

CRSSC: Salvage Reusable Samples from Noisy Data for Robust Learning

Zeren Sun

Nanjing University of Science and
Technology, Nanjing, China
zerens@njust.edu.cn

Xiu-Shen Wei

Nanjing University of Science and
Technology, Nanjing, China
weix.gm@gmail.com

Xian-Sheng Hua

Alibaba Group, Hangzhou, China
huaxiansheng@gmail.com

Guosheng Hu[†]

AnyVision, Belfast, United Kingdom
huguosheng100@gmail.com

Yazhou Yao^{*}

Nanjing University of Science and
Technology, Nanjing, China
yazhou.yao@njust.edu.cn

Jian Zhang

University of Technology Sydney,
Sydney, Australia
jian.zhang@uts.edu.au

ABSTRACT

Due to the existence of label noise in web images and the high memorization capacity of deep neural networks, training deep fine-grained (FG) models directly through web images tends to have an inferior recognition ability. In the literature, to alleviate this issue, loss correction methods try to estimate the noise transition matrix, but the inevitable false correction would cause severe accumulated errors. Sample selection methods identify clean (“easy”) samples based on the fact that small losses can alleviate the accumulated errors. However, “hard” and mislabeled examples that can both boost the robustness of FG models are also dropped. To this end, we propose a certainty-based reusable sample selection and correction approach, termed as CRSSC, for coping with label noise in training deep FG models with web images. Our key idea is to additionally identify and correct reusable samples, and then leverage them together with clean examples to update the networks. We demonstrate the superiority of the proposed approach from both theoretical and experimental perspectives. The source code, models, and data have been made available at <https://github.com/NUST-Machine-Intelligence-Laboratory/CRSSC>.

CCS CONCEPTS

• Computing methodologies → Computer vision.

KEYWORDS

label noise; noisy data; sample selection; robust learning

ACM Reference Format:

Zeren Sun, Xian-Sheng Hua, Yazhou Yao, Xiu-Shen Wei, Guosheng Hu, and Jian Zhang. 2020. CRSSC: Salvage Reusable Samples from Noisy Data for Robust Learning. In *Proceedings of the 28th ACM International Conference*

^{*}Co-corresponding author

[†]Co-corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

MM '20, October 12–16, 2020, Seattle, WA, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7988-5/20/10...\$15.00

<https://doi.org/10.1145/3394171.3413978>

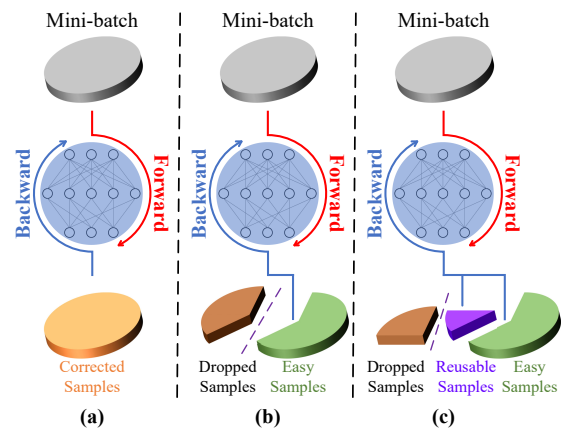


Figure 1: (a) Loss correction methods correct losses for all samples in each mini-batch before back-propagating. (b) Sample selection methods identify easy samples out of each mini-batch and then back-propagate using only these easy samples. (c) Our CRSSC is based on sample selection methods, but proposes to leverage additionally reusable samples (including “hard” and mislabeled ones) for boosting the model learning.

on Multimedia (MM '20), October 12–16, 2020, Seattle, WA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3394171.3413978>

1 INTRODUCTION

In the past few years, labeled image datasets have played a critical role in computer vision tasks [5, 28, 29, 37, 42]. To distinguish the subtle differences among fine-grained categories (e.g., birds [44], airplanes [30], or plants [43]), a large amount of well-labeled images are typically required. However, labeling objects at the subordinate level generally requires domain-specific expert knowledge, which is not always available for a human annotator from crowd-sourcing platforms like Amazon Mechanical Turk [46, 47].

