

Rethinking Confidence Calibration for Failure Prediction

Fei Zhu^{1,2}, Zhen Cheng^{1,2}, Xu-Yao Zhang^{1,2*}, and Cheng-Lin Liu^{1,2}

¹ NLPR, Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China

² University of Chinese Academy of Sciences, Beijing, 100049, China

{zhufei2018, chengzhen2019}@ia.ac.cn, {xyz, liuc1}@nlpr.ia.ac.cn

Abstract. Reliable confidence estimation for the predictions is important in many safety-critical applications. However, modern deep neural networks are often overconfident for their incorrect predictions. Recently, many calibration methods have been proposed to alleviate the overconfidence problem. With calibrated confidence, a primary and practical purpose is to detect misclassification errors by filtering out low-confidence predictions (known as failure prediction). In this paper, we find a general, widely-existed but actually-neglected phenomenon that most confidence calibration methods are useless or harmful for failure prediction. We investigate this problem and reveal that popular confidence calibration methods often lead to worse confidence separation between correct and incorrect samples, making it more difficult to decide whether to trust a prediction or not. Finally, inspired by the natural connection between flat minima and confidence separation, we propose a simple hypothesis: flat minima is beneficial for failure prediction. We verify this hypothesis via extensive experiments and further boost the performance by combining two different flat minima techniques. Our code is available at <https://github.com/Impression2805/FMFP>.

Keywords: Failure prediction · Confidence Calibration · Flat minima · Uncertainty · Misclassification Detection · Selective Classification

1 Introduction

Deep neural networks (DNNs), especially vision models, has been widely deployed in risk-sensitive applications such as computer-aided medical diagnosis [11, 44], autonomous driving [2, 29], and robotics [38]. For such applications, besides the prediction accuracy, another crucial requirement is to provide *reliable confidence* for users to make safe decisions. For example, an autonomous driving car should rely more on other sensors or trigger an alarm when the detection network is unable to confidently predict obstructions [29]. Another example is the control should be handed over to human doctors when the confidence of a disease diagnosis network is low [44]. Unfortunately, modern DNNs are generally *overconfident* for their predictions, and can easily assign high confidence for misclassified samples

* Corresponding author.

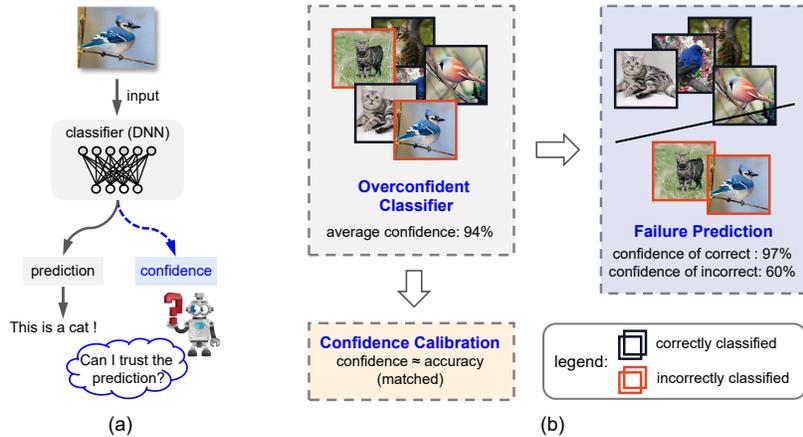


Fig. 1. Confidence calibration aims to reduce the mismatch between a model’s confidence and accuracy from the perspective of global statistics, while failure prediction is to distinguish correct from incorrect predictions from the perspective of confidence separability. They both focus on in-distribution data and share the same motivation to provide reliable confidence for trustworthy AI. Therefore, we explore a natural but ignored question: is calibrated confidence useful for failure prediction?

[7, 15, 17, 22]. The overconfident issue makes DNNs models untrustworthy, and therefore brings great concerns when DNNs are deployed in practical applications.

Recently, many approaches have been developed to alleviate the overconfidence problem by calibrating the confidence, *i.e.*, matching the accuracy and confidence scores to reflect the predictive uncertainty [43]. Specifically, one category of approaches [20, 41, 47, 48, 54, 61, 66, 68, 70, 75] aim to learn well-calibrated models during training. For instance, mixup [61], label smoothing [48] and focal loss [47] have been demonstrated to be effective for confidence calibration. Another class of approaches [15, 16, 35, 53, 56, 58] use post-processing techniques to calibrate DNNs. The most famous post-processing calibration method is temperature scaling [15] which learns a single scalar parameter to calibrate the probabilities.

In this paper, we study a natural but ignored question: *can we use calibrated confidence to detect misclassified samples by filtering out low-confidence predictions?* This, perhaps, is the most direct and practical way to evaluate the quality of the uncertainty. Actually, this problem is studied in the literature as *failure prediction* (also known as misclassification detection or selective classification) [7, 14, 22], whose purpose is to determine whether the prediction yielded by a classifier is correct or incorrect. Note that failure prediction aims to detect the erroneously classified natural example from seen class (*e.g.*, misclassified samples in test dataset), which is different from the widely studied out-of-distribution detection [23, 37, 39] that focuses on judging whether an input sample is from unseen classes. Compared with confidence calibration and out-of-distribution detection, failure prediction is far less explored in the literature.

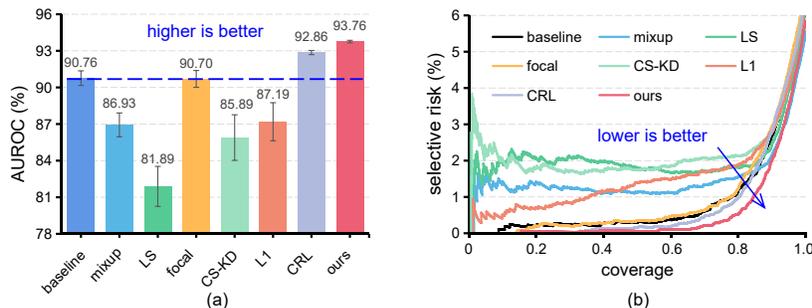


Fig. 2. A comparison of (a) AUROC and (b) risk-coverage curves. We observed that many popular confidence calibration methods are useless or harmful for failure prediction. We propose a simple flat minima based method that can outperform the state-of-the-art failure prediction method CRL [45]. ResNet110 [19] on CIFAR-10 [33].

As shown in Fig. 1, confidence calibration and failure prediction both focus on the confidence of in-distribution data and share the same motivation that enables the model to provide reliable confidence to make safe decisions. Therefore, common wisdom in the community suggests that calibrated confidence could be useful for failure prediction. However, we find a surprising pathology: many popular confidence calibration methods (including both training-time [24, 48, 54, 61, 68, 70] and post-processing [15] calibration methods) are more of a hindrance than a help for failure prediction, as illustrated in Fig. 2. Empirical study shows that those methods often reduce overconfidence by simply aligning the accuracy and average confidence. Such calibration could lead to worse separability between correct and misclassified samples, which is harmful for failure prediction. Consequently, one can not effectively detect misclassified samples by filtering out low-confidence predictions based on the calibrated confidence.

Finally, how can we improve the failure prediction performance of DNNs? Intuitively, failure prediction requires better discrimination between the confidence of correct and incorrect samples, which would increase the difficulty of changing the correct samples to be incorrect due to the larger confidence margins. Interestingly, this is closely related to the notion of “flatness” in DNNs, which reflects how sensitive the correct samples become misclassified when perturbing the model parameters [12, 27, 28]. Inspired by the natural connection between flat minima and confidence separation, we propose a simple hypothesis: *flat minima is beneficial for failure prediction*. We verify this hypothesis by extensive experiments and propose a simple and effective technique that combines different kinds of flat methods to achieve state-of-the-art performance on failure prediction.

Contributions. Motivated by the widely confirmed confidence calibration effect of recently proposed techniques, we rethink the confidence reliability by evaluating them on the challenging and practical failure prediction task. Surprisingly, we find that they often have negative effect on failure prediction. From a detailed analysis, we identify a compounding less-separability effect of training-time calibration methods [31, 40, 48, 61, 70], and further find that failure prediction can

not be improved by post-hoc calibration strategies like temperature scaling [15]. Finally, inspired by the connection between flat minima [12, 27, 28] and confidence separation, we propose to find flat minima to significantly reduce the confidence of misclassified samples while maintaining the confidence of correct samples. Extensive experiments show the strong performance of our method on both failure prediction and confidence calibration.

