Contents lists available at ScienceDirect







journal homepage: www.elsevier.com/locate/patcog

FocusNet: Classifying better by focusing on confusing classes

Xue Zhang^a, Zehua Sheng^a, Hui-Liang Shen^{a,*}

College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou 310027, China

ARTICLE INFO

Article history: Received 9 August 2021 Revised 14 March 2022 Accepted 10 April 2022 Available online 13 April 2022

Keywords: Image classification Inter-class correlations Confusing classes

ABSTRACT

Nowadays, most classification networks use one-hot encoding to represent categorical data because of its simplicity. However, one-hot encoding may affect the generalization ability as it neglects inter-class correlations. We observe that, even when a neural network trained with one-hot labels produces incorrect predictions, it still pays attention to the target image region and reveals which classes confuse the network. Inspired by this observation, we propose a confusion-focusing mechanism to address the class-confusion issue. Our confusion-focusing mechanism is implemented by a two-branch network architecture. Its baseline branch generates confusing classes, and its FocusNet branch, whose architecture is flexible, discriminates correct labels from these confusing classes. We also introduce a novel focus-picking loss function to improve classification accuracy by encouraging FocusNet to focus on the most confusing classes. The experimental results validate that our FocusNet is effective for image classification on common datasets, and that our focus-picking loss function can also benefit the current neural networks in improving their classification accuracy.

© 2022 Elsevier Ltd. All rights reserved.

1. Introduction

Neural network based image classification has received much attention in recent years due to the impressive performance of deep learning [1,3,4]. Most classification networks use one-hot encoding to represent categorical data. However, this encoding method neglects inter-class correlations, and may cause over-fitting as it tends to make the network produce overconfident predictions on the training sets [5,6]. To address this problem, label smoothing regularization [5] introduces a uniform distribution to encode correlations between classes, which encourages networks to be less confident. Considering the uniform distribution is independent of the training data, some approaches [4,7] use soft targets to replace it with a learned distribution. Alternatively, DropMax [8] drops non-target classes to select similar ones that confuse networks by using an additional multilayer perceptron. To our knowledge, there are few works exploring the confusing-class correlation for the purpose of image classification.

In this work, we propose the confusion-focusing mechanism and a focus-picking loss function to handle the class-confusion issue. Fig. 1 illustrates the overall architecture of our confusionfocusing mechanism, which consists of a baseline branch and a FocusNet branch. The baseline network is pre-trained to produce reference predictions. FocusNet, which can be implemented using common models such as ResNet [1] and MobileNet [3], identifies confusing classes and difficult samples based on the output of the baseline network. Furthermore, our proposed focus-picking loss function focuses more on these confusing classes and assigns larger losses to difficult samples in the training stage. It is worth noting that, only the FocusNet branch is needed for inference, so our approach does not add extra computing complexity. With this strategy, FocusNet can effectively improve the performance of the baseline network. As illustrated in Fig. 1(b) and (c), compared with the baseline network, FocusNet pays more attention to the target region and produces the correct prediction. To summarize, the main contributions of this work are threefold:

- We propose the confusion-focusing mechanism implemented by a two-branch network architecture to address the challenging class-confusion issue. The FocusNet branch can improve the performance by focusing on confusing classes derived from the baseline branch without adding extra computing complexity during inference.
- We introduce a focus-picking loss function for network training. Compared with normal cross-entropy, our loss function enables the network to focus on confusing classes and assign larger losses to difficult samples.
- We experimentally validate on various datasets that FocusNet outperforms the existing approaches in addressing the classconfusion issue, and that our focus-picking loss function can benefit the current knowledge distillation method in the accuracy improvement.

^{*} Corresponding author.

E-mail addresses: zxue2019@zju.edu.cn (X. Zhang), shengzehua@zju.edu.cn (Z. Sheng), shenhl@zju.edu.cn (H.-L. Shen).



Fig. 1. The overall architecture of our confusion-focusing mechanism. (a) Training and inference framework of our approach. After training, only the FocusNet branch is needed for inference. (b) Prediction of the baseline network. (c) Predictions of our FocusNet. The pictures next to the bar charts are Ridgeback's CAM (green border) and Dingo's CAM (red border). Note that the baseline and FocusNet can use the same or different network architectures. Here, for demonstration, the network architectures are both ResNet-18 [1], and the dataset is Imagewoof [2]. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

2. Related work

We briefly review the related works of network architecture design and loss function design, which are relevant to our novel network mechanism and loss function proposed in this work.

2.1. Network architecture design

Over the past decade, many advanced network architectures have been introduced to improve accuracy in image classification, including networks that increase in depth [1], width [5], cardinality [9] and scale [10]. Recently, inter-class correlations [4,5] and attention mechanisms [11,12], as two approaches to further improve the performance of existing networks, have attracted much attention.

It has been found that learning inter-class correlations is an effective way to improve the generalization ability because it prevents the networks from producing overconfident predictions [5]. There are two main ways to use the correlation between classes. The first way uses the manually designed distribution over classes, which does not change the original network architecture. For example, label smoothing regularization [5] mixes a uniform distribution with the ground-truth distribution by a smoothing factor. The work [13] introduces a random relationship between classes by encouraging networks to produce high entropy output distribution. The second way uses the learned distribution for each sample and typically designs a network architecture. For example, the label embedding network [14] generates a soft distribution to represent the continuous interactions between classes. RDCN [15] introduces a relation network to measure the relationship between the visual features and embedded semantics. Knowledge distillation [4] introduces soft targets to replace the hand-designed distribution. Standard knowledge distillation [4,16] is a combination of a pre-trained teacher model, and a smaller student model to be deployed. The recent self-knowledge distillation mechanism [17,18] refines knowledge using an auxiliary self-teacher model. Teach-free knowledge distillation [7] combines the knowledge distillation and the label smoothing regularization to produce a more accurate distribution over classes. Instead of considering the correlation between all classes, DropMax [8] learns the class retain probability of each sample to drop non-target classes and select the confusing ones.

Another effective approach to improve the accuracy of networks is to use the attention mechanism. For example, VINet [11] integrates human visual attention into neural networks in image classification and object detection. AE-Net [19] refines channel and spatial attention to retrieve fine-grained sketch-based images. DAAF [12] learns global and local attention aware features to improve the accuracy of person re-identification. Progressiveattention network [20] localizes discriminative parts at multiple scales progressively.

