



# Uncertainty Calibration with Energy Based Instance-wise Scaling in the Wild Dataset

Mijoo Kim<sup>1</sup>  and Junseok Kwon<sup>1</sup> 

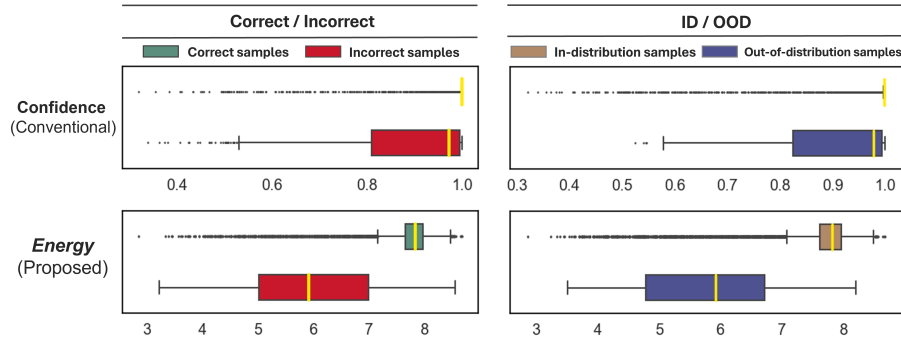
Chung-Ang University, Seoul, Korea  
{mijoo707, jskwon}@cau.ac.kr

**Abstract.** With the rapid advancement in the performance of deep neural networks (DNNs), there has been significant interest in deploying and incorporating artificial intelligence (AI) systems into real-world scenarios. However, many DNNs lack the ability to represent uncertainty, often exhibiting excessive confidence even when making incorrect predictions. To ensure the reliability of AI systems, particularly in safety-critical cases, DNNs should transparently reflect the uncertainty in their predictions. In this paper, we investigate robust post-hoc uncertainty calibration methods for DNNs within the context of multi-class classification tasks. While previous studies have made notable progress, they still face challenges in achieving robust calibration, particularly in scenarios involving out-of-distribution (OOD). We identify that previous methods lack adaptability to individual input data and struggle to accurately estimate uncertainty when processing inputs drawn from the wild dataset. To address this issue, we introduce a novel instance-wise calibration method based on an energy model. Our method incorporates energy scores instead of softmax confidence scores, allowing for adaptive consideration of DNN uncertainty for each prediction within a logit space. In experiments, we show that the proposed method consistently maintains robust performance across the spectrum, spanning from in-distribution to OOD scenarios, when compared to other state-of-the-art methods. The source code is available at <https://github.com/mijoo308/Energy-Calibration>.

**Keywords:** Uncertainty Calibration · Out-of-distribution · Energy based instance-wise scaling

## 1 Introduction

Despite the impressive performance demonstrated by recent AI systems, their deployment should be carefully considered, particularly in safety critical situations (*e.g.* autonomous driving, finance, health care, and medical diagnosis), because these systems cannot consistently ensure accurate predictions. For example, in the field of medical diagnosis, incorrect predictions have the potential to result in catastrophic outcomes. To address this concern, the system must transparently reveal the uncertainty associated with its prediction. Recently, DNNs rely on *confidence* as a way of expressing uncertainty, but they tend to assign higher



**Fig. 1: Conventional softmax confidence scores versus the proposed energy scores.** First column: Softmax confidence score (top) and negative energy score (bottom) for correct and incorrect samples. Second column: Softmax confidence score (top) and negative energy score (bottom) for in-distribution (CIFAR10) and OOD samples (SVHN). Our energy scores exhibit greater separability between correct and incorrect predictions, as well as between in-distribution and OOD samples in DenseNet201.

confidence scores than their actual accuracy [4]. This discrepancy stems from the inherent inability of DNNs to express appropriate uncertainty during inference.

To solve this problem, an increasing body of research has focused on refining confidence representations to reflect the uncertainty, a field known as uncertainty calibration. This effort aims to adjust confidence scores to align more closely with accuracy, ultimately improving the reliability of predictions. As a result of active researches in this field [4, 5, 19, 20, 25, 36], the discrepancy between confidence and accuracy has been significantly mitigated. However, these methods mainly deal with samples drawn from the same distribution on which DNNs were trained (*i.e.* in-distribution), often overlooking distribution shift scenarios. Consequently, they face difficulties in achieving calibration effects when confronted with distribution shift scenarios. When considering real-world deployment, calibration methods should demonstrate robustness in handling samples from unknown distributions. In this context, this issue has been considered by Tomani et al. [29] within the domain of post-hoc calibration. This method exhibits relatively effective performance in out-of-distribution (OOD) scenarios. However, conversely, it suffered from miscalibration in-distribution (ID) scenarios, exhibiting even greater miscalibration compared to the pre-calibration state.

In this paper, we propose a novel energy-based calibration method that exhibits robustness across the spectrum, from ID to OOD scenarios, including various distribution shift scenarios. We address uncertainty calibration in the context of multi-class classification, particularly in a post-hoc manner. The proposed method utilizes the *energy score* to adeptly capture the uncertainty in DNNs. Previous studies [17] have shown that the energy model produces scores that are discriminative between ID and OOD samples. Beyond this work, we demonstrate that the energy function can induce distinctive scores not only between ID and OOD samples but also between correct and incorrect samples, which can be effectively utilized for uncertainty calibration. Before delving into the mathematical

derivation for this in Section 4, we provide an intuitive overview of the proposed energy score in Fig. 1. The energy score exhibits a superior ability to produce distinctive scores between ID and OOD samples, as well as between correct and incorrect samples. This implies that the energy score can more accurately represent the uncertainty of DNNs compared to the confidence score. Motivated by this observation, we utilize the energy score as an uncertainty estimator for each prediction, which inspires us to propose an instance-wise robust calibration method that adjusts the calibration factor accordingly. In Section 5, we demonstrate that the proposed method achieves remarkable robust calibration performance across various baseline DNN models in the wild datasets. This includes scenarios involving OOD samples with semantic shift and varying degrees of covariate shift, as well as ID scenarios. To sum up, our main contributions are as follows:

- We demonstrate the effectiveness of utilizing energy scores for uncertainty calibration through both mathematical derivation and empirical validation.
- We introduce a novel post-hoc calibration method that utilizes the energy score to adaptively capture the uncertainty of predictions in DNNs for each individual input.
- We illustrate that the proposed calibration method shows robustness in the wild datasets across a wide range of distribution shifts such as covariate and semantic shifts, as well as in complete ID.