2 Problem Formulation and Background

Considering the multi-class classification problem, we assume a sample (\mathbf{x}, y) is drawn *i.i.d.* from an unknown joint distribution over $\mathcal{X} \times \mathcal{Y}$ where $\mathcal{X} = \mathbb{R}^d$ donates the feature space and $\mathcal{Y} = \{1, 2, \dots, k\}$ is a label space. Utilizing a standard softmax function, a deep neural network classifier $f : \mathcal{X} \rightarrow \mathcal{Y}$ produces a probability distribution over k classes. Specifically, given an input \mathbf{x} , f produces the predicted class probabilities $\hat{\mathbf{p}} = \hat{P}(y|\mathbf{x}, \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is the parameters of the classification model. With these probabilities, $\hat{y} = \arg \max_{y \in \mathcal{Y}} \hat{P}(y|\mathbf{x}, \boldsymbol{\theta})$ can be returned as the predicted class and the associated probability $\hat{p} = \max_{y \in \mathcal{Y}} \hat{P}(y|\mathbf{x}, \boldsymbol{\theta})$, *i.e.*, the maximum class probability, can be viewed as the predicted confidence.

2.1 Confidence Calibration

Definitions and notation. Intuitively, the predictive confidence of a well-calibrated model could be indicative of the actual likelihood of correctness [15]. For example, if a calibrated model predicts a set of inputs \mathbf{x} to be class y with 40% probability, then we expect 40% of the inputs indeed belong to class y . Formally, a model is perfectly calibrated if [15, 36]:

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

The most commonly used calibration metric is the Expected Calibration Error (ECE) [50], which approximates the miscalibration by binning the confidence in $[0, 1]$ under M equally-spaced intervals *i.e.*, $\{B_m\}_{m=1}^M$. Then the miscalibration is estimated by taking the expectation of the mismatch between the accuracy and averaged confidence in each bin: $\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{n} |\text{acc}(B_m) - \text{avgConf}(B_m)|$, where n is the number of all samples. Alternatives to ECE include the negative log likelihood (NLL) and brier score [3].

Improving calibration. Many strategies have been proposed to address the miscalibration of modern DNNs. (1) One category of approaches [47, 48, 54, 61, 66, 68, 70, 75] aim to learn well-calibrated models during training. For example, several works [61, 73, 74] found that the predicted scores of DNNs trained with mixup [72] are better calibrated. Muller *et al.* [48] showed the favorable calibration effect of label smoothing. Mukhoti *et al.* [47] demonstrated that focal loss [40] can automatically learn well-calibrated models. CS-KD [70] calibrates overconfident predictions by penalizing the predictive distribution between the samples within the same class. Recently, Joo *et al.* [31] explored the effect of explicit regularization

strategies (*e.g.*, L_p norm in the logits space) for calibration. (2) Another class of approaches [15, 16, 35, 53, 56, 58] rescal the predictions in a post-hoc manner. Among them, temperature scaling [15] is a effective and simple technique, which has inspired various post-processing approaches [34, 46, 58].

Empirical studies of calibration. In addition to the calibration strategies, there have been some empirical studies on calibration. Guo *et al.* [15] observed that larger networks tend to be less calibrated, even as classification accuracy is improved. Ovadia *et al.* [52] studied the calibration under distribution shift and empirically found the generally existing performance drop of different calibration methods under distribution shift. More recently, Wang *et al.* [65] found that it is harder to further calibrate the model with temperature scaling if it has been trained with regularization methods. Minderer *et al.* [43] found that the most recent non-convolutional models [10, 62] are well-calibrated, suggesting that architecture is a major factor of calibration performance. Differently from those works, we rethink the confidence calibration for failure prediction.

2.2 Failure Prediction

Definitions and notation. Failure prediction, also known as misclassification detection [22] or ordinal ranking [45], focus on distinguishing incorrect from correct predictions based on their confidence ranking. Intuitively, if the associated confidence of each misclassified sample is lower than that of any correctly classified samples, we can successfully predict each error made by the classification model at inference time. Formally, an optimal ordinal ranking model should reflect the following relationship for every two samples (\mathbf{x}_i, y_i) and (\mathbf{x}_j, y_j) :

$$\kappa(\hat{\mathbf{p}}_i|\mathbf{x}_i, \boldsymbol{\theta}) \geq \kappa(\hat{\mathbf{p}}_j|\mathbf{x}_j, \boldsymbol{\theta}) \iff P(\hat{y}_i = y_i|\mathbf{x}_i) \geq P(\hat{y}_j = y_j|\mathbf{x}_j), \quad (2)$$

where κ denotes a confidence-rate function (*e.g.*, the maximum class probability) that assess the degree of confidence of the predictions. Then, with a predefined threshold $\delta \in \mathbb{R}^+$, the users can reject the erroneously classification results based on the following decision function g :

$$g(\mathbf{x}) = \begin{cases} \text{accept, if } \kappa(\mathbf{x}) \geq \delta, \\ \text{reject, otherwise.} \end{cases} \quad (3)$$

Common metrics for failure prediction are the risk-coverage curve (AURC), the normalized AURC (E-AURC) [14, 45], the false positive rate at 95% true positive rate (FPR-95%TPR), and the area under the receiver operating characteristic curve (AUROC). In addition, there are some other metrics [22] to imply how the correct and incorrect predictions are separated.

Improving failure prediction. For DNNs, Hendrycks *et al.* [22] firstly established a standard **baseline** for failure prediction by using maximum softmax probability. Trust-Score [30] adopts the similarity between the classifier and a nearest-neighbor classifier as a confidence measure. The main drawback of Trust-Score is the lack of practicality and scalability in high-dimensional spaces. Some

works [7, 42] formulate the failure prediction as a supervised binary classification problem. Technically, ConfidNet [7, 8] and SS [42] train auxiliary models to predict confidence by learning the misclassified samples in training set. However, they may fail when the model has a high training accuracy, in which few or even no misclassified examples will exist in the training set. CRL [45] improves failure prediction by regularizing the model to learn an ordinal ranking relationship based on the historical correct rate during training. To the best of our knowledge, there are only a few works for failure prediction and no method can significantly or fairly outperform the baseline [22].

Why revisit confidence calibration for failure prediction? On the one hand, confidence calibration and failure prediction both focus on the confidence of in-distribution data and share the same motivation to provide reliable confidence for making safer decisions. From a practical perspective, with a calibrated classifier in hand, one natural way to verify its trustworthiness is to filter out predictions with low confidence. On the other hand, confidence calibration has drawn significant attention from the machine learning community, including improving calibration [15, 48, 61, 70], empirical studies of model calibration [15, 43, 52, 66] and measures of model calibration [3, 16, 51, 64]. However, there are few works for failure prediction, which is a practical, important, yet somewhat under-appreciated area of research. Therefore, revisiting calibration from the perspective of failure prediction not only helps understand the effect of calibration but also benefits the investigation of failure prediction.

3 Does Calibration Help Failure Prediction?

In recent years, there is a surge of research focused on alleviating the over-confidence problem of modern DNNs. As shown by many empirical results, existing methods do help the calibration of DNNs. In this section, we empirically investigate the reliability of the calibrated confidence for failure prediction.

3.1 Experimental Setup

Datasets and network architectures. We thoroughly conduct experiments on benchmark datasets CIFAR-10 and CIFAR-100 [33], and large-scale ImageNet [9] dataset. In terms of network architectures, we consider a range of models: PreAct-ResNet110 [19], WideResNet [71], DenseNet [26] and more recent architecture ConvMixer [63] for experiments on CIFAR-10 and CIFAR-100. For ImageNet, we used a ResNet-18 [18] model. Due to space limitation, we provide the results of more networks like MobileNet [25], EfficientNet [60] and dataset like Tiny-ImageNet [69] in the the supplementary material.

Evaluation metrics. We adopt the standard metrics in [22] and [14, 45] to measure failure prediction: AURC, E-AURC, AUROC, FPR-95%TPR, AUPR-Success and AUPR-Error. Lower values of AURC, E-AURC, FPR-95%TPR and higher values of AUROC, AUPR-Success, AUPR-Error indicate better failure prediction ability. Supplementary material provides definitions of these metrics.

