-6.8%). The highest increase in AUROC was observed in MixUp [62] at 75.7% (+3.3%).

For ResNet-101, the baseline performed with an accuracy of 65.3% and AUROC of 70.6%. In terms of accuracy, MixUp [62] and CutOut [17] presented increases while CutMix decreased the accuracy to 61.3% (-4%). The highest increase in accuracy was observed in MixUp [62] at an accuracy of 70.6% (+5.3%). For AUROC, all augmentations presented increases in performance the most significant observed in MixUp [62] at an AUROC of 76.5% (+5.9%).

4.4.2 Confidence Calibration

The ECE calibration results for the modern augmentations on the CT COVID-19 modality are shown in Table 2 (Row 7 and Row 8). For ResNet-50, the baseline ECE was 0.2866. In this scenario, MixUp and CutMix reduced the ECE while CutOut increased the ECE (+0.0501). The most significant ECE decrease was observed in CutMix at an ECE of 0.1909 (-0.0957). ResNet-101 had an ECE baseline of 0.3237. All augmentations reduced the ECE the most significant decrease was observed in MixUp at 0.1975 (-0.1262). Reliability plots for the CT modality are shown in Figure 2

4.5. Interpretation

In summary, it is evident that in certain scenarios of medical image analysis, modern image augmentations can increase performance and significantly improve confidence calibration of CNNs. It is also important to understand that certain modern augmentations can decrease the performance and lead to miscalibration of CNNs. Table 3 shows a numerical summary of the amount of times a specific modern augmentation increased or decreased the level of calibration across all experiments.

Augmentation	↑Calib	↓Calib
MixUp [62]	6	2
CutMix [60]	7	1
CutOut [17]	4	4

Table 3. Numerical summary of calibration results from experimentation for augmentations across all modalities.

From these results, it is evident that CutMix increased the level of calibration (decreased expected calibration error) most frequently (7 out of 8 times). MixUp also presented significant impact on calibration having increased the level of calibration for 6 out of 8 experiments. However, CutOut increased calibration only 4 out of 8 times. Such performance shows that not all modern augmentations can positively effect calibration. CutOut could potentially be detrimental to the calibration of CNNs for medical image

analysis tasks. We hypothesize that the reason for CutOut reducing performance is because it can potentially remove clinically relevant regions from the images. On the other hand, MixUp and CutMix modify visual information but do not remove regions completely.

5. Conclusion

In this paper, we have compared the effects of several modern data augmentations on the confidence calibration of CNNs for various medical image analysis tasks using open-source datasets. CNNs are often prone to overconfidence and unreliable uncertainty estimates leading to a low amount of awareness. Improper quantification of uncertainty can be a high risk in the clinical setting and could potentially lead to medical errors. Through our in-depth experiments on the calibration of CNNs for medical image analysis using both ECE and reliability plotting, it is evident that certain modern augmentations (e.g. MixUp and CutMix) present significant benefits in terms of calibration, while others (CutOut) could worsen performance. Additionally, through the use of conventional performance-based metrics, it is evident that modern augmentations can also significantly increase the accuracy of CNNs. In conclusion, the usage of modern augmentations in medical image analysis CNNs can be beneficial in improving the reliability of models to improve clinical decision-making, however they should be benchmarked before implementation.

References

- [1] Isic archive.
- [2] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [3] Syed Muhammad Anwar, Muhammad Majid, Adnan Qayyum, Muhammad Awais, Majdi Alnowami, and Muhammad Khurram Khan. Medical image analysis using convolutional neural networks: a review. *Journal of medical systems*, 42(11):1–13, 2018.
- [4] Valentina Bellemo, Zhan W Lim, Gilbert Lim, Quang D Nguyen, Yuchen Xie, Michelle YT Yip, Haslina Hamzah, Jinyi Ho, Xin Q Lee, Wynne Hsu, et al. Artificial intelligence using deep learning to screen for referable and vision-threatening diabetic retinopathy in africa: a clinical validation study. *The Lancet Digital Health*, 1(1):e35–e44, 2019.