The main difference between these methods and our work is that FocusNet establishes the relationship between confusing classes and network attention. In addition, our FocusNet does not change the architecture of the original network, and does not add extra computing complexity during inference.

2.2. Loss function design

In supervised learning, networks are trained through an optimization process that minimizes a loss function. In this case, most networks for multi-class classification use the softmax crossentropy as their loss function [1]. Minimizing the cross-entropy is equivalent to minimizing the Kullback-Leibler divergence (hereafter written as KL divergence) when measuring the degree of dissimilarity between empirical distribution and model distribution [21]. Currently, some variants of the softmax cross-entropy have been introduced to solve task-specific issues [4,8,22]. Focal loss [22] addresses the class imbalance problem by introducing a tunable weighting factor to the cross-entropy. Adjusting this weighting factor allows networks to reduce relative losses for wellclassified samples and focus more on hard and misclassified ones. DropMax [8] focuses on confusing classes by dropping non-target logits in the softmax layer, which uses the Bayesian inference to calculate class retain probabilities and uses variational inference to optimize parameters. Knowledge distillation [4] transfers knowledge from a large teacher to an easy-deployed student by minimizing the KL divergence of soft targets between the teacher and the student, which is mathematically equivalent to optimizing the mean square error of logits between the two models. LSA [23] considers the low-rank structure of the reconstruction data in tackling the domain bias problem.

Unlike the above methods, our focus-picking loss function mainly considers the most confusing inter-class interactions and assigns larger losses to more difficult samples.

3. Proposed method

In this section, we first introduce our motivation. Then, we elaborate on the proposed FocusNet and the focus-picking loss function. To better clarify the concept of focus-picking loss, we also analyze its mathematical mechanism.

3.1. Motivation

We notice that similar classes tend to confuse networks. As illustrated in Fig. 1(b), the baseline network yields an incorrect prediction. The ground-truth class is Dingo, but the top-1 predicted

one is Ridgeback. Nevertheless, the CAM of the class Dingo (red border) shows that the baseline network has paid attention to the target image region, *i.e.* the Dingo's body. In other words, the baseline network has noticed the correct class, but it is confused by incorrect classes and cannot make the correct prediction. Inspired by this, we aim to address the class-confusion issue by making networks discriminate the correct class from these confusing ones.

Fig. 1(a) illustrates our proposed confusion-focusing mechanism, which combines a pre-trained baseline network branch (for generating reference distributions at training time) and a Focus-Net branch (for making predictions). It is worth noting that the FocusNet branch is not restricted to a specific network architecture. It uses the reference distribution derived from the baseline network to determine confusing classes based on the above observation. In addition, FocusNet compares the probability on the correct class with the reference distribution to identify difficult samples. In this context, we design the focus-picking loss function to focus on picking the correct class from the confusing ones during training. The baseline network branch is removed during inference to keep the model size of our approach unchanged. With this approach, Fig. 1(c) shows that our FocusNet correctly predicts Dingo. Besides, the CAM of the class Dingo (red border) illustrates that our FocusNet pays more attention to the target image region.

From the perspective of requiring prior knowledge from a pretrained network, the network architecture of our approach is similar to standard knowledge distillation to some extent. But we have validated that our confusion-focusing mechanism is effective even when 1) the pre-trained baseline network uses various network scales, *i.e.*, the same as, or larger or smaller than FocusNet architecture, and 2) the pre-trained network is highly incorrect. In comparison, however, the standard knowledge distillation generally requires a larger and more accurate pre-trained model. Please refer to Section 3.6 for more detailed analysis.

3.2. Confusion-focusing mechanism

As mentioned, our confusion-focusing mechanism can obtain consistent performance improvements with a variety of network architecture combinations. Fig. 1(a) illustrates an example pipeline that adopts ResNet-18 [1] as the network architecture of the baseline network branch and the FocusNet branch.

3.2.1. Training and inference of the baseline network branch

Given an input, the baseline network branch in Fig. 1(a) produces logits \mathbf{z}^b through a sequence of convolutional layers and a fully-connected layer. Then the resulting logits are converted to the probability distribution $\mathbf{\hat{y}}^b$ using a softmax function. The process can be formulated as

$$\mathbf{z}^{b} = \left(\mathbf{W}^{b}\right)^{\mathsf{T}} \mathbf{h}^{b},\tag{1}$$

$$\hat{y}_n^b = \frac{\exp(z_n^b)}{\sum_{i=1}^N \exp(z_i^b)},\tag{2}$$

where matrix $\mathbf{W}^{b} = (\mathbf{w}_{1}^{b}, \mathbf{w}_{2}^{b}, \dots, \mathbf{w}_{N}^{b})$ contains the weights and biases of the fully-connected layer in the baseline network branch, and the column vector $\mathbf{w}_{n}^{b}, 1 \leq n \leq N$, denotes the template of the *n*-th class [6]. Vector \mathbf{h}^{b} contains the activations of the penultimate layer concatenated with "1" accounting for the bias. z_{n}^{b} and \hat{y}_{n}^{b} denote the *n*-th entries of the vectors \mathbf{z}^{b} and $\hat{\mathbf{y}}^{b}$, respectively.

The baseline network branch is trained by computing the crossentropy between the one-hot label **y** and probabilities $\hat{\mathbf{y}}^b$,

$$\mathcal{L}^{b} = H(\mathbf{y}, \hat{\mathbf{y}}^{b})$$

= $-\sum_{n=1}^{N} y_{n} \log(\hat{y}_{n}^{b}),$ (3)

where y_n denotes the *n*-th entry of **y**.

3.2.2. Training and inference of the FocusNet branch

Our FocusNet is trained on the same dataset by minimizing the proposed focus-picking loss function. Specifically, the green arrows in Fig. 1(a) illustrate that the training framework of our approach consists of a pre-trained baseline network branch and a Focus-Net branch. We note that the baseline branch is fixed when training our FocusNet. Each training sample is fed into both branches at the same time to produce two sets of logits and probabilities. Then, FocusNet uses the logits of the baseline branch to deduce the multi-warm label that represents the confusing classes. In addition, FocusNet compares the probabilities on the ground-truth class with the baseline prediction to determine if the training sample is difficult to classify. Our focus-picking loss function encourages FocusNet to learn the multi-warm label and adaptively tunes its value according to the classification difficulty of the training sample.