Table 1. Failure prediction performance on CIFAR-10 and CIFAR-100 datasets. AURC and E-AURC values are multiplied by 10^3 , and all remaining values are percentage.

| CIFAR-10 | | | | | | | |
|-----------|---------------|-------------------|-------------------|-------------------|-------------------|---------------------|-------------------|
| Network | Method | AURC (↓) | E-AURC (↓) | FPR-95% TPR(↓) | AUROC (↑) | AUPR- Success(↑) | AUPR- Error(↑) |
| ResNet110 | baseline [22] | 9.94±1.29 | 7.94±1.23 | 45.01±2.55 | 90.76±0.60 | 99.16±0.13 | 43.66±1.86 |
| | mixup [61] | 15.05±0.09 | 13.31±0.29 | 41.18±2.54 | 86.93±0.98 | 98.60±0.03 | 43.19±3.52 |
| | LS [48] | 24.50±2.73 | 22.57±2.77 | 44.76±3.76 | 81.89±1.64 | 97.63±0.29 | 39.57±1.00 |
| | Focal [47] | 10.70±1.32 | 8.30±1.25 | 45.85±0.75 | 90.70±0.69 | 99.12±0.13 | 44.32±1.40 |
| | CS-KD [70] | 18.02±3.47 | 15.81±3.43 | 43.98±3.86 | 85.89±1.87 | 98.33±0.36 | 41.17±4.09 |
| L1 [70] | 14.21±1.93 | 12.11±1.92 | 45.71±0.60 | 87.19±1.56 | 98.72±0.20 | 41.93±1.51 | |
| WRNet | baseline [22] | 4.89±0.25 | 3.93±0.22 | 32.85±0.36 | 93.24±0.15 | 99.59±0.02 | 43.38±1.12 |
| | mixup [61] | 6.24±0.79 | 5.52±0.80 | 31.27±0.63 | 91.18±0.66 | 99.43±0.08 | 41.86±0.79 |
| | LS [48] | 16.69±2.86 | 15.69±2.86 | 34.24±2.84 | 85.30±1.79 | 98.38±0.29 | 43.19±1.27 |
| | Focal [47] | 6.91±0.56 | 5.83±0.55 | 35.43±2.52 | 91.86±0.66 | 99.39±0.06 | 43.11±0.37 |
| | CS-KD [70] | 10.29±0.35 | 9.22±0.38 | 38.65±2.67 | 88.13±0.74 | 99.04±0.04 | 38.88±0.66 |
| L1 [70] | 7.01±1.63 | 5.99±1.66 | 34.19±0.78 | 91.08±1.91 | 99.38±0.17 | 42.25±1.51 | |
| DenseNet | baseline [22] | 6.2±0.29 | 4.66±0.27 | 38.20±2.48 | 92.87±0.44 | 99.51±0.03 | 43.74±2.44 |
| | mixup [61] | 8.57±0.54 | 7.08±0.51 | 37.88±3.73 | 91.17±0.84 | 99.26±0.05 | 44.35±1.18 |
| | LS [48] | 19.35±2.29 | 17.67±2.19 | 40.20±2.10 | 84.45±1.45 | 98.15±0.23 | 41.35±1.16 |
| | Focal [47] | 7.17±0.28 | 5.35±0.19 | 41.75±0.88 | 92.31±0.17 | 99.44±0.02 | 43.27±2.04 |
| | CS-KD [70] | 13.55±1.02 | 11.64±0.92 | 43.11±3.75 | 88.02±0.46 | 98.77±0.10 | 40.88±4.30 |
| L1 [70] | 8.34±1.24 | 6.94±1.07 | 37.02±0.97 | 91.40±0.63 | 99.27±0.13 | 44.19±0.55 | |
| ConvMixer | baseline [22] | 8.33±1.44 | 6.29±1.30 | 42.32±3.26 | 92.02±0.96 | 99.34±0.14 | 43.80±1.49 |
| | mixup [61] | 9.87±0.14 | 8.40±0.25 | 37.57±2.10 | 90.25±0.81 | 99.12±0.02 | 45.01±2.40 |
| | LS [48] | 18.45±1.12 | 16.41±1.08 | 40.99±1.91 | 86.01±0.24 | 98.27±0.11 | 43.32±0.70 |
| | Focal [47] | 9.59±1.02 | 7.17±1.03 | 46.18±1.66 | 91.32±0.87 | 99.24±0.11 | 44.03±0.57 |
| | CS-KD [70] | 13.62±0.86 | 11.69±0.93 | 43.06±0.35 | 88.02±0.53 | 98.77±0.10 | 41.35±0.78 |
| L1 [70] | 12.92±2.41 | 11.07±2.36 | 41.39±0.91 | 88.35±1.57 | 98.83±0.25 | 42.91±1.60 | |
| CIFAR-100 | | | | | | | |
| Network | Method | AURC (↓) | E-AURC (↓) | FPR-95% TPR(↓) | AUROC (↑) | AUPR- Success(↑) | AUPR- Error(↑) |
| ResNet110 | baseline [22] | 93.90±2.37 | 50.88±2.03 | 66.02±1.53 | 85.00±0.35 | 93.42±0.28 | 66.54±0.29 |
| | mixup [61] | 95.03±2.77 | 57.57±3.33 | 63.68±0.36 | 84.03±0.80 | 92.65±0.40 | 64.55±1.46 |
| | LS [48] | 111.18±0.36 | 69.46±1.08 | 63.93±0.65 | 82.85±0.34 | 91.00±0.11 | 65.19±0.32 |
| | Focal [47] | 96.60±2.81 | 52.97±1.96 | 66.60±0.64 | 84.22±0.17 | 93.16±0.28 | 65.23±0.53 |
| | CS-KD [70] | 100.68±2.83 | 58.89±1.94 | 66.15±1.69 | 83.98±0.21 | 92.39±0.27 | 65.09±0.55 |
| L1 [70] | 119.49±4.04 | 71.47±3.31 | 65.78±0.96 | 82.73±0.51 | 90.55±0.44 | 66.44±0.59 | |
| WRNet | baseline [22] | 51.97±1.74 | 30.18±1.09 | 59.19±0.20 | 87.75±0.30 | 96.38±0.14 | 62.99±0.98 |
| | mixup [61] | 50.54±0.98 | 32.00±0.92 | 57.29±1.57 | 87.47±0.43 | 96.21±0.10 | 61.72±1.25 |
| | LS [48] | 58.29±3.46 | 36.69±2.66 | 58.47±0.98 | 86.76±0.19 | 95.59±0.34 | 61.95±0.72 |
| | Focal [47] | 54.54±1.44 | 32.40±1.28 | 61.87±1.06 | 86.89±0.22 | 96.12±0.16 | 60.56±0.65 |
| | CS-KD [70] | 58.30±1.29 | 35.62±0.68 | 60.44±0.95 | 86.98±0.19 | 95.70±0.09 | 62.23±1.17 |
| L1 [70] | 61.63±0.46 | 38.78±0.45 | 59.48±2.07 | 86.21±0.28 | 95.31±0.05 | 62.55±0.77 | |
| DenseNet | baseline [22] | 67.41±0.67 | 35.82±0.54 | 61.55±2.01 | 86.46±0.31 | 95.55±0.05 | 65.60±1.26 |
| | mixup [61] | 64.84±4.26 | 37.06±2.74 | 62.94±2.55 | 86.26±0.63 | 95.47±0.36 | 62.64±0.81 |
| | LS [48] | 76.24±2.31 | 44.40±1.01 | 62.41±0.51 | 85.32±0.10 | 94.47±0.15 | 63.59±0.43 |
| | Focal [47] | 73.43±2.05 | 40.70±1.40 | 65.67±1.48 | 85.62±0.30 | 94.95±0.19 | 62.76±0.61 |
| | CS-KD [70] | 75.22±1.02 | 41.38±1.28 | 62.75±0.85 | 86.20±0.40 | 94.82±0.15 | 64.50±0.30 |
| L1 [70] | 68.73±1.10 | 40.47±1.37 | 63.39±0.92 | 85.46±0.50 | 95.06±0.22 | 61.90±0.44 | |
| ConvMixer | baseline [22] | 76.96±1.64 | 41.20±1.64 | 63.57±0.52 | 86.28±0.18 | 94.81±0.06 | 65.39±0.83 |
| | mixup [61] | 70.87±2.25 | 39.38±2.65 | 63.80±1.72 | 86.12±0.75 | 95.12±0.75 | 63.71±0.75 |
| | LS [48] | 83.50±3.88 | 47.37±2.45 | 66.49±1.50 | 84.69±0.62 | 94.03±0.33 | 62.87±1.24 |
| | Focal [47] | 83.79±2.03 | 45.51±0.88 | 66.26±1.50 | 85.17±0.34 | 94.24±0.12 | 64.36±1.76 |
| | CS-KD [70] | 74.02±1.18 | 42.18±0.56 | 64.83±1.46 | 85.44±0.14 | 94.76±0.09 | 62.64±0.69 |
| L1 [70] | 82.68±0.35 | 46.62±0.50 | 65.03±0.44 | 85.02±0.14 | 94.12±0.05 | 64.01±0.42 | |

Implementation details. All models are trained using SGD with a momentum of 0.9, an initial learning rate of 0.1, and a weight decay of $5e-4$ for 200 epochs with the mini-batch size of 128 on CIFAR-10 and CIFAR-100. The learning rate is reduced by a factor of 10 at 80, 130, and 170 epochs. We randomly sample 10% of training samples as a validation dataset for each task because it is a requirement for post-calibration methods like temperature scaling [15]. For each experiment, the mean and standard deviation over three random runs are reported.