## 2 Related work

### 2.1 Post-hoc Calibration

Confidence calibration can be divided into two categories. The first category is known as *training-time calibration* [16, 19, 20, 27, 37], such as focal loss [19] and label smoothing [20]. These methods train DNNs to exhibit calibrated behavior during training. The second category is referred to as *post-hoc calibration* [4, 7, 25, 31, 36]. In this approach, pre-trained neural networks are utilized along with hold-out validation datasets to learn calibration mappings in a post-hoc manner.

Post-hoc calibration can be further categorized into non-parametric and parametric approaches. Non-parametric approaches include Histogram Binning (HB) [33] and Isotonic Regression (IR) [34]. HB divided predicted probabilities into multiple intervals, each associated with representative confidences. IR utilized isotonic regression with uncalibrated confidences as the x-axis and the expected accuracy values as the y-axis. BBQ [21], non-parametric extension of HB, incorporated Bayesian model averaging to enhance calibration. On the other hand, the most common parametric calibration method is Temperature Scaling (TS) [4]. As the temperature increases, the distribution of logits becomes more uniformly distributed, resulting in a decrease in the confidence score associated with the predicted label. TS has a significant advantage in terms of accuracy preservation, as it maintains the originally predicted label with the highest confidence score unchanged. However, TS exhibits limited expressiveness as it relies on only a single parameter that is fixed on the validation set.

To address this issue, ensemble approaches, which combine both non-parametric and parametric methods, have been proposed. Ensemble TS [36] introduced additional parameters to enhance expressiveness, building upon TS. IRM [36] leveraged the accuracy-preserving property of the parametric approach and the expressiveness of the non-parametric approach, representing a multi-class extension of IR. Parameterized TS [28] employed a similar strategy to address the expressiveness limitation of TS. In other approaches, Beta calibration [14] was extended to Dirichlet calibration [13] and Spline calibration [5] utilized spline fitting to approximate the empirical cumulative distribution.

We employ a post-hoc calibration approach, however, our method differs from this body of work. While these methods only addressed scenarios where test samples are drawn from the same distribution on which DNNs were trained, our approach can handle more diverse situations. It ensures effective calibration for ID samples while maintaining the original classifier accuracy.

## 2.2 Beyond In-distribution Calibration

Conventional research has mainly focused on investigating post-hoc uncertainty calibration methods. However, these methods often overlook scenarios involving distribution shifts. Due to their dependence on a fixed calibration map optimized for ID validation sets, they struggle to effectively handle unknown test samples.

Recently, the importance of the robustness of post-hoc calibration method across various distribution shift scenarios has been emphasized [29, 30]. In [29], various degrees of Gaussian perturbations were injected into the ID validation dataset. The parameters of the calibration method were then adjusted using the perturbed validation data, resulting in enhanced robustness against shifted distributions. However, this method tends to achieve notable performance only within scenarios where a certain degree of distribution shift is present. Moreover, in a complete ID scenario, it exhibits even worse calibration compared to the pre-calibration state. To solve this problem, DAC [30] has been proposed as a pre-processing step before employing the existing post-hoc calibration methods. Notably, it leveraged additional output information for uncertainty estimation and enhanced the calibration performance in distribution shift scenarios.

Similar to these methods, we focus on a broad range of scenarios, ranging from ID to various OOD scenarios. However, unlike these methods, our approach can facilitate robust calibration without requiring any additional DNN’s layer information except for the last layer. In experiments, we demonstrate that our proposed method performs comparably to, and in some cases even surpasses, other state-of-the-art methods that utilize DAC as a preprocessing step.

## 3 Problem Setup

In this section, we establish preliminaries for uncertainty calibration. We define key notations in the context of multi-class classification and present a representative calibration metric derived from the concept of perfect calibration. Additionally, we address calibration in OOD scenarios.

### 3.1 Notation

Let  $\mathbf{x} \in \mathcal{X}$  and  $y \in \mathcal{Y} = \{1, \dots, K\}$  be random variables that denote  $d$ -dimensional inputs and labels, respectively, in multi-class classification tasks with  $K$  classes. These random variables follow a joint distribution  $\pi(\mathbf{x}, y) = \pi(y|\mathbf{x})\pi(\mathbf{x})$ . The dataset  $\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$  consists of the  $N$  number of i.i.d. samples drawn from  $\pi(\mathbf{x}, y)$ . Let  $f$  be a pre-trained neural network and  $f(\mathbf{x}) = (\hat{y}, \mathbf{z})$  be the output of the neural network, where  $\hat{y}$  is a predicted label and  $\mathbf{z}$  is an original non-probabilistic output of the network, referred to as the logit. The logit  $\mathbf{z}$  is converted into a probabilistic value  $\hat{p}$  using the softmax function  $\sigma_{SM}$ . Thus,  $\hat{p}$  represents a confidence score associated with the predicted label  $\hat{y}$ . To summarize, the output of the neural classifier  $f$ ,  $\hat{p}$  and  $\hat{y}$ , can be obtained as follows:

$$\hat{p} = \max_k \sigma_{SM}(\mathbf{z}), \text{ and } \hat{y} = \operatorname{argmax}_k \sigma_{SM}(\mathbf{z}) \quad \text{for } k \in \{1, \dots, K\}.$$

### 3.2 Calibration Metric

Perfect calibration is achieved when the predicted probability (*i.e.* confidence) matches with the actual likelihood of a correct prediction. If a neural network predicts a label as  $y$  with confidence  $p$ , the actual likelihood of the prediction should ideally be  $p$ . Thus, the perfect calibration in multi-class classification can be represented as follows:

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

The goal of uncertainty calibration is to minimize the gap between the ground-truth likelihood and the predicted confidence by calibrating the confidence value.