After completing the training stage, we only need the FocusNet branch for inference, as the red arrow in Fig. 1(a) indicates. As a result, our approach adds no extra computing complexity during inference.

3.3. Multi-warm label

To guide the computation of FocusNet attention, we employ the pre-trained baseline network branch to determine confusing classes for each training sample, mathematically represented as a multi-warm label. Specifically, the baseline network generates logits over classes based on Eq. (1). In this context, the logit $z_n^b =$ $(\mathbf{w}_n^b)^{\mathsf{T}} \mathbf{h}^b$ of the *n*-th class measures the correlation between the vector \mathbf{h}^b and the template \mathbf{w}_n^b . A positive or negative correlation value indicates that the vector and the template point in close or opposite directions, while a zero value indicates that they are orthogonal.

The *n*-th entry of the multi-warm label vector is defined as

$$l'_{n} = \begin{cases} 1, & \text{if } z_{n}^{b} > 0\\ 0, & \text{if } z_{n}^{b} \le 0. \end{cases}$$
(4)

Then, we add the ground-truth label of the training sample to the multi-warm label and clip it for learning stability [8],

$$l_n'' = \min(1, l_n' + y_n).$$
(5)

After normalizing each label entry by

$$l_n = \frac{l_n''}{\sum_{j=1}^N l_j''},$$
(6)

we obtain the final multi-warm label $\mathbf{l} = (l_1, l_2, \dots, l_N)^{\mathsf{T}}$. The non-zero entries of the label vector \mathbf{l} have the same probability and sum to 1.0.

The multi-warm label reflects the attention of the pre-trained baseline network. In other words, it converts the implicit attention of the baseline network into explicit knowledge that FocusNet can learn directly. In Fig. 2, we show the prediction examples of three scenes, each of which contains the CAMs of different classes, the corresponding logits, and the multi-warm labels. For Scene-A, the baseline network pays attention to Beagle (index: 2, ground truth), Samoyed (index: 8), and Dingo (index: 9). It generates positive logits on these classes as they are similar in appearance, implying that the multi-warm label reflects the attention of the baseline network. Scene-B and Scene-C both contain Dingoes, but with different fur colors (yellow and white, respectively). The baseline network generates significantly different CAMs and multi-warm labels, demonstrating that the multi-warm labels are adaptive. In this context, the multi-warm label can indicate which classes are confusing.



Fig. 2. The relationship between the attention of the pre-trained baseline network and our proposed multi-warm label. The pre-trained baseline network processes the input image in each scene to generate CAMs, logits, and the multi-warm label. CAMs highlight discriminative image regions for corresponding classes. The positive and negative logits are marked in red and gray, respectively. Non-zero entries of the multi-warm label and the CAMs of confusing classes are marked in red bars and red borders, respectively. The "Class" row presents categorical indexes and categorical names. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

3.4. Loss function

FocusNet converts logits \boldsymbol{z} into predicted probabilities $\boldsymbol{\hat{y}},$ computed by

$$\hat{y}_n = \frac{\exp(d_n + z_n)}{\sum_{j=1}^{N} \exp(d_j + z_j)}.$$
(7)

Here,

$$d_n = \frac{\exp(z_n)}{\sum_{j=1}^{N} \exp(z_j)} - \frac{\exp(z_n^b)}{\sum_{j=1}^{N} \exp(z_j^b)},$$
(8)

is the *n*-th entry of the vector **d**. z_n and \hat{y}_n are the logit and the predicted probability of FocusNet on the *n*-th class, respectively. Vector **d** denotes the probability difference over classes between FocusNet and the pre-trained baseline network. The normal softmax probability distribution in Eq. (2) is a special case of the reweighted probability distribution in Eq. (7) when $\mathbf{d} = \mathbf{0}$. For notation simplicity, we denote the normal and re-weighted probability distributions of FocusNet as $\hat{\mathbf{y}}(\mathbf{z}, \mathbf{d} = \mathbf{0})$ and $\hat{\mathbf{y}}(\mathbf{z}, \mathbf{d})$, respectively. A positive difference on the ground-truth class means that the input is a relatively easy sample since FocusNet generates a more confident prediction than the baseline network. In contrast, a negative difference corresponds to a relatively difficult sample. Here, the role of difference is to force FocusNet to tune the weights for logits so that easy samples and difficult samples can produce higher and lower probabilities on the ground-truth class, respectively. Then, we compute the cross-entropy between the re-weighted probabilities and the ground-truth label as the classification loss

$$\mathcal{L}_{cls} = H(\mathbf{y}, \hat{\mathbf{y}}(\mathbf{z}, \mathbf{d}))$$

= $-\sum_{n=1}^{N} y_n \log (\hat{y}_n(\mathbf{z}, \mathbf{d}))$
= $-\log \left(\frac{\exp(d_t + z_t)}{\sum_{i=1}^{N} \exp(d_i + z_i)}\right),$ (9)

where index *t* refers to the ground-truth class. This allows Focus-Net to take the prediction from the pre-trained baseline network as a reference and assigns larger losses to those relatively difficult samples (see Section 3.5 for more details). At test time, FocusNet computes the prediction by setting $\mathbf{d} = \mathbf{0}$, *i.e.* the normal softmax function.

To make FocusNet pay more attention to confusing classes during training, we propose to regularize it by minimizing the crossentropy between normal softmax probabilities and the multi-warm label. That is,

$$\mathcal{R}_{\text{attention}} = H(\mathbf{l}, \hat{\mathbf{y}}(\mathbf{z}, \mathbf{d} = \mathbf{0}))$$

$$= -\sum_{n=1}^{N} l_n \log \left(\hat{y}_n(\mathbf{z}, \mathbf{d} = \mathbf{0}) \right)$$

$$= -\sum_{n=1}^{N} l_n \log \left(\frac{\exp(z_n)}{\sum_{j=1}^{N} \exp(z_j)} \right).$$
(10)

This regularization encourages FocusNet to produce higher probabilities on confusing classes than the others. For example, *Scene-B* in Fig. 2 shows that the baseline network is confused by classes 1, 2, 4, and 9, respectively. Using our approach in Eq. (10), Focus-Net generates higher probabilities on these four classes than the others, as shown in the FocusNet branch of Fig. 1(c).