Evaluated calibration methods. We evaluate various calibration methods include training-time regularization like mixup [61], label-smoothing (LS) [48], focal loss [40], CS-KD [70], L_1 norm [31] and post-hoc method like temperature

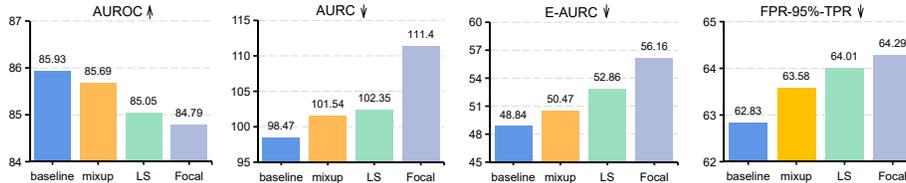


Fig. 3. Large-scale experiments on ImageNet. AURC and E-AURC values are multiplied by 10^3 for clarity, and all remaining values are percentage.

scaling (TS) [15]. Those methods have been verified to be effective to address the miscalibration problem of DNNs. Particularly, compared with post-hoc methods, the effect of training-time regularization is irreversible. Therefore, we mainly focus on their performance on failure prediction. Supplementary material provides detail introduction and hyperparameter setting of each method.

3.2 Results and Analysis

In our experiments, we confirmed the positive calibration effects of the evaluated methods. For example, on CIFAR-10, with mixup, the ECE (%) can be reduced from 4.14 to 2.97 for ResNet110 and from 2.96 to 1.39 for DenseNet; with focal loss, the ECE (%) can be reduced from 4.14 to 1.60 for ResNet110 and from 2.96 to 1.36 for DenseNet. These observations are consistent with that in [40, 61].

Popular calibration methods can harm failure prediction. In practice, users would naturally expect that the calibrated confidence can be used to filter out low-confidence predictions in risk-sensitive applications. However, if we shift focus to Table 1, it is evident that those methods generally lead to *worse* failure prediction performance under various metrics. For example, when training with mixup and LS on CIFAR-10/ResNet110, the AUROC (\uparrow) drops 3.83 and 9.07 percentages, respectively. And the AURC (\downarrow) increases 5.51 and 14.56 percentages, respectively. This is counter-intuitive as we expect those methods, which successfully calibrate the confidence, could be useful for failure prediction.

The same observations generalize to large-scale dataset. Here we verify our observation that calibration methods often harm failure prediction on ImageNet [9] dataset, which comprises 1000 classes and over 1.2 million images. We train a ResNet-18 [18] that achieve a 70.20% top-1 classification accuracy. The results are shown in Fig. 3, from which we can observe similar negative effect of calibration methods on failure prediction. More results on other networks, which exhibit similar pattern, can be found in supplementary material.

Selective risk analysis. To make intuitive sense of the effect of those calibration methods on failure prediction, Fig. 2(b) plots the risk-coverage curve. Specifically, selective risk is the empirical loss or error rate that trust the prediction, while coverage is the probability mass of non-rejected predictions [14, 45]. Intuitively, a better failure predictor should have low risk at a given coverage. As can be seen

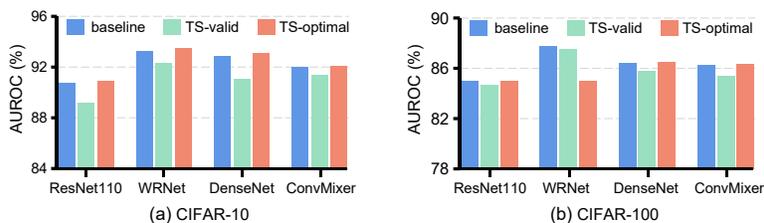


Fig. 4. Temperature scaling can hardly improve failure prediction.

from Fig. 2(b), the baseline has the lowest risks compared to other calibration methods, which indicates that using the confidence calibrated by those methods would unfortunately increase the risk when making decisions.

Does temperature scaling improve failure prediction? As a representative post-hoc calibration technique, TS [15] is simple and effective. Specifically, TS calibrates probabilities by learning a single scalar parameter T for all classes on a hold-out validation set. In Fig. 4, we show the failure prediction performance of TS on different networks and datasets. Specifically, using validation set and test set to learn the parameter T are denoted as *TS-valid* and *TS-optimal*, respectively. By directly using test set, TS-optimal yields the optimal T for failure prediction. As shown in Fig. 4, compared with baseline, TS-valid has negative effectiveness while TS-optimal has negligible improvement.

3.3 Calibration Harms Failure Prediction: A Closer Look

Illustrative experiments. To empirically understand the negative effect of those calibration methods, Fig. 5 shows the confidence distribution of test samples. It can be seen that LS leads to a more severe overlap between the confidence of correct and incorrect samples. Fig. 6(a) plots the average confidence of correctly classified samples during training, where their confidences are obviously reduced. This can also be seen from Fig. 6(b-c), in which mixup and LS lead to better overall ECE but worse ECE of correct samples. This indicates that those calibration methods yield under-confident correct prediction. In conclusion, those calibration methods reduce the overconfident of DNNs, but lead to worse separability between correct and misclassified samples, making it hard to detect misclassified samples based on the calibrated confidence.

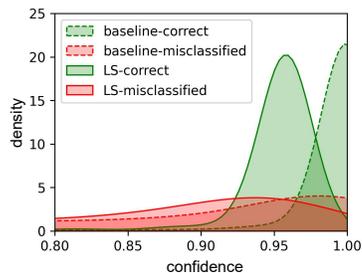


Fig. 5. LS results in under-confident of correct samples. ResNet110 on CIFAR-10.

Discussion on calibration for failure prediction. The best failure prediction is achieved when correct and wrong predictions are clearly separated according to the confidence level. However, calibration focuses on matching the average accuracy and confidence, and the best ECE score is achieved when correct and

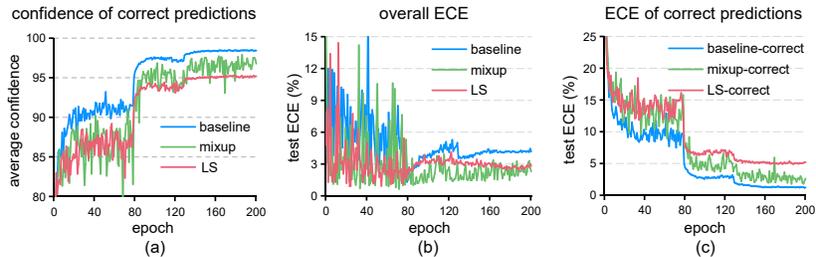


Fig. 6. (a) Comparison of average confidence of correctly classified samples during training. (b) Mixup and LS successfully reduce the overall ECE, (c) but result in worse ECE of correct samples. ResNet110 on CIFAR-10.

wrong predictions are “mixed” in the right way such that those with confidence at level $[c, c + w)$ (confidence window in a bin, where c donates the value of confidence) should have a mix of correct and wrong predictions with ratios $c : 1 - c$. Regularization methods such as mixup [61], LS [48], focal loss [40], CS-KD [70] and L_p norm [31] typically improve calibration by penalizing the confidence of the whole samples to a low level. However, this will lead to undesirable effects: erasing important information about the hardness of samples [59], which would introduce an undesired mixing in ranking, and thus result in drops in failure prediction qualities. Nevertheless, this does not mean that a better ECE must lead to worse failure prediction. A proper strategy might benefit both confidence-accuracy matching and confidence separability. As shown in Section 4, calibration and failure prediction could be improved concurrently.

4 Improving Failure Prediction by Finding Flat Minima

As reported in Section 3, none of those popular calibration methods seem to address failure prediction problem (stably) better than simple baseline [22]. Does there exist a more principled and hassle-free strategy to improve failure prediction?

4.1 Motivation and Methodology

Rationale: why. Confidence separability between correct and incorrect samples is crucial for failure prediction. Let us consider how confidence separability affects the confidence robustness of correct samples. Specifically, for a correctly classified sample, to become misclassified, it must reduce the probability on the ground truth class and increase its probability on another (wrong) class. During this process, the confidence margin plays a crucial role: a larger confidence margin could make it harder to change the predicted class label. Interestingly, flatness of a model reflects how sensitive the correctly classified samples become misclassified when perturbing the weights of a model [12, 27, 28]. As illustrated in Fig. 7, with flat minima, a correct sample is difficult to be misclassified under weight perturbations and vice versa. Therefore, we conjecture that the confidence gap

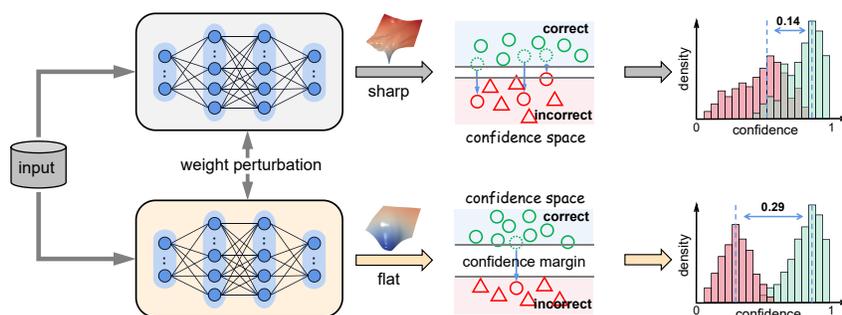


Fig. 7. An intuitive relationship between flatness and failure prediction. The stability of predictions to parameter perturbations can be seen as a *confidence margin* condition. This inspires us to improve failure prediction by finding flat minima.

between correct and incorrect samples of a flat minima is larger than that of a sharp minima.