Using this definition of perfect calibration, the calibration error can be computed by modifying (1):

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

Subsequently, the expected calibration error (ECE) [21] empirically approximates the calibration error in (2) using a binning technique. By discretizing the confidence interval into  $M$  equally sized bins, *i.e.*  $\{B_m\}_{m=1}^M$ , the ECE calculates a weighted average of differences between accuracy  $\text{acc}(\cdot)$  and confidence  $\text{conf}(\cdot)$  in each bin. With  $N$  samples, the ECE is defined as follows:

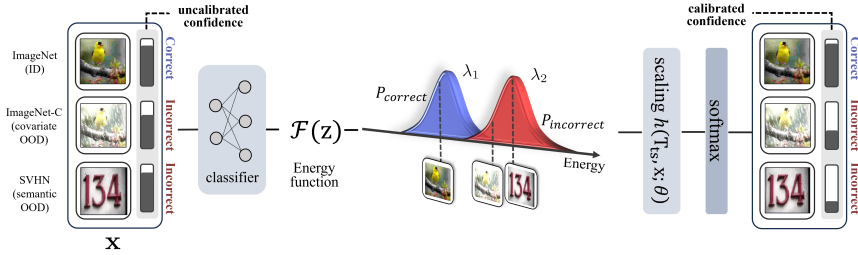
$$\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{N} |\text{acc}(B_m) - \text{conf}(B_m)|, \quad (3)$$

where all ECE values in this paper are calculated with  $M = 15$  and multiplied by 100. In addition to ECE, there are other metrics such as Maximum Calibration Error (MCE) [21], which represents the highest error among bins, Static Calibration Error (SCE) [23] that evaluates calibration errors on a classwise manner, and KDE-ECE [36] that utilizes Kernel Density Estimation (KDE). As ECE is the most representative metric, we primarily evaluate the proposed method using ECE, but we also employ other metrics.

### 3.3 Calibration in OOD Scenarios

In general, OOD refers to a distribution that differs from the training distribution [9, 32]. In this paper, the term OOD includes two types of distribution shifts: *covariate shift* and *semantic shift*. Covariate-shifted samples are drawn from a different joint distribution  $\pi_{ood}^{cov}(\mathbf{x}, y)$  such that  $\pi_{ood}^{cov} \neq \pi$ . In other words, while the samples may belong to the same class, they are presented in different forms [24]. In the case of semantic shift, the data is drawn from  $\pi_{ood}^{sem}(\mathbf{x}, \bar{y})$ , where  $\mathcal{Y} \cap \bar{\mathcal{Y}} = \emptyset$ , indicating that the data is from classes not present in the training set  $\mathcal{D}$  [9]. Therefore, in semantic shift scenarios, all predictions by a pre-trained classifier may be incorrect, as they may correspond to one of the  $K$  in-distribution classes in  $\mathcal{Y}$ . From the perspective of calibration, in such scenarios, the lower the confidence, the better the calibration.

## 4 Proposed Method



**Fig. 2: The overall pipeline for our calibration method.** An input image  $\mathbf{x}$  from the wild dataset is fed into a pre-trained classifier, producing the logit  $\mathbf{z}$ . Subsequently, the free energy  $\mathcal{F}$  defined in (6) is calculated for each logit  $\mathbf{z}$ . Then, two probability density functions (PDFs), *i.e.*  $P_{correct}$  and  $P_{incorrect}$ , are estimated based on the free energy  $\mathcal{F}$  of correct and incorrect instances, respectively. These PDFs are utilized to adjust the scaling factors  $\lambda_1$  and  $\lambda_2$  in (10). With the scaling factors, the parameters  $\theta_1$  and  $\theta_2$  are then trained through the optimization of the loss function in (12). Using the trained parameters, the calibrated confidence for a test image  $\mathbf{x}$  can be calculated by applying scaling in an instance-wise manner as in (13).

To overcome the limitations of conventional calibration methods when distribution shifts exist, we propose a robust calibration method that exhibits calibration improvements across various OOD scenarios. Previous approaches [29] that initially addressed shift scenarios in post-hoc calibration could not properly handle ID inputs. In contrast, our method achieves calibration improvements in both OOD and ID scenarios by adaptively capturing the uncertainty of pre-trained neural networks for each input. To accomplish this, our method leverages the concept of the energy model, which is technically derived from energy-based OOD detection methods [17]. Before introducing our calibration method, we lay out the mathematical motivation behind the proposed method by establishing a correlation between our method and the energy model. The overall pipeline for understanding the propose method is illustrated in Fig.2.

#### 4.1 Mathematical Motivation

The energy function makes scores on ID and OOD more distinguishable than the softmax function [17]. As demonstrated in [15], there is a connection between Gibbs distributions (*i.e.* Boltzmann distributions) and softmax functions:

$$P(y|\mathbf{x}) = \frac{e^{-\beta E(\mathbf{x}, y)}}{\int_{y \in \mathcal{Y}} e^{-\beta E(\mathbf{x}, y)}}, \text{ and } \sigma_{SM}(y|\mathbf{x}) = \frac{e^{f_y(\mathbf{x})}}{\sum_{i=1}^K e^{f_i(\mathbf{x})}}, \quad (4)$$

where  $P(y|\mathbf{x})$  denotes the Gibbs distribution with the energy  $E(\mathbf{x}, y) : \mathbb{R}^D \rightarrow \mathbb{R}$  and  $\sigma_{SM}(y|\mathbf{x})$  indicates the softmax function for the  $K$ -class classifier  $f(\mathbf{x}) : \mathbb{R}^D \rightarrow \mathbb{R}^K$ . In (4),  $f(\mathbf{x})$  outputs a vector of length  $K$  and  $f_i(\mathbf{x})$  denotes the  $i$ -th element of the vector.