Last but not least, minimizing the classification loss in Eq. (9) encourages FocusNet to make more confident predictions on ground-truth classes than the baseline network. This may produce overconfident distributions, *i.e.* low entropy distributions [13]. To reduce this potential degradation, we penalize the low-entropy distributions computed as

$$\mathcal{R}_{\text{entropy}} = H(\hat{\mathbf{y}}(\mathbf{z}, \mathbf{d} = \mathbf{0}))$$

$$= -\sum_{n=1}^{N} \hat{y}_n(\mathbf{z}, \mathbf{d} = \mathbf{0}) \log (\hat{y}_n(\mathbf{z}, \mathbf{d} = \mathbf{0}))$$

$$= -\sum_{n=1}^{N} \frac{\exp(z_n)}{\sum_{j=1}^{N} \exp(z_j)} \log \left(\frac{\exp(z_n)}{\sum_{j=1}^{N} \exp(z_j)}\right).$$
(11)

Adding negative entropy to the classification loss allows FocusNet to produce smoother probability distributions with higher entropy and thus contributes to better generalization, which is similar to label smoothing regularization [5].

To train FocusNet, we propose the focus-picking loss function, which consists of the three components described above. It is defined as

$$\mathcal{L} = \mathcal{L}_{cls} + \alpha \mathcal{R}_{attention} - \beta \mathcal{R}_{entropy}$$

= $H(\mathbf{y}, \mathbf{\hat{y}}(\mathbf{z}, \mathbf{d})) + \alpha H(\mathbf{l}, \mathbf{\hat{y}}(\mathbf{z}, \mathbf{d} = \mathbf{0})) - \beta H(\mathbf{\hat{y}}(\mathbf{z}, \mathbf{d} = \mathbf{0})).$ (12)

It encourages FocusNet to generate larger losses on relatively difficult samples, focus on confusing classes, and prevent the lowentropy output distributions at the same time. We experimentally find that setting $\alpha = \beta = 1$ works well.

3.5. Analysis of the focus-picking loss function

Different from traditional loss functions, our focus-picking loss function considers confusing inter-class correlations and assigns



Fig. 3. The impact of d_t on \mathcal{L}_{cls} . d_t denotes the probability difference on the correct class between FocusNet and the baseline network. The two descriptions in the illustration present our motivation for designing \mathcal{L}_{cls} .

larger loss values to those difficult samples. For instance, the widely used cross-entropy with one-hot labels does not consider the interaction between classes [5] and the difficulty of each sample. KL divergence with soft targets is usually used in knowledge distillation [4,7]. But the soft targets tend to be a uniform distribution when the temperature is high [24], so it can not emphasize the importance of confusing classes. The focal loss [22] is a dynamically scaled cross-entropy to address the class imbalance problem in object detection, but it also does not consider the class-confusion issue.

We first explain the data term \mathcal{L}_{cls} . It takes the probability distribution derived from the baseline network as a reference, downweights the contribution of easy samples during training, and rapidly enables FocusNet to focus on difficult samples. We rewrite Eq. (9) as

$$\mathcal{L}_{cls} = -\log\left(\frac{\exp(d_t + z_t)}{\sum_{j=1}^{N} \exp(d_j + z_j)}\right)$$

= $-\log\left(\frac{\exp(z_t)}{\exp(z_t) + \sum_{\substack{j \neq t \ \exp(d_t) \ \exp(d_t)}} \exp(z_j)}\right),$ (13)

where $d_t = \hat{y}_t(\mathbf{z}, \mathbf{d} = \mathbf{0}) - \hat{y}_t^b$ denotes the probability difference on the ground-truth class between FocusNet and the baseline network.

Fig. 3 shows that \mathcal{L}_{cls} becomes progressively smaller as d_t changes from negative to positive. Specifically, a negative d_t indicates that FocusNet has lower confidence on the correct class than the baseline network, which implies the sample is difficult, and thus \mathcal{L}_{cls} is large. A positive d_t means the sample is easy, and hence \mathcal{L}_{cls} is small.

 $\mathcal{R}_{attention}$ encourages FocusNet to focus on the classes that confuse the baseline network. One can compute the derivative of $\mathcal{R}_{attention}$ with respect to logits z as

$$\frac{\partial \mathcal{R}_{\text{attention}}}{\partial \mathbf{z}} = \hat{\mathbf{y}}(\mathbf{z}, \mathbf{d} = \mathbf{0}) - \mathbf{I}, \tag{14}$$

where **l** is the multi-warm label encoding confusing classes, $\hat{y}(z, d = 0)$ is the predicted probability distribution of FocusNet. Making the derivative equal to **0** enables FocusNet to generate higher probabilities over the confusing classes while penalizing the outputs of irrelevant classes.

 $\mathcal{R}_{entropy}$ prevents FocusNet from producing over-confident predictions [13]. It is necessary because \mathcal{L}_{cls} encourages FocusNet to

make more confident predictions on correct classes than the baseline network, as discussed in Section 3.4.

3.6. Differences from standard knowledge distillation

Our proposed confusion-focusing mechanism differs from standard knowledge distillation mainly in the following four aspects:

- 1) Motivations are different. As described in Section 3.1, our motivation is to make FocusNet discriminate the correct class from the most confusing ones. But standard knowledge distillation aims to transfer knowledge from a teacher (a larger model or an ensemble of models) to a student (an easy-to-deployed model). We have validated in Table 9 that a poorly pre-trained network has less influence on our FocusNet than on the knowledge distillation.
- 2) In the case of learning confusing classes, the network architectures are different. FocusNet only trains one model. It directly makes predictions without using the pre-trained baseline network at test time. However, in the case of learning confusing classes, standard knowledge distillation needs to train a set of specialist models. It first runs a pre-trained generalist model to decide which specialist models are relevant at test time. Then these specialist models are run to determine the final prediction (please refer to Section 5 and Section 7 in standard knowledge distillation [4] for more details). Therefore, our proposed FocusNet is more efficient for learning confusing classes.
- 3) Labels are different. Our proposed multi-warm label emphasizes the importance of confusing classes. Soft targets [4] of standard knowledge distillation considers connections between all classes, but they are difficult to reflect the relationship between confusing classes. A convincing example is that as the temperature increases, the soft target tends to be a uniform distribution [24], which means that the correct class has a close probability of being similar to all other classes.
- 4) Loss functions are different. The proposed focus-picking loss function enables FocusNet to focus on confusing classes and assign larger losses to more difficult samples (as described in Section 3.5). The distillation loss function enforces the student model to learn soft targets of the teacher model [4]. We have validated in Table 7 that replacing the distillation loss function with our focus-picking loss function can further improve the performance of the state-of-the-art distillation methods.