Reliable overfitting phenomenon.

As shown in Fig. 8, we observed an interesting phenomenon that the AUROC can be easily overfitting during the training of a model (ResNet110 on CIFAR-10). Concretely, the test accuracy continually increases while the AUROC decreases at the last phases, making it difficult for failure prediction.

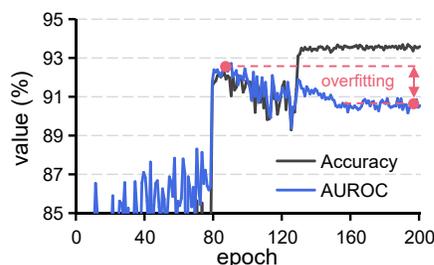


Fig. 8. *Reliable overfitting* phenomenon. We term this phenomenon as “*reliable overfitting*”, which exists on different model and dataset settings and somewhat similar to the *robust overfitting* [57] in adversarial robustness literature. Since flat minima has been verified to be effective for alleviating robust overfitting [6, 67], we expect that flat minima could also benefit failure prediction.

Approach: how. There are several methods have been proposed to seek flat minima for DNNs [5, 12, 28, 55]. We select *stochastic weight averaging* (SWA) [28] and *sharpness-aware minimization* (SAM) [12] as two representative methods due to the simplicity for proofs-of-concept. Specifically, SWA simply averages multiple parameters of the model along the training trajectory: $\theta_{swa}^t = \frac{\theta_{swa}^{t-1} \times n + \theta^t}{n+1}$, where n indexes the number of past checkpoints to be averaged, t is the training epoch, θ is the current weights, θ_{swa} is the averaged weights. While SAM finds the flat minima by directly perturbing the weights: $\min_{\theta} \max_{\|\epsilon\|_p \leq \rho} \mathcal{L}(\theta + \epsilon) + \frac{\lambda}{2} \|\theta\|$.

Although the SWA and SAM find flat minima based on different mechanism, we find they both improve the failure prediction performance. This also motivates us to combine them to get better performance. We refer the combine of them as **FMFP** (*Flat Minima for Failure Prediction*). Supplementary material presents pseudo-code for the FMFP algorithm.

Table 2. Confidence estimation results. AURC and E-AURC values are multiplied by 10^3 , and NLL are multiplied by 10 for clarity. Remaining values are percentages.

| CIFAR-10 | | | | | | | | | | |
|-----------|---------------|------------------|------------------|-------------------|-------------------|---------------------|-------------------|------------------|------------------|------------------|
| Network | Method | AURC (↓) | E-AURC (↓) | FPR-95% TPR(↓) | AUROC (↑) | AUPR- Success(↑) | AUPR- Error(↑) | ECE (↓) | NLL (↓) | Brier (↓) |
| ResNet110 | baseline [22] | 9.94±1.29 | 7.94±1.23 | 45.01±2.55 | 90.76±0.60 | 99.16±0.13 | 43.66±1.86 | 4.14±0.06 | 2.88±0.08 | 10.47±0.28 |
| | CRL [45] | 7.54±0.20 | 5.29±0.16 | 45.25±3.24 | 92.86±0.17 | 99.44±0.02 | 44.93±1.31 | 1.65±0.08 | 2.05±0.03 | 10.01±0.09 |
| | SAM | 5.31±0.13 | 3.74±0.17 | 39.44±2.75 | 93.73±0.45 | 99.61±0.02 | 45.23±2.88 | 1.86±0.12 | 1.76±0.02 | 8.55±0.19 |
| | SWA | 6.38±0.18 | 4.38±0.07 | 39.36±0.29 | 93.69±0.08 | 99.54±0.01 | 47.49±2.03 | 1.33±0.10 | 1.83±0.03 | 9.17±0.15 |
| WRNet | ours | 5.57±0.11 | 3.92±0.09 | 39.50±0.52 | 93.76±0.10 | 99.58±0.01 | 44.45±1.89 | 0.50±0.13 | 1.68±0.02 | 8.50±0.13 |
| | baseline [22] | 4.89±0.25 | 3.93±0.22 | 32.85±0.36 | 93.24±0.15 | 99.59±0.02 | 43.38±1.12 | 2.71±0.06 | 1.82±0.03 | 7.15±0.07 |
| | CRL [45] | 4.52±0.10 | 3.23±0.10 | 32.92±2.18 | 94.24±0.23 | 99.66±0.01 | 44.07±1.61 | 0.44±0.14 | 1.52±0.02 | 7.41±0.14 |
| | SAM | 2.85±0.04 | 2.15±0.07 | 26.07±0.96 | 95.10±0.19 | 99.78±0.01 | 44.50±1.20 | 1.58±0.05 | 1.25±0.01 | 5.76±0.07 |
| DenseNet | SWA | 2.84±0.04 | 2.06±0.05 | 28.60±0.08 | 95.31±0.12 | 99.79±0.01 | 44.62±2.14 | 1.22±0.05 | 1.22±0.01 | 5.92±0.05 |
| | ours | 2.60±0.06 | 1.90±0.03 | 27.80±2.19 | 95.43±0.07 | 99.80±0.01 | 44.36±3.46 | 0.43±0.14 | 1.12±0.01 | 5.57±0.03 |
| | baseline [22] | 6.20±0.29 | 4.66±0.27 | 38.20±2.48 | 92.87±0.44 | 99.51±0.03 | 43.74±2.44 | 2.96±0.12 | 2.10±0.06 | 8.77±0.17 |
| | CRL [45] | 6.17±0.26 | 4.29±0.09 | 39.76±0.24 | 93.64±0.18 | 99.55±0.01 | 46.51±1.88 | 0.95±0.18 | 1.81±0.06 | 8.98±0.35 |
| ConvMixer | SAM | 4.56±0.17 | 3.06±0.14 | 33.69±1.68 | 94.55±0.20 | 99.68±0.01 | 47.32±1.03 | 1.37±0.13 | 1.53±0.05 | 7.45±0.20 |
| | SWA | 5.02±0.20 | 3.48±0.04 | 37.87±2.71 | 94.31±0.24 | 99.64±0.01 | 44.48±1.38 | 1.13±0.11 | 1.61±0.03 | 8.12±0.24 |
| | ours | 4.48±0.13 | 3.01±0.12 | 31.41±1.29 | 94.92±0.19 | 99.69±0.01 | 48.12±0.88 | 0.52±0.08 | 1.51±0.01 | 7.68±0.06 |
| | baseline [22] | 8.33±1.44 | 6.29±1.30 | 42.32±3.26 | 92.02±0.96 | 99.34±0.14 | 43.80±1.49 | 3.43±0.51 | 2.53±0.27 | 10.15±0.65 |
| ConvMixer | CRL [45] | 6.89±0.51 | 4.78±0.35 | 41.47±1.38 | 93.27±0.21 | 99.50±0.04 | 45.99±1.38 | 0.99±0.22 | 1.98±0.07 | 9.57±0.31 |
| | SAM | 5.52±0.22 | 3.88±0.16 | 36.16±1.35 | 93.92±0.32 | 99.59±0.02 | 46.66±2.60 | 2.16±0.21 | 1.84±0.06 | 8.55±0.29 |
| | SWA | 4.68±0.26 | 3.35±0.17 | 34.73±2.04 | 94.54±0.16 | 99.65±0.02 | 45.62±1.17 | 1.31±0.17 | 1.59±0.07 | 7.90±0.30 |
| | ours | 4.98±0.23 | 3.47±0.13 | 32.88±1.77 | 94.75±0.19 | 99.63±0.02 | 49.02±1.36 | 0.89±0.20 | 1.58±0.05 | 7.86±0.28 |