- [5] Yoshua Bengio, Frédéric Bastien, Arnaud Bergeron, Nicolas Boulanger-Lewandowski, Thomas Breuel, Youssouf Chherawala, Moustapha Cisse, Myriam Côté, Dumitru Erhan, Jeremy Eustache, et al. Deep learners benefit more from out-of-distribution examples. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 164–172. JMLR Workshop and Conference Proceedings, 2011.
- [6] Debnath Bhattacharyya and Tai-hoon Kim. Brain tumor detection using mri image analysis. In *International Conference on Ubiquitous Computing and Multimedia Applications*, pages 307–314. Springer, 2011.
- [7] Nicholas Bien, Pranav Rajpurkar, Robyn L Ball, Jeremy Irvin, Allison Park, Erik Jones, Michael Bereket, Bhavik N Patel, Kristen W Yeom, Katie Shpanskaya, et al. Deep-learning-assisted diagnosis for knee magnetic resonance imaging: development and retrospective validation of mrnet. *PLoS medicine*, 15(11):e1002699, 2018.
- [8] Luigi Carratino, Moustapha Cissé, Rodolphe Jenatton, and Jean-Philippe Vert. On mixup regularization. *arXiv preprint arXiv:2006.06049*, 2020.
- [9] Heang-Ping Chan, Ravi K Samala, Lubomir M Hadjiiski, and Chuan Zhou. Deep learning in medical image analysis. *Deep Learning in Medical Image Analysis*, pages 3–21, 2020.
- [10] Olivier Chapelle, Jason Weston, Léon Bottou, and Vladimir Vapnik. Vicinal risk minimization. *Advances in neural information processing systems*, 13, 2000.
- [11] Phillip Chlap, Hang Min, Nym Vandenberg, Jason Dowling, Lois Holloway, and Annette Haworth. A review of medical image data augmentation techniques for deep learning applications. *Journal of Medical Imaging and Radiation Oncology*, 65(5):545–563, 2021.
- [12] Sanghyuk Chun, Seong Joon Oh, Sangdoon Yun, Dongyoon Han, Junsuk Choe, and Youngjoon Yoo. An empirical evaluation on robustness and uncertainty of regularization methods. *arXiv preprint arXiv:2003.03879*, 2020.
- [13] Leda Cosmides and John Tooby. Are humans good intuitive statisticians after all? rethinking some conclusions from the literature on judgment under uncertainty. *cognition*, 58(1):1–73, 1996.
- [14] Deepan Das, Haley Massa, Abhimanyu Kulkarni, and Theodoros Rekatsinas. An empirical analysis of the impact of data augmentation on knowledge distillation. *arXiv preprint arXiv:2006.03810*, 2020.
- [15] Morris H DeGroot and Stephen E Fienberg. The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 32(1-2):12–22, 1983.
- [16] John Denker and Yann LeCun. Transforming neural-net output levels to probability distributions. *Advances in neural information processing systems*, 3, 1990.
- [17] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- [18] Zach Eaton-Rosen, Felix Bragman, Sebastien Ourselin, and M Jorge Cardoso. Improving data augmentation for medical image segmentation. 2018.
- [19] Andre Esteva, Katherine Chou, Serena Yeung, Nikhil Naik, Ali Madani, Ali Mottaghi, Yun Liu, Eric Topol, Jeff Dean, and Richard Socher. Deep learning-enabled medical computer vision. *NPJ digital medicine*, 4(1):1–9, 2021.
- [20] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016.
- [21] Adrian Galdran, Gustavo Carneiro, and Miguel A González Ballester. Balanced-mixup for highly imbalanced medical image classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 323–333. Springer, 2021.
- [22] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017.
- [23] James A Hanley and Barbara J McNeil. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143(1):29–36, 1982.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [25] Xuehai He, Xingyi Yang, Shanghang Zhang, Jinyu Zhao, Yichen Zhang, Eric Xing, and Pengtao Xie. Sample-efficient deep learning for covid-19 diagnosis based on ct scans. *medrxiv*, 2020.
- [26] Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012.
- [27] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [28] Shih-Cheng Huang, Tanay Kothari, Imon Banerjee, Chris Chute, Robyn L Ball, Norah Borus, Andrew Huang, Bhavik N Patel, Pranav Rajpurkar, Jeremy Irvin, et al. Penet—a scalable deep-learning model for automated diagnosis of pulmonary embolism using volumetric ct imaging. *NPJ digital medicine*, 3(1):1–9, 2020.
- [29] Anthony F Jerant, Jennifer T Johnson, Catherine Demastes Sheridan, and Timothy J Caffrey. Early detection and treatment of skin cancer. *American family physician*, 62(2):357–368, 2000.
- [30] Alain Jungo and Mauricio Reyes. Assessing reliability and challenges of uncertainty estimations for medical image segmentation. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II 22*, pages 48–56. Springer, 2019.
- [31] Daniel Kermany, Kang Zhang, Michael Goldbaum, et al. Labeled optical coherence tomography (oct) and chest x-ray images for classification. *Mendeley data*, 2(2), 2018.
- [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- [33] Padmavathi Kora, Chui Ping Ooi, Oliver Faust, U Raghavendra, Anjan Gudigar, Wai Yee Chan, K Meenakshi, K Swaraja, Pawel Plawiak, and U Rajendra Acharya. Transfer learning techniques for medical image analysis: A review. *Biocybernetics and Biomedical Engineering*, 2021.
- [34] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.