To reduce the cost of fine-grained annotation, many methods have been proposed which mainly focus on a semi-supervised learning paradigm [8, 32, 48, 54]. For example, Xu *et al.* [48] proposed to utilize detailed annotations and transfer as much knowledge as possible from existing strongly supervised datasets to weakly supervised web images for fine-grained recognition. Niu *et al.* [32] proposed a new learning scenario which only required experts to label a few fine-grained subcategories and can predict all the

remaining subcategories by virtue of web data. Nevertheless, semi-supervised methods involve various forms of human intervention and have relatively limited scalability.

To further reduce the demand of manual annotation, leveraging web images to train FGVC models has attracted broad attention [51, 52, 55, 58]. Owing to the error-prone automatic tagging system or non-expert annotation, web images for fine-grained categories are usually associated with massive label noise. Therefore, along with how to find discriminative regions in fine-grained images, how to handle label noise is another pivotal problem for training deep FGVC models with web images, which is also the focus of this paper. Statistical learning has contributed significantly to cope with label noise, especially in theoretical aspects [53]. However, in this work, we mainly focus on deep learning based methods.

One typical method, as illustrated in Fig. 1 (a), is using “loss correction” to correct the loss of training samples based on the estimated noise transition matrix [33, 34]. The problem is that it is extremely difficult to get an accurate estimation of the noise transition matrix, thus inevitable false correction will lead to severe accumulated errors in the training process. Alternatively, as described in Fig. 1 (b), another popular training schema endeavors to adopt “sample selection” to identify and remove samples with label noise, and only use clean samples to update the networks [14, 18, 31]. Despite promising results that have been achieved in these approaches, the loss-based selection strategy favors “easy” examples. “Hard” and mislabeled samples (e.g., a “Laysan Albatross” image is labeled as “Black Footed Albatross”) are ignored although they are surprisingly beneficial in making FG models more robust.

As shown in Fig. 1 (c), our idea is to re-utilize informative images (“hard” and mislabeled examples) by selecting and correcting employable instances from high-loss samples. To be specific, we first split samples via the loss-based criterion. Low-loss instances are deemed to be clean, “easy” examples and their labels remain unaltered. Then, we perform a further partition on high-loss samples. In this separation process, we follow a simple but intuitive observation: for images of irrelevant categories, since they do not belong to any categories involved in the task, the network tends to be more confused when predicting their label probabilities. On the contrary, “hard” and mislabeled ones lean to obtain a relatively more certain prediction. After hard samples and mislabeled ones are selected out of high-loss instances, we manage to correct their labels and then feed them together with clean samples into the network for updating parameters. Extensive experiments and ablation studies on tasks of fine-grained image categorization demonstrate the superiority of our proposed approach over existing webly supervised state-of-the-art methods. The primary contributions of this work can be summarized as follows:

(1) A webly supervised deep model CRSSC is proposed to bridge the gap between FGVC tasks and numerous web images. Comprehensive experiments demonstrate that our approach outperforms existing state-of-the-art webly supervised methods by a large margin.

(2) Three types of FG web images (*i.e.*, clean, reusable, and irrelevant) which inherently exist in collected web images are successfully identified and then separated by CRSSC. Compared with existing methods, our approach can further leverage reusable samples to boost the model learning.

(3) A novel label correction method which utilizes the prediction history of the network is proposed to re-label the reusable samples. Our experiments show that our label correction method is better than the existing one which uses the current epoch prediction results.

2 RELATED WORK

Fine-grained Visual Classification Fine-grained visual classification (FGVC) aims to distinguish similar subcategories belonging to the same basic category. Generally, existing approaches can be roughly grouped into three categories: 1) strongly supervised methods, 2) weakly supervised methods, and 3) semi-supervised methods. Strongly supervised methods tend to require not only image-level labels but also manually annotated bounding boxes or part annotations [2, 17, 45]. Different from strongly supervised methods, weakly supervised methods cease to use bounding boxes and part annotations. Instead, methods in this group only require image-level labels during training [6, 9–12, 19, 24, 26, 27, 63]. The third group involves leveraging web images in training the FGVC model [8, 32, 48]. However, these approaches still contain a certain level of human intervention, making them not purely web-supervised.