| CIFAR-100 | | | | | | | | | | |
|-----------|---------------|-------------------|-------------------|-------------------|-------------------|---------------------|-------------------|------------------|------------------|-------------------|
| Network | Method | AURC (↓) | E-AURC (↓) | FPR-95% TPR(↓) | AUROC (↑) | AUPR- Success(↑) | AUPR- Error(↑) | ECE (↓) | NLL (↓) | Brier (↓) |
| ResNet110 | baseline [22] | 93.90±2.37 | 50.88±2.03 | 66.02±1.53 | 85.00±0.35 | 93.42±0.28 | 66.34±0.29 | 15.71±0.15 | 13.54±0.12 | 42.88±0.27 |
| | CRL [45] | 81.02±2.06 | 41.82±0.94 | 64.37±0.43 | 86.32±0.24 | 94.68±0.13 | 66.26±0.32 | 10.96±0.53 | 11.02±0.09 | 38.82±0.57 |
| | SAM | 78.35±1.32 | 41.01±0.61 | 63.86±0.52 | 86.38±0.14 | 94.82±0.09 | 66.21±0.72 | 10.58±0.15 | 10.54±0.18 | 37.89±0.44 |
| | SWA | 71.94±0.60 | 37.23±0.67 | 63.98±0.68 | 86.86±0.21 | 95.35±0.08 | 65.59±0.88 | 5.38±0.13 | 8.81±0.02 | 34.95±0.08 |
| WRNet | ours | 69.51±0.56 | 35.51±0.08 | 62.57±0.81 | 87.24±0.12 | 95.57±0.02 | 66.52±1.04 | 3.41±0.35 | 8.51±0.07 | 34.10±0.19 |
| | baseline [22] | 51.97±1.74 | 30.18±1.09 | 59.19±0.20 | 87.75±0.30 | 96.38±0.14 | 62.99±0.98 | 7.19±0.36 | 8.43±0.11 | 29.67±0.45 |
| | CRL [45] | 46.17±0.20 | 25.28±0.34 | 58.32±0.48 | 88.77±0.35 | 97.00±0.03 | 63.34±0.98 | 3.93±0.06 | 7.33±0.06 | 27.98±0.17 |
| | SAM | 43.21±0.22 | 24.51±0.65 | 57.56±0.91 | 88.82±0.26 | 97.11±0.07 | 62.69±0.72 | 4.97±0.38 | 7.15±0.03 | 26.99±0.16 |
| DenseNet | SWA | 41.62±0.18 | 22.73±0.02 | 57.19±0.87 | 89.23±0.03 | 97.33±0.01 | 62.17±0.27 | 7.61±0.14 | 7.29±0.05 | 27.64±0.06 |
| | ours | 40.80±0.31 | 22.00±0.15 | 56.13±0.44 | 89.53±0.10 | 97.41±0.01 | 63.11±0.12 | 6.07±0.34 | 6.83±0.01 | 26.78±0.20 |
| | baseline [22] | 67.41±0.67 | 35.82±0.54 | 61.55±2.01 | 86.46±0.31 | 95.55±0.05 | 65.60±1.26 | 9.04±0.31 | 9.60±0.08 | 35.11±0.39 |
| | CRL [45] | 64.30±1.26 | 34.59±0.33 | 61.42±1.74 | 87.19±0.25 | 95.74±0.06 | 64.67±1.45 | 5.38±0.34 | 8.50±0.14 | 32.85±0.32 |
| ConvMixer | SAM | 63.52±2.41 | 34.51±1.07 | 61.46±0.32 | 87.01±0.13 | 95.76±0.15 | 64.41±1.04 | 5.91±0.39 | 8.53±0.18 | 32.75±0.71 |
| | SWA | 59.88±1.40 | 32.08±0.25 | 63.11±0.73 | 87.34±0.18 | 96.09±0.05 | 62.88±1.39 | 4.89±0.28 | 7.82±0.10 | 31.63±0.47 |
| | ours | 56.62±0.18 | 30.33±0.42 | 61.34±1.59 | 87.75±0.20 | 96.32±0.05 | 63.45±0.70 | 3.16±0.08 | 7.53±0.04 | 30.55±0.10 |
| | baseline [22] | 76.96±1.64 | 41.20±1.64 | 63.57±0.52 | 86.28±0.18 | 94.81±0.06 | 65.39±0.83 | 7.42±1.20 | 9.98±0.33 | 36.27±0.83 |
| ConvMixer | CRL [45] | 64.20±0.50 | 33.57±0.56 | 60.11±2.62 | 87.59±0.31 | 95.85±0.06 | 65.78±1.73 | 3.99±0.43 | 8.55±0.02 | 32.84±0.06 |
| | SAM | 64.47±2.10 | 34.05±1.19 | 64.58±1.44 | 87.11±0.24 | 95.82±0.16 | 63.52±0.66 | 7.96±0.10 | 8.34±0.10 | 33.36±0.38 |
| | SWA | 69.73±2.51 | 37.47±1.14 | 63.43±1.24 | 86.56±0.15 | 95.35±0.17 | 64.57±0.64 | 6.29±0.35 | 9.13±0.24 | 34.54±0.67 |
| | ours | 57.15±1.92 | 29.96±0.59 | 60.49±2.22 | 88.04±0.35 | 96.35±0.09 | 64.57±2.05 | 6.40±0.62 | 8.05±0.18 | 31.52±0.63 |

4.2 Experiments Results and Analysis

Experimental setup. We conduct experiments on CIFAR-10, CIFAR-100, and Tiny-ImageNet with various network architectures. For comparison methods, we mainly compare our method with baseline [22] and CRL [45], which is the state-of-the-art approach for failure prediction that outperforms representative bayesian methods [13, 32]. Due to the limitation of space, we provide implementation details and experimental results on Tiny-ImageNet in supplementary material.

Flat minima do improve failure prediction. Comparative results are summarized in Table 2. We observe that flat minima based methods: SAM, SWA, and FMFP (ours) consistently outperform the strong baseline and CRL on various metrics of failure prediction. Particularly, FMFP generally yields the best results. For example, in the case of ResNet110, our method has 3.00% and 2.24% higher

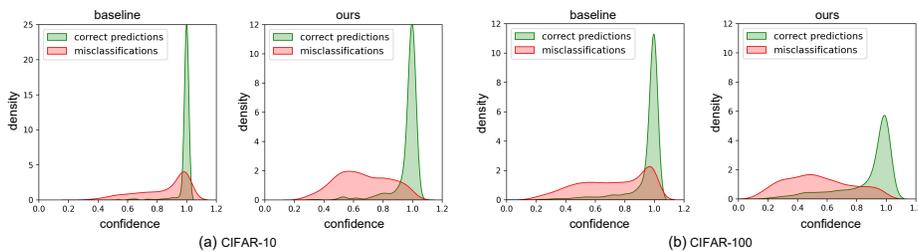


Fig. 9. Confidence distribution on correct and misclassified samples. Our method leads to a better separation for failure prediction.

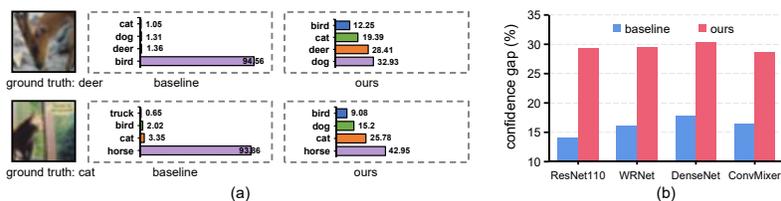


Fig. 10. (a) Predictive distribution on misclassified samples. (b) Our method (FMFP) significantly enlarges the average confidence gap between correct and incorrect samples.

values of AUROC on CIFAR-10 and CIFAR-100, respectively. In addition, flat minima based methods, especially the proposed FMFP, can achieve effective gains over confidence calibration (the last three columns in Table 2).

In Fig. 9, we observe that correct predictions and erroneous predictions overlap severely, making it difficult to distinguish them. Our method remarkably shifts the errors’ confidence distributions to smaller values and maintains the confidence of correct samples, leading to a better separation for failure prediction. Fig. 10(a) presents some examples of misclassified samples and their corresponding confidence distribution. Ours outputs much lower confidence on the erroneously predicted class. In Fig. 10(b), our method significantly enlarges the confidence gap between correct and incorrect samples. Besides, the risk-coverage curves in Fig. 11 and Fig. 2(b) also demonstrate the confidence reliability of our method.

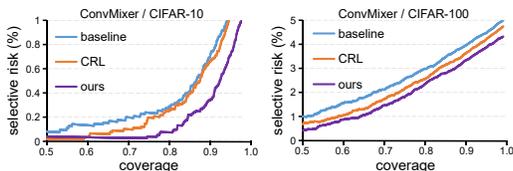


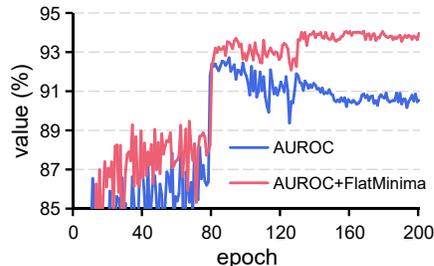
Fig. 11. Comparison of risk-coverage curves.