By comparing  $P(y|\mathbf{x})$  and  $\sigma_{SM}(y|\mathbf{x})$  in (4), we can derive the energy as

$$E(\mathbf{x}, y) = -f_y(\mathbf{x}), \quad (5)$$

where the positive constant value  $\beta$  is set to 1. The denominator of  $P(y|\mathbf{x})$  in (4) is a partition function, transforming each energy value corresponding to  $y$  into a probability value within the range of  $[0, 1]$ . In particular, Helmholtz free energy is defined as a log partition function [15]. Then, the free energy  $\mathcal{F}$  can be represented using the connection between the Gibbs distribution and the softmax function:

$$\mathcal{F}(\mathbf{x}) = -\log \sum_{i=1}^K e^{f_i(\mathbf{x})}. \quad (6)$$

A mathematical relationship between the energy and the negative log likelihood (NLL) loss has been derived in [15]. Based on this, it is demonstrated that the NLL loss inherently decreases the energy for ID samples, while increasing the energy for OOD samples [17]. From these findings, we define the NLL loss,  $\mathcal{L}_{NLL} = -\log \sigma_{SM}(y|\mathbf{x})$ , as a combination of  $E(\mathbf{x}, y)$  in (5) and  $-\mathcal{F}(\mathbf{x})$  in (6).

$$\mathcal{L}_{NLL} = -\log \frac{e^{f_y(\mathbf{x})}}{\sum_{i=1}^K e^{f_i(\mathbf{x})}} = -f_y(\mathbf{x}) + \log \sum_{i=1}^K e^{f_i(\mathbf{x})} = E(\mathbf{x}, y) - \mathcal{F}(\mathbf{x}), \quad (7)$$

where the free energy  $\mathcal{F}$  can be interpreted as a contrastive term that aggregates the energies for all classes of  $i \in \{1, \dots, K\}$ . From the third equation in (7), we can see that the NLL loss inherently lowers the energy for the correct label  $y$  and raises the energy for the other labels. Additionally, the derivative of  $\mathcal{L}_{NLL}$  over the network parameter  $\theta$  is calculated as follows.

$$\begin{aligned} \frac{\partial \mathcal{L}_{NLL}}{\partial \theta} &= \frac{\partial E(\mathbf{x}, y)}{\partial \theta} - \frac{\partial \mathcal{F}(\mathbf{x})}{\partial \theta} = \frac{\partial E(\mathbf{x}, y)}{\partial \theta} - \sum_{i=1}^K \frac{\partial E(\mathbf{x}, i)}{\partial \theta} \frac{e^{-E(\mathbf{x}, i)}}{\sum_{j=1}^K e^{-E(\mathbf{x}, j)}} \\ &= \frac{\partial E(\mathbf{x}, y)}{\partial \theta} - \sum_{i=1}^K \frac{\partial E(\mathbf{x}, i)}{\partial \theta} P(i|\mathbf{x}) \\ &= \frac{\partial E(\mathbf{x}, y)}{\partial \theta} (1 - P(y|\mathbf{x})) - \sum_{i \neq y} \frac{\partial E(\mathbf{x}, i)}{\partial \theta} P(i|\mathbf{x}), \end{aligned} \quad (8)$$

where the second equality holds by using the Gibbs distribution in (4) ( $\beta = 1$ ). From the last equation in (8), we can see that the energy function is weighted by each probability, pushing down the energy for the correct label and pulling up the energy for the incorrect labels. Liu et al. [17] explained that increasing the energy, except for the correct labels, inherently boosts the energy of OOD samples. Following this, the free energy  $\mathcal{F}$  can serve as the *energy score*, which can make distinctive values between ID and OOD. This is because it is a smooth approximation of  $E$  with the dominance of the ground-truth label  $y$  over all other labels. They focused on the distinguishability of the energy score on ID and OOD samples in the OOD detection task, particularly in cases of semantic shift. To further expand this interpretation, we bring the concept into the perspective of uncertainty calibration.

In our context, we focus on not only the distinguishability between *ID* and *OOD* samples but also between *correct* and *incorrect* predictions. Interpreting the final equation in (8) in a simpler manner, it reveals that the energy score has the capability to differentiate between correct and incorrect samples. This implies that it can make distinctive scores between correct and incorrect predictions not only in ID but also in covariate shifts, as  $\mathcal{Y}$  remains consistent. Furthermore, for semantic shift cases, as demonstrated in [17], all predictions are considered incorrect, indicating distinguishability since there is no overlap in labels ( $\mathcal{Y} \cap \bar{\mathcal{Y}} = \emptyset$ ). Building upon the motivation outlined so far, we introduce the incorporation into post-hoc uncertainty calibration in the next section.

## 4.2 Robust Instance-wise Calibration

Most existing calibration methods are limited by the assumption of the same distribution on which the classifier has been trained. As the parameters are optimized using the consistent distribution of the validation set, these methods lack the adaptiveness to effectively address distribution shift scenarios. To solve this problem, an uncertainty calibration method should possess the capability to capture the uncertainty of the neural network for each individual sample. In this context, the energy score can effectively fulfill this role within the framework of post-hoc calibration. As shown in the motivation and Fig.1, it is evident that the energy score is capable of producing distinct values for both cases: ID and OOD samples, as well as correct and incorrect samples.

By facilitating our motivation, we adjust the scaling factor for each input samples to achieve uncertainty calibration. Our method is fundamentally built upon the temperature scaling (TS) technique introduced in [4] to incorporate the advantages of accuracy-preserving property. The proposed scaling factor is defined as follows:

$$h(T_{ts}, \mathbf{x}; \theta) = \underbrace{T_{ts}}_{\text{Fixed on validation set}} \underbrace{-\lambda_1 \theta_1 + \lambda_2 \theta_2}_{\text{Adaptive for each input}}, \quad (9)$$

where  $T_{ts}$  denotes the temperature parameter obtained by the TS technique [4], which is fixed on the validation set, and  $\theta = \{\theta_1, \theta_2\}$  denotes trainable

**Algorithm 1** Training Calibration Parameters

---

**Require:** In-distribution validation data:  $\mathcal{D}_{in}$ , semantic OOD data  $\mathcal{D}_{out}$ , Pre-trained  $K$ -class classifier:  $f(\mathbf{x})$ , Temperature value from temperature scaling:  $T_{ts}$