Webly Supervised Learning Since learning directly from web images requires no human annotation, this learning scenario is becoming popular [16, 40, 52, 59–62]. However, training deep FGVC models directly with web images usually leads to poor performance due to the existence of label noise and memorization effects [1, 57] of neural networks. Existing deep methods for overcoming label noise can be categorized into two sets [39]: 1) loss correction and 2) sample selection. Loss correction methods choose to correct the loss of training samples based on an estimated noise transition matrix [3, 13, 33–35, 56]. However, due to the difficulty in accurately estimating the noise transition matrix, accumulated errors induced by false correction are inevitable [14, 18]. Sample selection methods identify clean samples out of mini-batches based on their losses and leverage them to update the network [14, 18, 22, 31, 39]. Nevertheless, loss-based selection methods would cause the domination of easy samples in the training procedure while hard ones get ignored substantially [3, 25, 36, 39]. Our proposed CRSSC can leverage additional reusable samples to boost the deep FGVC models.

3 THE PROPOSED APPROACH

3.1 Overview

Fig. 2 presents the architecture of our proposed model. Generally, we can train a deep FGVC model through a well-labeled dataset $\mathcal{D} = \{(x_i, y_i) | 1 \leq i \leq M_L\}$, in which x_i is the i -th training sample and y_i is the corresponding ground-truth label. In the conventional training schema, the model parameters are updated by optimizing a cross-entropy loss L_{CE} as follows:

$$\theta_{t+1} = \theta_t - \alpha \nabla \left[\frac{1}{|\mathcal{D}|} \sum_{x_i \in \mathcal{D}} L_{CE}(f(x_i; \theta_t), y_i) \right], \quad (1)$$

where θ_t is the network parameter in the t -th training epoch and $f(x_i; \theta_t)$ is the network output of sample x_i .

However, for web images $\mathcal{D}_W = \{(x_i, \hat{y}_i) | 1 \leq i \leq M_W\}$, reliable labels are not always available and they are usually associated with

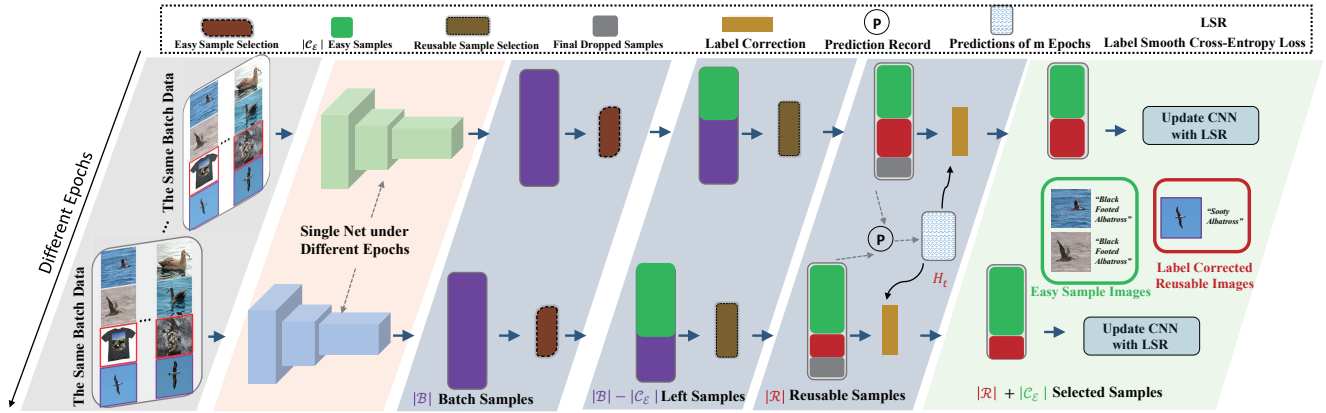


Figure 2: The architecture of our proposed model. The deep network feeds forward a mini-batch of web images to predict labels and calculate losses for each image. Then the drop module sorts and selects small-loss instances as the easy sample set C_E . Furthermore, we take a reuse module to leverage part of the high-loss instances. Specifically, the reuse module calculates the prediction certainty for each high-loss instance and selects high-certainty instances as reusable set \mathcal{R} . To use samples in \mathcal{R} , our model utilizes predictions of previous epochs to re-label reusable samples. Finally, training instances in $C_E \cup \mathcal{R}$ are leveraged together to promote the optimization of model.