Failure prediction under distribution shift. In real-world applications, the model may encoder inputs subject to various kinds of distributional shifts. Thus, it becomes necessary to evaluate the confidence estimation performance to distributional shifts. The model is trained on CIFAR-10 and evaluated on corrupted dataset CIFAR-10-C [21]. The average results for 15 kinds of corruption under 5 different levels of perturbation severity are reported in Table 3. Our method consistently performs better than baseline and CRL methods.

Table 3. Failure prediction performance under distributional shifts.

| Method | AUROC \uparrow | | | | AURC \downarrow | | | | FPR-95%TPR \downarrow | | | |
|---------------|------------------|--------------|--------------|--------------|-------------------|--------------|---------------|---------------|-------------------------|--------------|--------------|--------------|
| | ResNet110 | WRNet | DenseNet | ConvMixer | ResNet110 | WRNet | DenseNet | ConvMixer | ResNet110 | WRNet | DenseNet | ConvMixer |
| baseline [22] | 79.45 | 83.81 | 81.97 | 81.28 | 157.46 | 112.46 | 148.91 | 168.39 | 71.29 | 64.05 | 69.26 | 71.07 |
| CRL [45] | 82.54 | 85.91 | 83.51 | 82.46 | 133.73 | 104.95 | 133.58 | 163.68 | 68.86 | 63.35 | 67.73 | 69.34 |
| ours | 84.72 | 87.34 | 84.90 | 84.93 | 119.79 | 94.34 | 130.42 | 145.21 | 65.51 | 58.87 | 64.95 | 65.29 |

Flat minima mitigates the reliable overfitting. Fig. 12 plots the AUROC curves during training. Although the failure prediction performance can be improved by early stopping, the classification accuracy of early checkpoint is much lower (Fig. 8). We can clearly observe that with flat minima, reliable overfitting has been diminished significantly, and the AUROC curves robustly improve until the end. Flat minima further leads to better classification accuracy, avoiding the trade-off between AUROC and accuracy when applying early stopping.

**Fig. 12.** Flat minima effectively mitigates the reliable overfitting.

Further understanding flat minima for failure prediction. ① *Representation learning.* Misclassification with high confidence is often attributed to spurious correlations appearing in the sample and the wrong class. Flat minima has been theoretically proved to result in invariant and disentangled representations [1], which is effective for spurious representations mitigation [4]. Therefore, with fewer spurious or irrelevant representations, the misclassified sample would be near the decision boundary with low confidence and less activated for wrong classes. ② *Uncertainty.* It has been shown that flat minima corresponds to regions in parameter space with rich posterior uncertainty [49]. Therefore, flat minima has the advantage to indicate the uncertainty of an input.

5 Concluding Remarks

Failure prediction is an important yet far less explored problem for safety-critical applications. This paper evaluates the effect of popular calibration methods for failure prediction. To our surprise, they have no or negative effect on failure prediction. We further find the under-confident issue of correctly classified samples, which leads to worse separation between the confidence of correct and incorrect samples. Finally, we propose to enlarge the confidence gap by finding flat minima, which yields strong performance on extensive experiments.

Acknowledgement. This work has been supported by the National Key Research and Development Program under Grant No. 2018AAA0100400, the National Natural Science Foundation of China grants U20A20223, 62076236, 61721004, the Key Research Program of Frontier Sciences of CAS under Grant ZDBS-LY-7004, and the Youth Innovation Promotion Association of CAS under Grant 2019141.

References

1. Achille, A., Soatto, S.: Emergence of invariance and disentanglement in deep representations. *J. Mach. Learn. Res.* **19**, 50:1–50:34 (2018) [14](#)
2. Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L.D., Monfort, M., Muller, U., Zhang, J., et al.: End to end learning for self-driving cars. arXiv preprint arXiv:1604.07316 (2016) [1](#)
3. Brier, G.W., et al.: Verification of forecasts expressed in terms of probability. *Monthly weather review* **78**(1), 1–3 (1950) [4](#), [6](#)
4. Cha, J., Chun, S., Lee, K., Cho, H.C., Park, S., Lee, Y., Park, S.: Swad: Domain generalization by seeking flat minima. In: *NeurIPS* (2021) [14](#)
5. Chaudhari, P., Choromanska, A., Soatto, S., LeCun, Y., Baldassi, C., Borgs, C., Chayes, J., Sagun, L., Zecchina, R.: Entropy-sgd: Biasing gradient descent into wide valleys. *Journal of Statistical Mechanics: Theory and Experiment* **2019**(12), 124018 (2019) [11](#)
6. Chen, T., Zhang, Z., Liu, S., Chang, S., Wang, Z.: Robust overfitting may be mitigated by properly learned smoothening. In: *ICLR* (2021) [11](#)
7. Corbière, C., Thome, N., Bar-Hen, A., Cord, M., Pérez, P.: Addressing failure prediction by learning model confidence. In: *NeurIPS*. pp. 2898–2909 (2019) [2](#), [6](#)
8. Corbière, C., Thome, N., Saporta, A., Vu, T.H., Cord, M., Perez, P.: Confidence estimation via auxiliary models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021) [6](#)
9. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *CVPR*. pp. 248–255 (2009) [6](#), [8](#)
10. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *ICLR* (2020) [5](#)
11. Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S.: Dermatologist-level classification of skin cancer with deep neural networks. *nature* **542**(7639), 115–118 (2017) [1](#)
12. Foret, P., Kleiner, A., Mobahi, H., Neyshabur, B.: Sharpness-aware minimization for efficiently improving generalization. In: *ICLR* (2020) [3](#), [4](#), [10](#), [11](#)
13. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: *ICML*. vol. 48, pp. 1050–1059 (2016) [12](#)
14. Geifman, Y., El-Yaniv, R.: Selective classification for deep neural networks. In: *NeurIPS*. pp. 4878–4887 (2017) [2](#), [5](#), [6](#), [8](#)
15. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: *ICML*. pp. 1321–1330 (2017) [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [9](#)
16. Gupta, K., Rahimi, A., Ajanthan, T., Mensink, T., Sminchisescu, C., Hartley, R.: Calibration of neural networks using splines. In: *ICLR* (2020) [2](#), [5](#), [6](#)
17. Havasi, M., Jenatton, R., Fort, S., Liu, J.Z., Snoek, J., Lakshminarayanan, B., Dai, A.M., Tran, D.: Training independent subnetworks for robust prediction. In: *ICLR* (2020) [2](#)
18. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR*. pp. 770–778 (2016) [6](#), [8](#)
19. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: *ECCV*. pp. 630–645 (2016) [3](#), [6](#)
20. Hebbalaguppe, R., Prakash, J., Madan, N., Arora, C.: A stitch in time saves nine: A train-time regularizing loss for improved neural network calibration. In: *CVPR*. pp. 16081–16090 (June 2022) [2](#)