- [35] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- [36] June-Goo Lee, Sanghoon Jun, Young-Won Cho, Hyunna Lee, Guk Bae Kim, Joon Beom Seo, and Namkug Kim. Deep learning in medical imaging: general overview. *Korean journal of radiology*, 18(4):570–584, 2017.
- [37] Gongbo Liang, Yu Zhang, Xiaoqin Wang, and Nathan Jacobs. Improved trainable calibration method for neural networks on medical imaging classification. *arXiv preprint arXiv:2009.04057*, 2020.
- [38] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghahfoori, Jeroen Awm Van Der Laak, Bram Van Ginneken, and Clara I Sánchez. A survey on deep learning in medical image analysis. *Medical image analysis*, 42:60–88, 2017.
- [39] David JC MacKay. A practical bayesian framework for back-propagation networks. *Neural computation*, 4(3):448–472, 1992.
- [40] Alireza Mehrtash, William M Wells, Clare M Tempany, Purang Abolmaesumi, and Tina Kapur. Confidence calibration and predictive uncertainty estimation for deep medical image segmentation. *IEEE transactions on medical imaging*, 39(12):3868–3878, 2020.
- [41] Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [42] Alexandru Niculescu-Mizil and Rich Caruana. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd international conference on Machine learning*, pages 625–632, 2005.
- [43] Allison Park, Chris Chute, Pranav Rajpurkar, Joe Lou, Robyn L Ball, Katie Shpanskaya, Rashad Jabarkheel, Lily H Kim, Emily McKenna, Joe Tseng, et al. Deep learning–assisted diagnosis of cerebral aneurysms using the headxnet model. *JAMA network open*, 2(6):e195600–e195600, 2019.
- [44] Zhi Zhen Qin, Melissa S Sander, Bishwa Rai, Collins N Titahong, Santat Sudrungrot, Sylvain N Laah, Lal Mani Adhikari, E Jane Carter, Lekha Puri, Andrew J Codlin, et al. Using artificial intelligence to read chest radiographs for tuberculosis detection: A multi-site evaluation of the diagnostic accuracy of three deep learning systems. *Scientific reports*, 9(1):1–10, 2019.
- [45] Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding transfer learning for medical imaging. *Advances in neural information processing systems*, 32, 2019.
- [46] Pranav Rajpurkar, Jeremy Irvin, Robyn L Ball, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis P Langlotz, et al. Deep learning for chest radiograph diagnosis: A retrospective comparison of the chexnext algorithm to practicing radiologists. *PLoS medicine*, 15(11):e1002686, 2018.
- [47] Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis Langlotz, Katie Shpanskaya, et al. Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint arXiv:1711.05225*, 2017.
- [48] Pranav Rajpurkar, Allison Park, Jeremy Irvin, Chris Chute, Michael Bereket, Domenico Mastrodicasa, Curtis P Langlotz, Matthew P Lungren, Andrew Y Ng, and Bhavik N Patel. Appendixnet: Deep learning for diagnosis of appendicitis from a small dataset of ct exams using video pretraining. *Scientific reports*, 10(1):1–7, 2020.
- [49] Hariharan Ravishankar, Prasad Sudhakar, Rahul Venkataramani, Sheshadri Thiruvankadam, Pavan Annangi, Narayanan Babu, and Vivek Vaidya. Understanding the mechanisms of deep transfer learning for medical images. In *Deep learning and data labeling for medical applications*, pages 188–196. Springer, 2016.
- [50] Axel-Jan Rousseau, Thijs Becker, Jeroen Bertels, Matthew B Blaschko, and Dirk Valkenburg. Post training uncertainty calibration of deep networks for medical image segmentation. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 1052–1056. IEEE, 2021.
- [51] Sartaj. Brain tumor classification (mri), May 2020.
- [52] Dinggang Shen, Guorong Wu, and Heung-Il Suk. Deep learning in medical image analysis. *Annual review of biomedical engineering*, 19:221–248, 2017.
- [53] 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.
- [54] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-first AAAI conference on artificial intelligence*, 2017.
- [55] Jayaraman J Thiagarajan, Kowshik Thopalli, Deepta Rajan, and Pavan Turaga. Training calibration-based counterfactual explainers for deep learning models in medical image analysis. *Scientific reports*, 12(1):597, 2022.
- [56] Sunil Thulasidasan, Gopinath Chennupati, Jeff A Bilmes, Tanmoy Bhattacharya, and Sarah Michalak. On mixup training: Improved calibration and predictive uncertainty for deep neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- [57] Vladimir N Vapnik. An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5):988–999, 1999.
- [58] Maya Varma, Mandy Lu, Rachel Gardner, Jared Dunnmon, Nishith Khandwala, Pranav Rajpurkar, Jin Long, Christopher Beaulieu, Katie Shpanskaya, Li Fei-Fei, et al. Automated abnormality detection in lower extremity radiographs using deep learning. *Nature Machine Intelligence*, 1(12):578–583, 2019.

- [59] Xingyi Yang, Xuehai He, Jinyu Zhao, Yichen Zhang, Shanghang Zhang, and Pengtao Xie. Covid-ct-dataset: a ct scan dataset about covid-19. *arXiv preprint arXiv:2003.13865*, 2020.
- [60] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019.
- [61] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021.
- [62] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- [63] Linjun Zhang, Zhun Deng, Kenji Kawaguchi, and James Zou. When and how mixup improves calibration. In *International Conference on Machine Learning*, pages 26135–26160. PMLR, 2022.