label noise. Then we can divide \mathcal{D}_W into three subsets:

$$\mathcal{D}_W = \mathcal{N} \cup \mathcal{M} \cup \mathcal{C}, \quad (2)$$

where \mathcal{N} indicates noisy samples, \mathcal{M} represents mislabeled ones, and \mathcal{C} stands for the clean set. More specifically, \mathcal{C} can be further separated into easy example set C_E and hard example set C_H :

$$\mathcal{C} = C_E \cup C_H. \quad (3)$$

It should be noted that, as the training proceeds, hard samples will gradually become “easy”. Thus, the split of the clean set \mathcal{C} changes in the training process. In this work, we aim to train a robust deep FGVC model through noisy web data. Our main idea is to properly select and then re-label informative training samples for boosting the robustness of the FGVC model. Based on the division described in Eq. (2) and (3), we regard the union of \mathcal{C} and \mathcal{M} as informative training set $\mathcal{I} = \mathcal{C} \cup \mathcal{M}$, though samples in \mathcal{M} have to be corrected before being fed into the model for further network optimization.

3.2 Drop and Reuse

To optimize the network with only useful knowledge from the web images, we have to 1) eliminate the negative influence from samples which belong to \mathcal{N} , 2) reduce the misleading impact of samples belonging to \mathcal{M} . Therefore, two key challenges of tackling label noise in web images \mathcal{D}_W are: 1) how to select samples which belong to \mathcal{I} and prevent the network from learning irrelevant samples, and 2) how to correct labels of mislabeled ones and reuse them as part of the informative knowledge.

Memorization effects [1, 57] indicate that, on noisy datasets, CNNs tend to first learn clean and easy patterns in initial epochs. As the number of epochs increases, CNNs will eventually overfit on noisy samples. Our key idea is to drop these noisy instances before they are memorized. A widely used sample selection strategy is to separate instances based on losses [14, 18]. These methods typically select a human-defined proportion $(1 - \tau) \times 100\%$ of low-loss instances as clean samples and directly drop the rest ones. Although significant improvements have been achieved in these works for dealing with label noise, the way they separate instances would

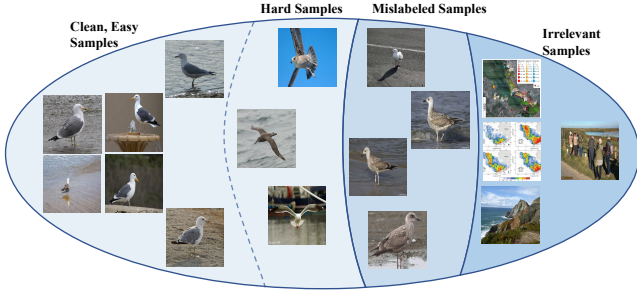
lead to the mistaken deletion of samples belonging to C_H as hard examples also tend to produce high losses. Moreover, mislabeled samples would also be dropped by these sample selection methods due to their high losses.

To tackle drawbacks of these loss-based sample selection methods, we design a drop and reuse mechanism to further select useful instances from high-loss samples. Furthermore, for the purpose of avoiding involving the human-defined noise rate τ , we modify the conventional loss-based sample selection as in Definition 3.1. Through adopting our proposed loss-based drop module as well as the reuse module, we can effectively avoid mistakenly deleting hard examples and can also make full use of mislabeled samples.

Definition 3.1. In a mini-batch \mathcal{B} , a sample $x \in \mathcal{B}$ belongs to C_E only if its loss $L_{CE}(x, \hat{y}; \theta) < \frac{1}{|\mathcal{B}|} \sum_{x_i \in \mathcal{B}} L_{CE}(x_i, \hat{y}_i; \theta)$

We use the average loss as the selection threshold for dynamically separating informative samples from irrelevant ones. Specifically, due to limited robustness in initial epochs, more samples tend to have high losses. The selection threshold will be large and CNNs will learn easy patterns from as many samples as possible. As the training proceeds, CNNs gain more robust ability and more samples tend to have low losses. The selection threshold will be small and more samples will be dropped. In this situation, CNNs will discard as many samples as possible for ensuring the data learned by CNN is informative. In this way, our method can reduce the negative impact of error accumulation.