4. Experiments

4.1. Datasets

Table 1 lists the datasets used in our experiments, including both fine-grained datasets (Nos. 6 and 7) and datasets of different scales. We choose the fine-grained dataset considering that it contains a large number of confusing classes.

Table 1

The datasets we used in this work. "*N*" is the number of classes, "# Training" and "# Validation" represent the number of images in the training and validation sets, respectively.

Datasets	Ν	# Training	# Validation
1. MNIST [25]	10	60,000	10,000
2. CIFAR-10 [26]	10	50,000	10,000
3. Imagewoof [2]	10	9,025	3,929
4. CIFAR-100 [26]	100	50,000	10,000
5. ImageNet Dogs [27]	118	147,873	5,900
6. Stanford Dogs [28]	120	12,000	8,580
7. CUB-200 [29]	200	5,994	5,794
8. Tiny ImageNet [30]	200	100,000	10,000
9. ImageNet [31]	1,000	1,281,167	50,000

MNIST [25] is a handwritten digits dataset with images of size 28×28 . It consists of 10 classes, with 60,000 training images and 10,000 validation images.

CIFAR-10 and CIFAR-100 [26] have 10 and 100 classes, respectively. They both contain 50,000 training images and 10,000 validation images. The size of the images is $32 \times 32 \times 3$.

Imagewoof [2] contains 10 dog breeds that are *not* easy to classify. It contains 9025 training images and 3929 validation images.

ImageNet Dogs [27] is composed of all dog breeds in the ImageNet dataset [31]. It has 118 classes with 147,873 images for training and 5900 images for validation. All images have been down-sampled to $64 \times 64 \times 3$ pixels.

Stanford Dogs [28] is a fine-grained dog breeds dataset. It contains 120 classes and 20,580 images, with 12,000 images for training and 8580 ones for validation.

CUB-200 [29] is a fine-grained bird species dataset. It contains 200 classes with 5994 training images and 5794 validation images.

Tiny ImageNet [30] is a subset of the ImageNet dataset [31]. This dataset contains 200 classes, each of which has 500 training images and 50 validation images. All images have been down-sampled to $64 \times 64 \times 3$ pixels.

ImageNet [31] is a large-scale dataset. It has 1000 classes, 1.2 million training images and 50,000 validation images.

4.2. Implementation details

We use LeNet-5 [25] for experiments on MNIST, use the standard ResNet-18 [1] and MobileNetV2 [3] for experiments on Imagewoof, Stanford Dogs and CUB-200, and use the standard ResNet-18 and ResNet-34 for experiments on ImageNet. Because the standard ResNet-18 and MobileNetV2 down-sample input images by a factor of 32. To adapt them to small image sizes, we modify the first convolutional layer of ResNet-18 by setting the kernel size to 3 \times 3, the stride to 1, and the padding size to 1 as well. The max-pooling layer is removed. For MobileNetV2, we adjust the stride of the first and fourth convolutional layers to 1. As a result, the modified ResNet-18 and MobileNetV2 down-sample input images by a factor of 8. Therefore, we use the modified ResNet-18 and MobileNetV2 for experiments on CIFAR-10, CIFAR-100, Tiny ImageNet, and ImageNet Dogs. We follow the standard operation in MobileNetV2 [3] and uniformly multiply the width by 0.5 at each layer, which leads to a smaller and faster MobileNetV2 (0.94 M Parameters and 0.02 G FLOPs) than ResNet-18 (11.27 M Parameters and 4.46 G FLOPs). We use the identical network architecture for the baseline network and FocusNet if not otherwise specified.

For all experiments, we use stochastic gradient descent (SGD) with the initial learning rate of 0.1, the momentum of 0.9, and the weight decay of 0.0001. For MNIST, CIFAR-10, CIFAR-100, Imagewoof, Stanford Dogs, and CUB-200, we set the total number of epochs to 200 and divide the learning rate by 10 at epochs 100 and 150. For ImageNet Dogs, Tiny ImageNet, and ImageNet, we set the total number of epochs to 90 and divide the learning rate by 10 at epochs 30 and 60. We do not use data augmentation for MNIST, but standard augmentation for the others, *i.e.* random cropping and flipping.

Pattern Recognition 129 (2022) 108709

Table 3

Classification accuracy (%) on ImageNet. The numbers in parentheses represent the accuracy difference between our FocusNet and the baseline network.

Model	Methods	Top-1	Top-5
ResNet-18	Baseline	69.76	89.08
	FocusNet (Ours)	(+0.66) 70.42	(+0.48) 89.56
ResNet-34	Baseline	73.31	91.42
	FocusNet (Ours)	(+1.04) 74.35	(+0.47) 91.89

Table 4

Model complexity comparison of different network architectures.

Methods	MobileNetV2	ResNet-18	ResNet-34
Params (M)	0.94	11.27	21.38
FLOPs (G)	0.02	4.46	9.29

4.3. Performance evaluation

In the following, we evaluate the performance of FocusNet in terms of classification accuracy, compared with DropMax [8] and a number of state-of-the-art knowledge distillation methods includ-ing [7,17,18,32–34], and [35].

4.3.1. Comparisons with DropMax

We validate our proposed FocusNet on all datasets and compare it with DropMax [8], because they both aim to focus on the confusing classes and ignore the irrelevant ones. For a fair comparison, we use the same network architecture for DropMax, the baseline network, and FocusNet.

Table 2 shows the top-1 validation accuracy. We have two observations. (1) Both DropMax and our FocusNet improve the performances of the baseline network on all datasets, which implies that focusing on confusing classes is an effective strategy for classification. (2) Our FocusNet is more accurate than DropMax, especially on the fine-grained datasets.