21. Hendrycks, D., Dietterich, T.G.: Benchmarking neural network robustness to common corruptions and perturbations. In: ICLR (2019) [13](#)
22. Hendrycks, D., Gimpel, K.: A baseline for detecting misclassified and out-of-distribution examples in neural networks. In: ICLR (2017) [2, 5, 6, 7, 10, 12, 14](#)
23. Hendrycks, D., Mazeika, M., Dietterich, T.G.: Deep anomaly detection with outlier exposure. In: ICLR (2019) [2](#)
24. Hendrycks, D., Mu, N., Cubuk, E.D., Zoph, B., Gilmer, J., Lakshminarayanan, B.: Augmix: A simple data processing method to improve robustness and uncertainty. In: ICLR (2020) [3](#)
25. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017) [6](#)
26. Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: CVPR. pp. 2261–2269 (2017) [6](#)
27. Huang, W.R., Emam, Z.A.S., Goldblum, M., Fowl, L.H., Terry, J., Huang, F., Goldstein, T.: Understanding generalization through visualizations. In: "I Can't Believe It's Not Better!" NeurIPS 2020 workshop (2020) [3, 4, 10](#)
28. Izmailov, P., Wilson, A., Podoprikin, D., Vetrov, D., Garipov, T.: Averaging weights leads to wider optima and better generalization. In: UAI. pp. 876–885 (2018) [3, 4, 10, 11](#)
29. Janai, J., Güney, F., Behl, A., Geiger, A., et al.: Computer vision for autonomous vehicles: Problems, datasets and state of the art. Foundations and Trends® in Computer Graphics and Vision **12**(1–3), 1–308 (2020) [1](#)
30. Jiang, H., Kim, B., Gupta, M.R.: To trust or not to trust a classifier. In: NeurIPS (2018) [5](#)
31. Joo, T., Chung, U.: Revisiting explicit regularization in neural networks for well-calibrated predictive uncertainty. arXiv preprint arXiv:2006.06399 (2020) [3, 4, 7, 10](#)
32. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? In: NeurIPS. pp. 5574–5584 (2017) [12](#)
33. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. Tech. rep., Citeseer (2009) [3, 6](#)
34. Kull, M., Perelló-Nieto, M., Kängsepp, M., de Menezes e Silva Filho, T., Song, H., Flach, P.A.: Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with dirichlet calibration. In: NeurIPS. pp. 12295–12305 (2019) [5](#)
35. Kull, M., de Menezes e Silva Filho, T., Flach, P.A.: Beta calibration: a well-founded and easily implemented improvement on logistic calibration for binary classifiers. In: AISTATS. pp. 623–631 (2017) [2, 5](#)
36. Kumar, A., Liang, P.S., Ma, T.: Verified uncertainty calibration. NeurIPS (2019) [4](#)
37. Lee, K., Lee, K., Lee, H., Shin, J.: A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In: NeurIPS. pp. 7167–7177 (2018) [2](#)
38. Leidner, D., Borst, C., Dietrich, A., Beetz, M., Albu-Schäffer, A.: Classifying compliant manipulation tasks for automated planning in robotics. In: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1769–1776 (2015) [1](#)
39. Liang, S., Li, Y., Srikant, R.: Enhancing the reliability of out-of-distribution image detection in neural networks. In: ICLR (2018) [2](#)
40. Lin, T., Goyal, P., Girshick, R.B., He, K., Dollár, P.: Focal loss for dense object detection. IEEE Trans. Pattern Anal. Mach. Intell. pp. 318–327 (2020) [3, 4, 7, 8, 10](#)

41. Liu, B., Ben Ayed, I., Galdran, A., Dolz, J.: The devil is in the margin: Margin-based label smoothing for network calibration. In: CVPR. pp. 80–88 (June 2022) [2](#)
42. Luo, Y., Wong, Y., Kankanhalli, M.S., Zhao, Q.: Learning to predict trustworthiness with steep slope loss. NeurIPS (2021) [6](#)
43. Minderer, M., Djolonga, J., Romijnders, R., Hubis, F., Zhai, X., Houlsby, N., Tran, D., Lucic, M.: Revisiting the calibration of modern neural networks. NeurIPS (2021) [2](#), [5](#), [6](#)
44. Miotto, R., Li, L., Kidd, B.A., Dudley, J.T.: Deep patient: an unsupervised representation to predict the future of patients from the electronic health records. Scientific reports p. 26094 (2016) [1](#)
45. Moon, J., Kim, J., Shin, Y., Hwang, S.: Confidence-aware learning for deep neural networks. In: ICML. pp. 7034–7044 (2020) [3](#), [5](#), [6](#), [8](#), [12](#), [14](#)
46. Mozafari, A.S., Gomes, H.S., Leão, W., Gagné, C.: Unsupervised temperature scaling: An unsupervised post-processing calibration method of deep networks. arXiv: Computer Vision and Pattern Recognition (2019) [5](#)
47. Mukhoti, J., Kulharia, V., Sanyal, A., Golodetz, S., Torr, P.H.S., Dokania, P.K.: Calibrating deep neural networks using focal loss. In: NeurIPS (2020) [2](#), [4](#), [7](#)
48. Müller, R., Kornblith, S., Hinton, G.: When does label smoothing help? In: NeurIPS. pp. 4696–4705 (2019) [2](#), [3](#), [4](#), [6](#), [7](#), [10](#)
49. Murphy, K.P.: Probabilistic Machine Learning: An introduction. MIT Press (2022), [probl.ai](#) [14](#)
50. Naeini, M.P., Cooper, G.F., Hauskrecht, M.: Obtaining well calibrated probabilities using bayesian binning. In: AAAI. pp. 2901–2907 (2015) [4](#)
51. Nixon, J., Dusenberry, M.W., Zhang, L., Jerfel, G., Tran, D.: Measuring calibration in deep learning. In: CVPR Workshops. vol. 2 (2019) [6](#)
52. Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D., Nowozin, S., Dillon, J., Lakshminarayanan, B., Snoek, J.: Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. NeurIPS (2019) [5](#), [6](#)
53. Patel, K., Beluch, W.H., Yang, B., Pfeiffer, M., Zhang, D.: Multi-class uncertainty calibration via mutual information maximization-based binning. In: ICLR (2020) [2](#), [5](#)
54. Pereyra, G., Tucker, G., Chorowski, J., Kaiser, L., Hinton, G.: Regularizing neural networks by penalizing confident output distributions. arXiv preprint arXiv:1701.06548 (2017) [2](#), [3](#), [4](#)
55. Pittorino, F., Lucibello, C., Feinauer, C., Perugini, G., Baldassi, C., Demyanenko, E., Zecchina, R.: Entropic gradient descent algorithms and wide flat minima. Journal of Statistical Mechanics: Theory and Experiment (12), 124015 (2021) [11](#)
56. Rahimi, A., Shaban, A., Cheng, C., Hartley, R., Boots, B.: Intra order-preserving functions for calibration of multi-class neural networks. In: NeurIPS (2020) [2](#), [5](#)
57. Rice, L., Wong, E., Kolter, Z.: Overfitting in adversarially robust deep learning. In: ICML. pp. 8093–8104 (2020) [11](#)
58. Shehzad, M.N., Bashir, Q., Farooq, U., Ahmed, G., Raza, M., Kumar, P.M., Khalid, M.: Threshold temperature scaling: Heuristic to address temperature and power issues in mpsoes. Microprocess. Microsystems p. 103124 (2020) [2](#), [5](#)
59. Shen, Z., Liu, Z., Xu, D., Chen, Z., Cheng, K.T., Savvides, M.: Is label smoothing truly incompatible with knowledge distillation: An empirical study. In: ICLR (2020) [10](#)
60. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: ICML. pp. 6105–6114 (2019) [6](#)

61. Thulasidasan, S., Chennupati, G., Bilmes, J., Bhattacharya, T., Michalak, S.: On mixup training: Improved calibration and predictive uncertainty for deep neural networks. In: NeurIPS. pp. 13888–13899 (2019) [2](#), [3](#), [4](#), [6](#), [7](#), [8](#), [10](#)
62. Tolstikhin, I.O., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., et al.: Mlp-mixer: An all-mlp architecture for vision. NeurIPS (2021) [5](#)
63. Trockman, A., Kolter, J.Z.: Patches are all you need? arXiv preprint arXiv:2201.09792 (2022) [6](#)
64. Vaicenavicius, J., Widmann, D., Andersson, C., Lindsten, F., Roll, J., Schön, T.: Evaluating model calibration in classification. In: AISTATS. pp. 3459–3467 (2019) [6](#)
65. Wang, D., Feng, L., Zhang, M.: Rethinking calibration of deep neural networks: Do not be afraid of overconfidence. In: NeurIPS (2021) [5](#)
66. Wen, Y., Jerfel, G., Muller, R., Dusenberry, M.W., Snoek, J., Lakshminarayanan, B., Tran, D.: Combining ensembles and data augmentation can harm your calibration. In: ICLR (2020) [2](#), [4](#), [6](#)
67. Wu, D., Xia, S., Wang, Y.: Adversarial weight perturbation helps robust generalization. In: NeurIPS (2020) [11](#)
68. Xing, C., Arik, S.Ö., Zhang, Z., Pfister, T.: Distance-based learning from errors for confidence calibration. In: ICLR (2020) [2](#), [3](#), [4](#)
69. Yao, L., Miller, J.: Tiny imagenet classification with convolutional neural networks. CS 231N [6](#)
70. Yun, S., Park, J., Lee, K., Shin, J.: Regularizing class-wise predictions via self-knowledge distillation. In: CVPR. pp. 13873–13882 (2020) [2](#), [3](#), [4](#), [6](#), [7](#), [10](#)
71. Zagoruyko, S., Komodakis, N.: Wide residual networks. In: BMVC (2016) [6](#)
72. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: Mixup: Beyond empirical risk minimization. In: ICLR (2018) [4](#)
73. Zhang, L., Deng, Z., Kawaguchi, K., Zou, J.: When and how mixup improves calibration. In: ICML. pp. 26135–26160 (2022) [4](#)
74. Zhang, W., Vaidya, I.: Mixup training leads to reduced overfitting and improved calibration for the transformer architecture. CoRR (2021) [4](#)
75. Zhong, Z., Cui, J., Liu, S., Jia, J.: Improving calibration for long-tailed recognition. In: CVPR. pp. 16489–16498 (2021) [2](#), [4](#)