Why can we distinguish reusable samples from irrelevant ones? Intuitively, the predicted label probability of a reusable sample has a completely different pattern from that of an instance which belongs to an irrelevant category. Samples belonging to \mathcal{N} can be distinguished based on the confusion of prediction. For example, as shown in Fig. 3, in a bird classification task, when we feed a bird image (e.g., hard or mislabeled samples) into the model, the network tends to produce a certain prediction although it may have a high loss due to incorrect labeling. However, if we feed an irrelevant sample (e.g., a bird distribution map), which apparently belongs to



Collected Web Images Using Query “California Gull”

Figure 3: The easy, hard, mislabeled, and irrelevant web images collected through query “California Gull”, in which hard examples and mislabeled ones can be leveraged to boost the robustness of the deep FGVC model.

\mathcal{N} in this task, the network would get confused and thus produce an uncertain prediction. Inspired by this observation, we formalize a new criterion to further select reusable samples from the ones with high-loss:

Definition 3.2. In a mini-batch \mathcal{B} , a sample $x \in \mathcal{B} \cap \mathcal{C}_{\mathcal{E}}^{-1}$ is reusable if its prediction certainty $V(x; \theta)$ satisfies the condition: $V(x; \theta) \geq \frac{1}{|\mathcal{B} \cap \mathcal{C}_{\mathcal{E}}|} \sum_{x_i \in \mathcal{B} \cap \mathcal{C}_{\mathcal{E}}} V(x_i; \theta)$.

[4] demonstrated that the prediction variance can be used to measure the uncertainty of each sample in classification tasks. Therefore, in order to quantify the certainty of prediction, we simply adopt the standard deviation (the square root of the variance) of predicted probabilities for sample x :

$$V(x; \theta) = \sqrt{\frac{1}{K} \sum_{j=1}^K (P_j(x; \theta) - \mu)^2}, \quad (4)$$

where K is the number of categories, $P_j(x; \theta)$ is the softmax result of $f(x; \theta_t)$, which is deemed as the (pseudo) predicted probability of sample x belonging to the j -th category, and μ is the mean value of predicted probabilities. Since $\sum_{j=1}^K P_j(x; \theta) = 1$, we can easily rewrite Eq. (4) as:

$$V(x; \theta) = \sqrt{\frac{1}{K} \sum_{j=1}^K P_j(x; \theta)^2 - \frac{1}{K^2}}. \quad (5)$$

The standard deviation of predicted probabilities is consistent with prediction certainty. From Eq. (5) and the constraint of $\sum_{j=1}^K P_j(x; \theta) = 1$, we can have the following observations. 1) The standard deviation of predicted probabilities $V(x; \theta)$ is bounded. 2) $V(x; \theta)$ gets larger when one label’s probability (e.g., $P_q(x, \theta)$) gets notably higher. 3) If $P_q(x, \theta)$ is significantly higher than others’ value, the network is more certain about its prediction on this sample. $V(x; \theta)$ reaches its maximum value $\frac{\sqrt{K-1}}{K}$ when

$$\exists q \in \{1, 2, 3, \dots, K\}, P_j(x, \theta) = \begin{cases} 1, & j = q \\ 0, & j \neq q \end{cases}. \quad (6)$$

In this case, the prediction certainty also reaches its maximum. On the other hand, $V(x; \theta)$ gets smaller when all labels’ probabilities get closer to each other. In this case, the prediction certainty also

drops. This observation is consistent with the intuition. It results from a fact that if the network produces a prediction in which each label’s probability is close to others, it means the network is highly confused about this sample. $V(x; \theta)$ reaches its minimum value 0 when

$$\forall j \in \{1, 2, 3, \dots, K\}, P_j(x, \theta) = 1/K. \quad (7)$$

In this situation, the model fails as all labels are predicted equally.

3.3 Label Correction

The reusable samples selected by the certainty-based criterion defined in Definition 3.2 are assembled into a sample set \mathcal{R} . To be specific, \mathcal{R} includes two types of images: 1) mislabeled instances and 2) hard examples that have correct labels.

In order to leverage these informative samples for training, their noisy labels have to be corrected. [49] proposed to use the current prediction to replace original labels. Nevertheless, due to a lack of robustness for CNN predictions on noisy datasets, using a single prediction to relabel mislabeled instances may result in error accumulation. Therefore, different from [49], we propose to correct noisy labels using the prediction history. This results from a fact that averaging distributions over classifier iterations can increase stability and reduce the influence of misleading predictions. Empirical results show that our label correction method works better.