$\mathcal{D} = \mathcal{D}_{in} + \mathcal{D}_{out}$   
 $(\mathbf{x}, y) \sim \mathcal{D}$   
 $\mathbf{z} = f(\mathbf{x})$   
 $\mathcal{F}(\mathbf{z}) \leftarrow -\log \text{sumexp}(\mathbf{z})$

Estimate Gaussian probability density functions  
 $P_1 \leftarrow \text{fit}(\mathcal{F}(\mathbf{z}_{correct}))$   
 $P_2 \leftarrow \text{fit}(\mathcal{F}(\mathbf{z}_{incorrect}))$

**for**  $i \in [1, N]$  **do**  
  Find  $\theta$  that minimizes the MSE loss  
   $T = T_{ts} - P_1(\mathcal{F}(\mathbf{z}_i))\theta_1 + P_2(\mathcal{F}(\mathbf{z}_i))\theta_2$   
   $\mathcal{L} = \text{MSE}(\text{softmax}(\mathbf{z}_i/T), y_i)$   
**end for**

**return**  $\theta, P_1, P_2$

---

parameters that are optimized using the loss function described in (12). The term  $-\lambda_1\theta_1$  is designed to lower the temperature, whereas  $\lambda_2\theta_2$  is included to raise the temperature. Raising the temperature makes the distribution of the logits to be more uniformly distributed, ultimately reducing confidence in predictions. These adaptive terms adjust the scaling factor depending on whether a given sample is likely to be correctly classified or not.

We obtain  $\lambda_1$  and  $\lambda_2$  in (9), as follows:

$$\lambda_1 = P_{correct}(\mathcal{F}(\mathbf{x})), \text{ and } \lambda_2 = P_{incorrect}(\mathcal{F}(\mathbf{x})), \quad (10)$$

where  $\mathcal{F}$  denotes the energy score function in (6), and  $P_{correct}$  and  $P_{incorrect}$  denote the probability density functions of  $\mathcal{F}(\mathbf{x})$  fitted to Gaussian distributions for correct and incorrect samples, respectively. By leveraging the energy scores that characterize a specific classifier on a given dataset, we can establish a distribution for these energy scores and subsequently calculate the corresponding probability density function. The energy scores corresponding to correct instances are utilized to construct the distribution of  $P_{correct}$ , whereas the energy scores associated with incorrect instances are used to form the distribution of  $P_{incorrect}$ . We utilize an ID validation set and a semantic OOD dataset to make the aforementioned probability density functions, in which the predictions made using semantic OOD data are categorized as incorrect instances, since all predictions correspond to one of the  $K$  in-distribution classes.

By designing the proposed scaling factor in the manner described above, our method can gain the ability to instance-wisely distinguish between incorrect and correct samples. Ensuring the adaptability of our calibration method to each

**Algorithm 2** Proposed Instance-wise Calibration Method

---

**Require:** Test data :  $\mathcal{D}$ , Pre-trained  $K$ -class classifier  $f(\mathbf{x})$ , Temperature value from temperature scaling:  $T_{ts}$ , Trained parameters:  $\theta$ , PDFs:  $P_1, P_2$

$(\mathbf{x}, y) \sim \mathcal{D}$   
 $\mathbf{z} = f(\mathbf{x})$   
 $\mathcal{F}(\mathbf{z}) \leftarrow -\log\text{sumexp}(\mathbf{z})$   
 $T = T_{ts} - P_1(\mathcal{F}(\mathbf{z}))\theta_1 + P_2(\mathcal{F}(\mathbf{z}))\theta_2$   
 $\hat{q} = \max(\text{softmax}(\mathbf{z}/T))$

**return** calibrated confidence  $\hat{q}$

---

input sample is crucial for capturing the uncertainty of a particular prediction, especially in the presence of distribution shifts. This is the main reason on the robust calibration performance exhibited by our method across various test data, including ID samples and various types of distribution shifted samples (Fig. 4), which is empirically demonstrated in the experiment section. Then, using our instance-wise scaling with the scaling factor in (9), the calibrated probability of sample  $\mathbf{x}$  for  $K$  classes,  $\hat{p}_\theta \in \mathbb{R}^K$ , can be expressed as follows:

$$\hat{p}_\theta = \sigma_{SM}(f(\mathbf{x})/h(T_{ts}, \mathbf{x}; \theta)), \quad (11)$$

where  $h(T_{ts}, \mathbf{x}; \theta)$  denotes the proposed scaling factor defined in (9). To train the parameters  $\theta = \{\theta_1, \theta_2\}$  in (11), we design a mean squared error loss function:

$$\mathcal{L}_\theta = \frac{1}{N} \sum_i \|y^{(i)} - \hat{p}_\theta^{(i)}\|_2^2, \quad (12)$$

where  $y^{(i)} \in \mathbb{R}^K$  denotes the one-hot encoded ground-truth label for the  $i$ -th sample,  $\hat{p}_\theta^{(i)} \in \mathbb{R}^K$  denotes  $\hat{p}_\theta$  in (11) for the  $i$ -th sample, and  $N$  is the total number of training samples. Using the trained parameter  $\theta$ , the calibrated confidence for a test sample  $\mathbf{x}$  can be calculated in an instance-wise manner:

$$\hat{q} = \max(\sigma_{SM}(f(\mathbf{x})/h(T_{ts}, \mathbf{x}; \theta))). \quad (13)$$

Algorithm 1 describes the entire procedure for training calibration parameters, while Algorithm 2 outlines the procedure of implementing the instance-wise calibration method with the trained parameters.