Table 3 shows the classification performance of our FocusNet on the ImageNet dataset. We did not compare it with DropMax because our experiments show that FocusNet constantly outperforms DropMax. Our FocusNet using ResNet-18 and ResNet-34 as network architectures consistently achieves better performance, which means that FocusNet scales well on different models.

4.3.2. Comparisons with knowledge distillation

We compare our FocusNet with knowledge distillation [4] because they both use knowledge from another network during training. However, as described in Section 3.6, FocusNet is essentially different from the knowledge distillation.

In our confusion-focusing mechanism, the architectures of the baseline branch and the FocusNet branch can be the same or different. Similarly, knowledge distillation can also use different network combinations. We also evaluate the classification performance in an extreme case where a poorly-trained and small model guides a large model.

We select MobileNetV2 [3], ResNet-18 [1] or ResNet-34 [1] as the network architecture in the following experiments. Table 4 compares their model complexity. We choose three network combinations on most datasets, *i.e.*, "ResNet-18 \rightarrow MobileNetV2",

Table 2

Comparisons of top-1 validation accuracy (%) with baseline and DropMax. The best ones are bolded, and the second-best ones are underlined. For MNIST, the network architecture is LeNet-5 [25]. For the other datasets, the network architecture is ResNet-18 [1].

Methods	MNIST	CIFAR-10	Imagewoof	CIFAR-100	CUB-200	Stanford Dogs	Tiny ImageNet	ImageNet Dogs
Baseline	99.31	92.43	84.43	74.52	50.29	63.32	59.50	70.76
DropMax [8]	<u>99.43</u>	<u>93.70</u>	86.87	74.59	<u>57.68</u>	64.02	<u>60.93</u>	71.34
FocusNet (Ours)	99.56	94.86	86.94	78.29	63.43	71.05	64.49	73.92

Table 5

Comparisons of top-1 validation accuracy (%) with standard knowledge distillation. The best ones are bolded, and the second-best ones are underlined. T and S denote teacher and student models, respectively.

Methods	CIFAR-10	Imagewoof	CIFAR-100	CUB-200	Stanford Dogs	Tiny ImageNet	ImageNet Dogs
T: ResNet-18	92.43	84.43	74.52	50.29	63.32	59.50	70.76
S: MobileNetV2	87.74	81.99	69.06	52.71	60.52	52.77	52.81
Standard KD	88.20	82.76	71.44	58.03	64.73	56.60	64.20
FocusNet (Ours)	90.45	83.56	74.49	64.27	66.71	58.96	65.78

Table 6

Model complexity comparisons with self-knowledge distillation methods. The network architecture is ResNet-18.

Methods	Baseline	DDGSD	CS-KD	SLA-SD	Tf-KD	ONE	BYOT	FRSKD	FocusNet (Ours)
Params (M)	11.27	11.27	11.27	11.58	11.27	11.27	12.49	17.75	11.27
FLOPs (G)	4.46	4.46	4.46	4.46	4.46	4.46	4.68	7.62	4.46

Table 7

 $\label{eq:comparisons of top-1 validation accuracy (\%) with seven self-distillation methods. The best ones are bolded, and the second-best ones are underlined. The network architecture is ResNet-18.$

Methods	CIFAR-100	CUB-200	Stanford Dogs	Tiny ImageNet	ImageNet Dogs
Baseline [1]	74.52	50.29	63.32	59.50	70.76
Tf-KD [7]	76.86	57.23	66.55	60.02	71.64
ONE [35]	76.67	54.71	65.39	62.33	72.56
DDGSD [32]	76.61	58.49	69.00	61.59	71.98
BYOT [17]	76.68	58.66	68.82	63.10	72.90
CS-KD [33]	77.19	64.34	68.91	60.44	72.76
SLA-SD [34]	77.52	56.17	67.30	60.81	72.78
FRSKD\F [18]	77.64	62.29	69.48	63.84	73.53
FRSKD [18]	77.71	65.39	70.77	64.35	74.31
FocusNet (Ours)	78.29	63.43	71.05	64.49	73.92
OurLoss+FRSKD	78.45	67.19	71.49	64.92	74.69

"ResNet-18 \rightarrow ResNet-18", and "a poorly-trained MobileNetV2 \rightarrow ResNet-18", corresponding to standard, self and defectivereverse knowledge distillation, respectively. In addition, we also use "ResNet-34 \rightarrow ResNet-34" for experiments on ImageNet to compare with the state-of-the-art self-knowledge distillation method.

We use ResNet-18 to guide MobileNetV2. Table 5 compares the proposed FocusNet and standard knowledge distillation. Because we only evaluate the student model, comparisons among MobileNetV2, standard knowledge distillation and FocusNet are fair. Top-1 validation accuracy demonstrates that FocusNet is superior to standard knowledge distillation, especially on the fine-grained datasets, CUB-200, for example.

We adopt ResNet-18 to guide ResNet-18, and compare the proposed FocusNet with seven representative self-knowledge distillation methods including three data-augmentation-based methods (DDGSD [32], CS-KD [33], SLA-SD [34]) and four methods using auxiliary networks (Tf-KD [7], ONE [35], BYOT [17], FRSKD [18]). Table 6 compares the model complexity between FocusNet and these self-distillation methods.

Table 7 compares the classification performance of FocusNet with self-knowledge distillation methods. Results show that our FocusNet outperforms state-of-the-art self-knowledge distillation methods on the CIFAR-100, Stanford Dogs, and Tiny ImageNet datasets. In addition, we notice that FRSKD [18] uses an efficient learning paradigm that simultaneously trains the classifier network and an auxiliary self-teacher network, and passes feature maps to each other. We adopt this paradigm and replace its distillation loss function with our proposed focus-picking loss function, called OurLoss+FRSKD. The results show that OurLoss+FRSKD achieves the best performances on all datasets, validating the advantage of our focus-picking loss function.

Table 8 compares our FocusNet with FRSKD on the ImageNet dataset. Results show that our FocusNet outperforms the state-of-the-art method in the top-1 accuracy.

Table 8

Comparisons of validation accuracy (%) with FRSKD on ImageNet. The best ones are bolded, and the second-best ones are underlined.