Specifically, we record the label prediction as well as its corresponding predicted probability for each training sample $x \in \mathcal{D}_{\mathcal{W}}$. A history list $H_t(x)$ is defined as follows:

$$H_t(x) = \{(h_{t-i}(x), p_{t-i}(x)) | 1 \leq i \leq m, x \in \mathcal{D}_{\mathcal{W}}\}, \quad (8)$$

in which $h_t(x)$ is the label prediction of the sample x in the t -th epoch, and $p_t(x)$ is its corresponding predicted probability. The history list $H_t(x)$ is maintained to memorize each sample’s prediction of the previous m epochs.

When calculating forward losses, we correct samples belonging to \mathcal{R} by replacing their original labels with corrected ones defined in Definition 3.3. For samples belonging to $\mathcal{C}_{\mathcal{E}}$, we directly use their original labels. The leftover samples are regarded as irrelevant data and are excluded for robust training.

Definition 3.3. For $x \in \mathcal{R}$, its corrected label y^{corr} is the label prediction who has the highest accumulated probability in the previous m epochs:

$$y^{corr} = \arg \max_{1 \leq j \leq K} \sum_{i=1, h_{t-i}(x)=j}^m p_{t-i}(x). \quad (9)$$

Why don’t we need to separate mislabeled and hard examples? As stated above, both mislabeled samples and hard ones are included in $\mathcal{R} = \mathcal{C}_{\mathcal{H}} \cup \mathcal{M}$. It is difficult to make a reliable split on \mathcal{R} to distinguish hard samples and mislabeled ones. As a result, the label correction might also relabel hard samples using their previous predictions. However, both hard examples and mislabeled ones tend to produce consistent predictions on their true labels, though their predicted probabilities might be lower than those of easy samples. Therefore, using prediction history to relabel hard samples would not compromise the model.

Algorithm 1: The proposed CRSSC algorithm

Input: Initialized network f , warm-up epochs T_w , maximum epoch T_{\max} , and training set \mathcal{D}_W .

```
for  $T = 1, 2, \dots, T_{\max}$  do
  Randomly draw a mini-batch  $\mathcal{B}$  from  $\mathcal{D}_W$ .
  if  $T_w \leq T \leq T_{\max}$  then
    Construct  $\mathcal{B}_{C_E} = \mathcal{B} \cap C_E$  through Definition 3.1.
    Construct  $\mathcal{B}_R = \mathcal{B} \cap \mathcal{R}$  by Definition 3.2.
    Relabel samples in  $\mathcal{B}_R$  with Definition 3.3.
    Update the network  $f$  with Eq. (10).
  else
    Update the network  $f$  according to Eq. (1).
  end
end
Output: Updated network  $f$ .
```

3.4 Summary of CRSSC

Based on the Definition 3.2 and 3.3, we can reformulate the network update function in Eq. (1) as

$$\theta_{t+1} = \theta_t - \alpha \nabla \left[\frac{1}{|\mathcal{C}_E \cup \mathcal{R}|} \sum_{x_i \in \mathcal{C}_E \cup \mathcal{R}} L_{CE}(f(x_i; \theta_t), y_i) \right], \quad (10)$$

where

$$y_i^* = \begin{cases} \hat{y}_i, & x_i \in \mathcal{C}_E; \\ y_i^{corr}, & x_i \in \mathcal{R}. \end{cases} \quad (11)$$

Here, to further enhance the generalization performance of CRSSC, we adopt the Label Smoothing Regularization [41] when calculating the cross-entropy loss. That is to say, for input image x_i , we adopt the following smoothed ground-truth probability $\{q_j | 1 \leq j \leq K\}$ in the loss calculation:

$$q_j = \begin{cases} 1 - \epsilon, & j = y_i^* \\ \epsilon / (K - 1), & j \neq y_i^* \end{cases}. \quad (12)$$

For CRSSC, we first train the network on whole training set in a conventional manner to warm up the network. The reason is that deep CNN has memorization effects [1, 57] and will learn clean and easy patterns in the initial epochs. With this warm-up step, CNN will get equipped with an initial learning capacity. Then, we perform a two-step sample selection in each mini-batch \mathcal{B} : 1) select low-loss instances from \mathcal{B} using the criterion defined in the Definition 3.1, and 2) identify reusable samples, which have high prediction certainty in high-loss instances based on the Definition 3.2. Subsequently, reusable samples are relabeled based on the Definition 3.3. Finally, parameters of the network are updated using the clean, easy sample set \mathcal{B}_{C_E} along with the reusable sample set \mathcal{B}_R . The detailed process of our proposed CRSSC is shown in Algorithm 1.