## 5 Experiment

### 5.1 Experimental Settings

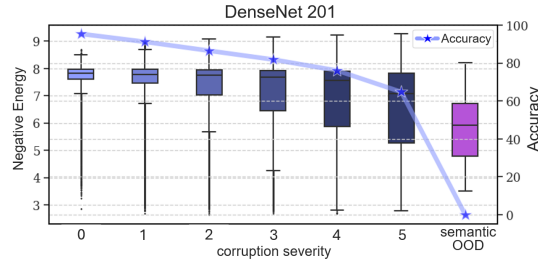
For experiments, we trained classification DNNs including VGGNet [26], ResNet [6], WideResNet [35], DenseNet [11] and SE-ResNet [10] on the CIFAR10/CIFAR100 datasets [12]. We employed pre-trained weights implemented in PyTorch for

ImageNet-1k [2]. We utilized two types of datasets for training: one intended for ID scenarios and the other for semantic OOD scenarios. For tuning  $\theta$ , we utilized ID validation images (5,000 for CIFAR10/CIFAR100, 12,500 for ImageNet-1k) and semantic OOD samples (100/400/3,500 SVHN [22] or Texture [1] images for CIFAR10/CIFAR100/ImageNet-1k, respectively).

To evaluate our method, we employed 10,000/10,000/12,500 images each from CIFAR10/CIFAR100/ImageNet-1k as test ID samples. For covariate OOD test data, we utilized the corrupted dataset CIFAR10-C, CIFAR100-C, and ImageNet-C in [8]. These corrupted datasets contain five severity levels for 19 corruption types (*e.g.* blur, contrast, and frost). We utilized 10,000 images for each severity level of corruption types, maintaining the same approach across all datasets. For the test semantic OOD scenarios, we used Textures or SVHN dataset, which were not used during the tuning time.

To demonstrate the effectiveness of our method, we compared it with five baseline post-hoc calibration methods, which are TS [4], ETS [36], IRM [36], IROvA/IROvATS [36], and SPLINE [5]. Among these methods, TS, ETS, IRM, SPLINE, and our method are considered as accuracy-preserving methods. In addition, we compared our method with the state-of-the-art post-hoc calibration method, DAC, proposed in [30]. Please note that additional experimental results based on the type of semantic OOD dataset can be found in the **supplementary materials**.

## 5.2 Ablation Study on Energy Score



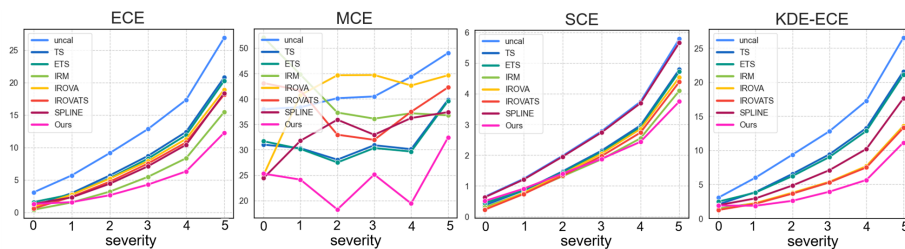
**Fig. 3: Negative energy scores and accuracy from ID to semantic OOD data.** The x-axis shows the severity levels of corruptions, while the left y-axis shows the negative energy scores (box plot) and the right y-axis shows accuracy (line plot). Severity level 0 (light blue) indicates the negative energy score for ID data (CIFAR10), levels 1-5 (deepening shades of blue) represent increasing degrees of corruption (CIFAR10-C), and the negative energy score for semantic OOD data (SVHN) is depicted in purple.

We conducted an ablation study on the energy score in (6). To demonstrate the capability of our method in capturing network uncertainty, we evaluated the energy score on various samples from complete ID dataset (CIFAR10), corrupted dataset (CIFAR10-C) and semantic OOD dataset (SVHN). As shown in Fig. 3, the energy scores tend to decrease with higher variances as the degree of distribution shift increases. Since the tendency for energy scores to decrease follows the trend of diminishing accuracy, it implies that the energy scores can indeed efficiently capture the uncertainty of DNNs.

### 5.3 Comparison with Baseline Methods

**Table 1: Averaged ECE over all severity levels of corruption.** The lowest ECE results were highlighted in bold and the parentheses indicates the OOD dataset used for tuning our method. It can be observed that our method has a notable calibration effect across various DNN architectures and all corruption levels from 0 to 5.

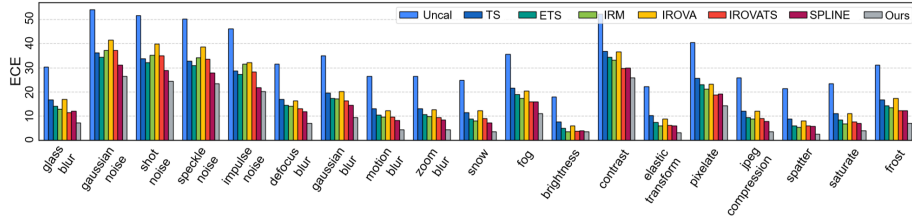
		ECE ( $\downarrow$ )							
ID (OOD)	Network	Uncalibrated	TS	ETS	IRM	IROvA	IROvATS	SPLINE	Ours
CIFAR10 (SVHN)	DenseNet201	12.90	9.30	8.95	5.78	6.76	6.05	7.40	<b>4.47</b>
	VGG19	16.17	10.59	10.34	7.18	7.92	7.33	7.76	<b>5.78</b>
	WideResNet40	12.49	8.57	8.33	5.72	7.80	7.35	7.30	<b>4.71</b>
	SE-ResNet50	14.80	10.17	9.94	8.13	8.51	8.34	9.09	<b>6.28</b>
CIFAR100 (SVHN)	DenseNet201	22.79	11.21	9.73	8.14	10.88	8.02	7.52	<b>5.94</b>
	VGG19	32.49	12.54	11.99	8.65	11.42	8.55	8.23	<b>7.35</b>
	WideResNet40	23.83	13.32	11.17	10.50	13.88	9.69	9.90	<b>8.74</b>
	SE-ResNet101	30.27	15.24	12.89	10.17	13.59	10.67	10.46	<b>8.95</b>
ImageNet-1k (Texture)	DenseNet201	6.99	5.49	5.49	6.37	10.67	8.95	6.44	<b>4.92</b>
	ResNet18	7.32	5.57	5.57	6.77	9.58	8.06	6.83	<b>5.31</b>



**Fig. 4: Calibration errors across different levels of corruption severity.** We can observe that the our method also benefits from the calibration effect in complete ID (severity level 0), while outperforming other approaches in the remaining severity levels. We used WideResNet40 trained on CIFAR10.