Model	Methods	Top-1	Top-5
ResNet-18	Baseline	69.76	89.08
	FRSKD	<u>70.17</u>	89.78
	FocusNet (Ours)	70.42	<u>89.56</u>
ResNet-34	Baseline	73.31	91.42
	FRSKD	<u>73.75</u>	92.11
	FocusNet (Ours)	74.35	<u>91.89</u>

We use the poorly-trained MobileNetV2 to guide ResNet-18 and compare the proposed FocusNet with the defective-reverse knowledge distillation. The poorly-trained MobileNetV2 means that we only train it a few epochs and do not decay the learning rate during training. Specifically, we train 1 epoch for CIFAR-10, 30 epochs for Imagewoof, 5 epochs for CIFAR-100, 50 epochs for CUB-200 and Stanford Dogs, and 10 epochs for Tiny ImageNet and ImageNet Dogs. Table 9 shows that knowledge distillation is not suitable for this kind of network combination. However, FocusNet significantly improves the baseline ResNet-18. The differences in performance are mainly due to the difference in design motivations. As described in Section 3.6, the knowledge distillation aims to train an easy-deployed student model to match the teacher's performance. However, FocusNet aims to focus on confusing classes. The credibility of confusing classes mainly depends on the top-k accuracy instead of the top-1 accuracy, so the poorly-trained MobileNetV2 has less influence on our FocusNet than on the knowledge distillation.

4.4. Impacts of the inaccurate baseline network on FocusNet

We investigate the impact of confusing classes credibility on FocusNet by using a poorly-trained baseline network to guide Focus-

Table 9

Comparisons of top-1 validation accuracy (%) with defective-reverse knowledge distillation. The best ones are bolded, and the second-best ones are underlined. Pt-MobileNetV2 and ResNet-18 are the poorly-trained teacher and student models, respectively.

Methods	CIFAR-10	Imagewoof	CIFAR-100	CUB-200	Stanford Dogs	Tiny ImageNet	ImageNet Dogs
T: Pt-MobileNetV2	46.93	41.33	35.65	35.43	37.98	32.14	33.29
S: ResNet-18	<u>92.43</u>	84.43	74.52	<u>50.29</u>	<u>63.32</u>	<u>59.50</u>	<u>70.76</u>
De-reverse KD	55.50	51.69	46.22	42.89	45.50	38.45	39.39
FocusNet (Ours)	93.78	86.13	77.11	62.34	66.88	62.65	72.10

Table 10

Ablation study results. The results are top-1 validation accuracy (%). The network architecture is ResNet-18. The best performing strategies are in bold.

Methods	Cross-entropy	\mathcal{L}_{cls}	$\mathcal{R}_{attention}$	$\mathcal{R}_{entropy}$	CUB-200	Stanford Dogs
Baseline	\checkmark				50.29	63.32
		\checkmark			56.82	64.07
Ablations		\checkmark		\checkmark	61.10	68.76
		\checkmark	\checkmark		61.27	67.21
FocusNet (Ours)		\checkmark	\checkmark	\checkmark	63.43	71.05



Fig. 4. The impact of inaccurate baseline branch on FocusNet. The poorly-trained baseline network denotes that we stop training at epochs 1, 10, 20, 30, 40, 50 and 60. The blue dash line and solid red line both indicate accuracy at epoch 200. The network architecture is ResNet-18. The benchmark datasets are CUB-200 and Stanford Dogs in (a) and (b), respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Net. In the following experiments, we select ResNet-18 as the architecture of both the baseline and FocusNet. The poorly-trained baseline ResNet-18 means that we only train it a few epochs and do not decay the learning rate during training. FocusNet denotes we train ResNet-18 with our focus-picking loss function. Fig. 4 shows the impact of the inaccurate baseline network branch on FocusNet. The performance of FocusNet generally improves as the accuracy of the baseline network increases. Furthermore, even with a highly inaccurate pre-trained network (1.29% for CUB-200 and 1.22% for Stanford Dogs), FocusNet still outperforms the baseline (blue dashed line). We attribute this phenomenon to FocusNet, which emphasizes the contribution of difficult samples on the one hand, and focuses on the most confusing classes on the other. The credibility of confusing classes mainly depends on the top-*k* accuracy rather than the top-1 accuracy.

4.5. Ablation study

In this section, we investigate the necessity of each term of our proposed focus-picking loss function. Table 10 shows the top-1 validation accuracy of ResNet-18 on the CUB-200 and Stanford Dogs datasets using different loss function combinations. \mathcal{L}_{cls} improves the baseline by emphasizing the importance of difficult samples. $\mathcal{R}_{attention}$ and $\mathcal{R}_{entropy}$ both contribute to better performances by encouraging FocusNet to focus on confusing classes and prevent over-fitting. Our FocusNet outperforms the baseline by a large margin, demonstrating the importance of each term of our focus-picking loss function.

5. Conclusions

In this work, we observe that a classification network trained with one-hot labels can pay attention to the target image region even if it makes an incorrect prediction. To enable the network to better discriminate between confusing classes, we propose a confusion-focusing mechanism and a focus-picking loss function by considering confusing inter-class interactions. The confusionfocusing mechanism is implemented in a two-branch network architecture, i.e. a baseline branch and a FocusNet branch. The Focus-Net branch does not restrict to specific network architectures and can adopt common models. Its main advantage is that it improves the performance of networks without increasing model complexity. Fine-grained classification and computational resource-limited scenes can benefit from our method because our method encourages networks to focus on the most confusing classes and does not impose additional deployment burdens at test time. In addition, we have demonstrated that FocusNet can still improve performance even with a highly-inaccurate pre-trained network. Experimental results have validated that our FocusNet performs well on a number of challenging datasets, and that the focus-picking loss function can further improve the state-of-the-art techniques.