We compared our method with the aforementioned baseline methods by evaluating them across various datasets and backbone architectures in terms of ECE. Because our method emphasizes robust calibration performance on diverse datasets, we comprehensively conducted experiments on a variety of distribution shift scenarios, spanning from complete in-distribution to heavily corrupted scenarios. For this purpose, we employed a corrupted dataset comprising severity levels ranging from 1 to 5, with complete ID test data added as severity level 0.

Table 1 shows the averaged ECE across all severity levels. Our method outperforms other baseline methods for various backbone networks and datasets. Furthermore, as shown in Fig. 4, our method surpassed other approaches in most individual severity levels. It is noteworthy that our method shows consistent performance not only in scenarios involving corruptions but also in cases of complete ID scenarios. This is in contrast to [29], which exhibited greater miscalibration even than the uncalibrated one in the context of complete in-distribution data. Fig. 5 shows a comparison of ECE by corruption type, similar to [7, 29]. Our method demonstrates robust calibration across various corruption types. Additional results on diverse calibration metrics [23, 36] and transformer-based models [3, 18] are available in the **supplementary materials**.



**Fig. 5: Expected Calibration Error (ECE) for various corruption types at severity level 5.** It is evident that our proposed method exhibits superior performance across various types of corruption when compared to other methods, using DenseNet201 trained on the CIFAR100 dataset.

#### 5.4 Exploring Synergies with Applicable State-of-the-art Method

We analyzed the results of applying an applicable post-hoc calibration method to our proposed method. DAC [30], similar to our objective, aims for robust calibration performance even in OOD scenarios. Unlike most methods, including ours, which use only the output of the last layer of the DNN, DAC additionally utilized the output of other layers. Since DAC is designed to be used alongside post-hoc calibration methods, we applied it to our proposed method. We followed the DAC’s layer selection method proposed by Tomani et al. [30]. Table 2 shows the averaged ECE for each corrupted dataset. We compared our method with ETS+DAC and SPLINE+DAC, both of which primarily achieved state-of-the-art in [30]. While our method showed good synergy with DAC, it alone achieved the best performance even without DAC. Notably, our approach can attain these results without needing the additional output information from each classifier layer used by DAC.

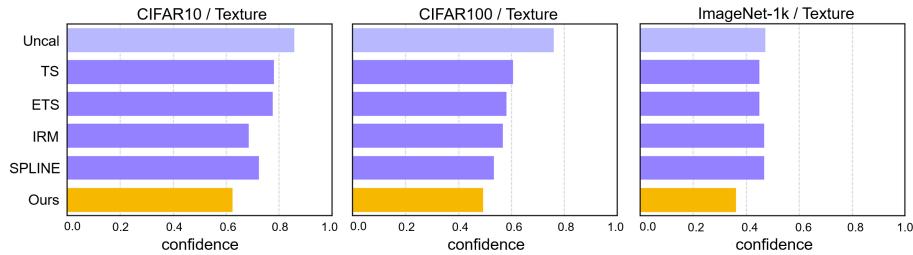
**Table 2: ECE comparison on state-of-the-art methods.** The best results were highlighted in bold, while the second best results were emphasized using underline. Our method exhibits comparable performance on its own, even without appending DAC. VGG19 was used for CIFARs, and DenseNet201 for ImageNet.

Dataset	Uncal.	ETS		SPLINE		Ours	
		w/o DAC	w/ DAC	w/o DAC	w/ DAC	w/o DAC	w/ DAC
CIFAR10-C	15.49	12.33	8.59	15.12	8.45	<b>5.98</b>	<u>7.81</u>
CIFAR100-C	31.83	22.21	11.97	27.97	12.28	<u>11.92</u>	<b>8.99</b>
ImageNet-C	6.84	7.34	<u>5.07</u>	8.05	5.96	<b>4.96</b>	5.72

#### 5.5 Evaluating Robustness on Semantic OOD

We measured the calibrated confidence scores for semantic OOD test samples, which are different from OOD train samples used to adjust our calibration parameters. For this experiment, we utilized DenseNet201. Fig. 6 demonstrates that our method produces the lowest confidence scores for OOD samples, because all predictions for OOD samples are incorrect.

Furthermore, we conducted an experiment to investigate the potential extension into OOD detection. To accomplish this, we utilized key evaluation metrics commonly employed in OOD detection, such as AUROC, AUPR-in, and



**Fig. 6: Calibrated confidence scores for semantic OOD samples.** Each bar denotes the confidence score on semantic OOD samples. The shorter bar represents better-calibrated confidence because all predictions for semantic OOD samples are incorrect; thus, the lower confidence score, the better. For this experiment, we used CIFARs and ImageNet-1k as ID data, while Texture served as the OOD test data.

**Table 3: OOD detection comparison.** In (a) and (b),  $[\cdot/\cdot]$  represents the ID/OOD data. In particular, we obtained the results using DenseNet201 trained on CIFAR100 and ImageNet-1k as ID data, while Texture served as the OOD test data.

(a) [CIFAR100 / Texture]				(b) [ImageNet-1k / Texture]			
	AU-ROC( $\uparrow$ )	AUPR-IN( $\uparrow$ )	AUPR-OUT( $\uparrow$ )		AU-ROC( $\uparrow$ )	AUPR-IN( $\uparrow$ )	AUPR-OUT( $\uparrow$ )
Uncal.	74.21	70.62	71.25	Uncal.	81.43	83.50	79.17
TS	74.84	<b>70.89</b>	72.27	TS	81.68	83.68	79.50
ETS	74.86	<b>70.89</b>	72.27	ETS	81.68	83.68	79.50
IRM	71.73	69.57	72.16	IRM	80.25	<b>84.23</b>	79.81
SPLINE	73.76	70.27	70.98	SPLINE	81.43	83.52	79.17
Ours	<b>75.03</b>	<b>70.89</b>	<b>72.37</b>	Ours	<b>82.25</b>	83.82	<b>80.14</b>

AUPR-out. Our method demonstrates superior performance compared to other approaches in most cases, as observed in Table 3, thus affirming the potential for extension into OOD detection with our proposed approach.

## 6 Conclusion

In this paper, we addressed the limitations of existing post-hoc calibration methods on the wild datasets, including in-distribution, covariate shift, and semantic shift. Conventional methods could not consider all these scenarios in achieving robust calibration. To solve this problem, we introduced a novel instance-wise calibration method using the energy score. Our method adaptively captured uncertainty for each instance by leveraging the energy score. Through experiments conducted across various networks and datasets, we demonstrated that our method outperforms existing calibration methods in scenarios involving various types of distribution shifts, while consistently maintaining calibration effect in the complete in-distribution dataset.

As the reliability of AI in safety-critical situations becomes increasingly important, we believe that our method can contribute to the safer deployment of AI systems in real-world scenarios. By offering a promising direction with our method, we hope to inspire future research efforts for enhancing trustworthy AI.

**Acknowledgements.** This work was partly supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (NRF2020R1C1C1004907) and partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (RS-2022-00143911, AI Excellence Global Innovative Leader Education Program and 2021-0-01341, Artificial Intelligence Graduate School Program (Chung-Ang university)).

## References

1. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: CVPR (2014) 11
2. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009) 11
3. 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. arXiv preprint arXiv:2010.11929 (2020) 12
4. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: ICML (2017) 2, 3, 8, 11
5. Gupta, K., Rahimi, A., Ajanthan, T., Mensink, T., Sminchisescu, C., Hartley, R.: Calibration of neural networks using splines. arXiv preprint arXiv:2006.12800 (2020) 2, 4, 11
6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016) 10
7. Hekler, A., Brinker, T.J., Buettner, F.: Test time augmentation meets post-hoc calibration: Uncertainty quantification under real-world conditions. In: AAAI (2023) 3, 12
8. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. In: ICLR (2019) 11
9. Hsu, Y.C., Shen, Y., Jin, H., Kira, Z.: Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data. In: CVPR (2020) 6
10. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: CVPR (2018) 10
11. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: CVPR (2017) 10
12. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009) 10
13. Kull, M., Perello Nieto, M., Kängsepp, M., Silva Filho, T., Song, H., Flach, P.: Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with dirichlet calibration. In: NIPS (2019) 4
14. Kull, M., Silva Filho, T., Flach, P.: Beta calibration: a well-founded and easily implemented improvement on logistic calibration for binary classifiers. In: AISTAT (2017) 4
15. LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., Huang, F.: A tutorial on energy-based learning. Predicting structured data (2006) 7
16. Liu, B., Ben Ayed, I., Galdran, A., Dolz, J.: The devil is in the margin: Margin-based label smoothing for network calibration. In: CVPR (2022) 3

17. Liu, W., Wang, X., Owens, J., Li, Y.: Energy-based out-of-distribution detection. In: NeurIPS (2020) [2](#), [6](#), [7](#), [8](#)
18. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: ICCV (2021) [12](#)
19. Mukhoti, J., Kulharia, V., Sanyal, A., Golodetz, S., Torr, P., Dokania, P.: Calibrating deep neural networks using focal loss. In: NeurIPS (2020) [2](#), [3](#)
20. Müller, R., Kornblith, S., Hinton, G.E.: When does label smoothing help? In: NIPS (2019) [2](#), [3](#)
21. Naeini, M.P., Cooper, G., Hauskrecht, M.: Obtaining well calibrated probabilities using bayesian binning. In: AAAI (2015) [3](#), [5](#)
22. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning. In: NIPS Workshop (2011) [11](#)
23. Nixon, J., Dusenberry, M.W., Zhang, L., Jerfel, G., Tran, D.: Measuring calibration in deep learning. In: CVPR workshops (2019) [5](#), [12](#)
24. Patel, V.M., Gopalan, R., Li, R., Chellappa, R.: Visual domain adaptation: A survey of recent advances. IEEE signal processing magazine (2015) [6](#)
25. Rahimi, A., Shaban, A., Cheng, C.A., Hartley, R., Boots, B.: Intra order-preserving functions for calibration of multi-class neural networks. In: NeurIPS (2020) [2](#), [3](#)
26. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014) [10](#)
27. Thulasidasan, S., Chennupati, G., Bilmes, J.A., Bhattacharya, T., Michalak, S.: On mixup training: Improved calibration and predictive uncertainty for deep neural networks. In: NIPS (2019) [3](#)
28. Tomani, C., Cremers, D., Buettner, F.: Parameterized temperature scaling for boosting the expressive power in post-hoc uncertainty calibration. In: ECCV (2022) [4](#)
29. Tomani, C., Gruber, S., Erdem, M.E., Cremers, D., Buettner, F.: Post-hoc uncertainty calibration for domain drift scenarios. In: CVPR (2021) [2](#), [4](#), [6](#), [12](#)
30. Tomani, C., Waseda, F.K., Shen, Y., Cremers, D.: Beyond in-domain scenarios: Robust density-aware calibration. In: ICML (2023) [4](#), [11](#), [13](#)
31. Wenger, J., Kjellström, H., Triebel, R.: Non-parametric calibration for classification. In: AISTAT (2020) [3](#)
32. Yang, J., Zhou, K., Li, Y., Liu, Z.: Generalized out-of-distribution detection: A survey. arxiv. arXiv preprint arXiv:2110.11334 (2021) [6](#)
33. Zadrozny, B., Elkan, C.: Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. In: ICML (2001) [3](#)
34. Zadrozny, B., Elkan, C.: Transforming classifier scores into accurate multiclass probability estimates. In: ACM SIGKDD (2002) [3](#)
35. Zagoruyko, S., Komodakis, N.: Wide residual networks. arXiv preprint arXiv:1605.07146 (2016) [10](#)
36. Zhang, J., Kailkhura, B., Han, T.Y.J.: Mix-n-match: Ensemble and compositional methods for uncertainty calibration in deep learning. In: ICML (2020) [2](#), [3](#), [4](#), [5](#), [11](#), [12](#)
37. Zhong, Z., Cui, J., Liu, S., Jia, J.: Improving calibration for long-tailed recognition. In: CVPR (2021) [3](#)