A limitation of our current confusion-focusing mechanism is its ineffective two-stage training strategy. In the future we will work towards an optimized one-stage scheme that can further improve training efficiency and classification accuracy.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
- [2] J. Howard, Imagewoof, https://www.github.com/fastai/imagenette.
- [3] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, L.-C. Chen, MobileNetV2: Inverted residuals and linear bottlenecks, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018.
- [4] G. Hinton, O. Vinyals, J. Dean, Distilling the knowledge in a neural network, Advances in Neural Information Processing Systems Workshop, 2015.
- [5] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, Rethinking the inception architecture for computer vision, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 2818–2826.
- [6] R. Müller, S. Kornblith, G.E. Hinton, When does label smoothing help? Advances in Neural Information Processing Systems, 2019.
- [7] L. Yuan, F.E. Tay, G. Li, T. Wang, J. Feng, Revisiting knowledge distillation via label smoothing regularization, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2020, pp. 3903–3911.
- [8] H.B. Lee, J. Lee, S. Kim, E. Yang, S.J. Hwang, DropMax: Adaptive variational softmax, Advances in Neural Information Processing Systems, 2018.
- [9] S. Xie, R. Girshick, P. Dollár, Z. Tu, K. He, Aggregated residual transformations for deep neural networks, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 1492–1500.
- [10] S.-H. Gao, M.-M. Cheng, K. Zhao, X.-Y. Zhang, M.-H. Yang, P. Torr, Res2net: a new multi-scale backbone architecture, IEEE Trans Pattern Anal Mach Intell 43 (2) (2019) 652–662.
- [11] A.M. Obeso, J. Benois-Pineau, M.S.G. Vázquez, A.Á.R. Acosta, Visual vs internal attention mechanisms in deep neural networks for image classification and object detection, Pattern Recognit 123 (2022) 108411.
- [12] Y. Chen, H. Wang, X. Sun, B. Fan, C. Tang, H. Zeng, Deep attention aware feature learning for person re-identification, Pattern Recognit 126 (2022) 108567.
- [13] G. Pereyra, G. Tucker, J. Chorowski, Ł. Kaiser, G. Hinton, Regularizing neural networks by penalizing confident output distributions, International Conference on Learning Representations Workshop (2017).
- [14] S. Liu, X. Song, Z. Ma, E.D. Ganaa, X. Shen, MoRE: multi-output residual embedding for multi-label classification, Pattern Recognit (2022) 108584.
- [15] Y. Liu, X. Gao, Q. Gao, J. Han, L. Shao, Relation-based discriminative cooperation network for zero-shot classification, Pattern Recognit 118 (2021) 108024.
- [16] C. Shi, L. Fang, Z. Lv, M. Zhao, Explainable scale distillation for hyperspectral image classification, Pattern Recognit 122 (2022) 108316.
- [17] L. Zhang, J. Song, A. Gao, J. Chen, C. Bao, K. Ma, Be your own teacher: Improve the performance of convolutional neural networks via self distillation, in: Proceedings of the IEEE International Conference on Computer Vision, 2019.
- [18] M. Ji, S. Shin, S. Hwang, G. Park, I.-C. Moon, Refine myself by teaching myself: Feature refinement via self-knowledge distillation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2021.

- [19] Y. Chen, Z. Zhang, Y. Wang, Y. Zhang, R. Feng, T. Zhang, W. Fan, AE-NEt: finegrained sketch-based image retrieval via attention-enhanced network, Pattern Recognit 122 (2022) 108291.
- [20] H. Zheng, J. Fu, Z.-J. Zha, J. Luo, T. Mei, Learning rich part hierarchies with progressive attention networks for fine-grained image recognition, IEEE Trans. Image Process. 29 (2019) 476–488.
- [21] W. Chen, Y. Liu, E.M. Bakker, M.S. Lew, Integrating information theory and adversarial learning for cross-modal retrieval, Pattern Recognit 117 (2021) 107983.
- [22] T.-Y. Lin, P. Goyal, R. Girshick, K. He, P. Dollár, Focal loss for dense object detection, in: Proceedings of the IEEE International Conference on Computer Vision, 2017.
- [23] Y. Liu, X. Gao, J. Han, L. Liu, L. Shao, Zero-shot learning via a specific rank-controlled semantic autoencoder, Pattern Recognit 122 (2022) 108237.
- [24] C. Biswas, D. Ganguly, P.S. Mukherjee, U. Bhattacharya, Y. Hou, Privacy-aware supervised classification: an informative subspace based multi-objective approach, Pattern Recognit 122 (2022) 108301.
- [25] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, Proc. IEEE 86 (11) (1998) 2278–2324.
- [26] A. Krizhevsky, G. Hinton, et al., Learning multiple layers of features from tiny images, Technical Report (2009).
- [27] O. Vinyals, C. Blundell, T. Lillicrap, k. kavukcuoglu, D. Wierstra, Matching networks for one shot learning, Advances in Neural Information Processing Systems, 2016.
- [28] A. Khosla, N. Jayadevaprakash, B. Yao, F.-F. Li, Novel dataset for fine-grained image categorization: Stanford dogs, in: IEEE Conference on Computer Vision and Pattern Recognition Workshop, 2011.
- [29] C. Wah, S. Branson, P. Welinder, P. Perona, S. Belongie, The Caltech-UCSD Birds-200-2011 dataset, Technical Report, 2011.
- [30] Y. Le, X. Yang, Tiny ImageNet visual recognition challenge, CS 231N (2015).
- [31] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al., ImageNet large scale visual recognition challenge, Int J Comput Vis 115 (3) (2015) 211–252.
- [32] T.-B. Xu, C.-L. Liu, Data-distortion guided self-distillation for deep neural networks, in: Proceedings of the AAAI Conference on Artificial Intelligence, 2019.
- [33] S. Yun, J. Park, K. Lee, J. Shin, Regularizing class-wise predictions via selfknowledge distillation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2020.
- [34] H. Lee, S.J. Hwang, J. Shin, Self-supervised label augmentation via input transformations, in: International Conference on Machine Learning, 2020.
- [35] X. Lan, X. Zhu, S. Gong, Knowledge distillation by on-the-fly native ensemble, Advances in Neural Information Processing Systems, 2018.

Xue Zhang received the B.S. and M.S. degrees from Shandong Jianzhu University and Shandong University, China, in 2016 and 2019, respectively. He is currently pursuing the Ph.D. degree with the College of Information Science and Electronic Engineering, Zhejiang University. His research interests include image classification and material recognition.

Zehua Sheng received his B.E. degree from Zhejiang University in 2017. Currently, he is studying for his Ph.D degree in Zhejiang University, China. His current working interests are image denoising and multi-modal image processing.

Hui-Liang Shen received the B.Eng. and Ph.D. degrees in electronic engineering from Zhejiang University, China, in 1996 and 2002, respectively. He is currently a Full Professor with the College of Information Science and Electronic Engineering, Zhejiang University. His research interests include multispectral imaging, image processing, computer vision, and machine learning.