4 EXPERIMENTS

4.1 Datasets and Evaluation Metrics

Datasets: We evaluate our approach on three popular benchmark fine-grained datasets: CUB200-2011 [44], FGVC-Aircraft [30] and Stanford Cars [20].

Evaluation Metric: Average Classification Accuracy (ACA) is taken as the default evaluation metric.

4.2 Implementation Details

Data preparation: To collect the web training set, we follow [32] and retrieve images from image search engine using the category labels in benchmark datasets. For ensuring no overlaps between the training and testing set, we additionally perform a PCA near-duplicate removal [64] between collected web images and test images in the benchmark datasets. Finally, we regard filtered web images (13503 for FGVC-Aircraft, 18388 for CUB200-2011, and 21448 for Stanford Cars) as the training set and adopt testing images from original benchmark datasets.

CRSSC learning: We use a pre-trained model (e.g., VGG-16 [38]) as the basic CNN network. The number of warm-up epochs T_w is tested in $\{5, 8, 10, 15, 20\}$. The history list length limit l is selected from $\{0, 5, 10, 15, 20\}$. The LSR parameter ϵ is chosen from $\{0.0, 0.2, 0.5, 0.6, 0.8\}$. During the network optimization, we adopt a SGD optimizer with momentum = 0.9. The learning rate, batch size, and weight decay are set to be 0.01, 32, and 0.0003, respectively.

4.3 Baseline Methods

Strongly supervised methods require bounding boxes or part annotations during training. This set of baselines includes Part-Stacked [17], Coarse-to-fine [50], HSnet [23], and Mask-CNN [45]. **Weakly supervised methods** require image-level labels, including Parts Model [12], iSQRT-COV [24], TASN [63], and DCL [6]. **Semi-supervised methods** leverage web images but remain involving human intervention, including Cui *et al.* [8], Xu *et al.* [48], Niu *et al.* [32], and Cui *et al.* [7]. For strongly, weakly, and semi-supervised methods, we report the performances in their papers. **Webly supervised methods** directly leverage web images without human involvement, including VGG-16 [38], ResNet-50 [15], B-CNN [27], Decoupling [31] and Co-teaching [14]. To be fair, we use the same backbone B-CNN [27] in Decoupling, Co-teaching and our CRSSC. For basic networks VGG-16, ResNet-50, and B-CNN, we fine-tune them with noisy web images.

4.4 Experimental Results

Table 1 presents the comparison of ACA results on three benchmark datasets. It should be noted that the results of webly methods are all produced from experiments using exactly the same training data. By observing Table 1, we can notice that our approach performs better than other webly supervised methods on all three benchmark datasets. Compared with basic networks VGG-16, ResNet-50, and B-CNN, our CRSSC (with backbone B-CNN) can effectively alleviate the influence of label noise in the process of model training. Compared with state-of-the-art webly supervised methods Decoupling and Co-teaching, our approach can additionally identify reusable samples and salvage them by performing a label correction. Thus, our CRSSC can efficiently explore more useful samples to boost the robustness of the FGVC model.

5 ABLATION STUDIES

5.1 Training Loss and Prediction Certainty

The prediction loss and certainty are two fundamental criteria for selecting informative samples. To investigate the distribution of prediction loss and certainty for clean, reusable, and dropped

Table 1: The ACA (%) results on three benchmark datasets. BBox/Part means bounding box or part annotation is required during training. Training set means the training data is manually labeled (anno.) or collected from the web (web). iNat refers to the iNat2017 dataset.

Types	Methods	BBox/Part	Training Set	Performance		
				FGVC-Aircraft [30]	CUB200-2011 [44]	Stanford Cars [20]
Strongly	Part-Stacked [17]	✓	anno.	-	76.6	-
	Coarse-to-fine [50]	✓	anno.	87.7	82.9	-
	HSnet [23]	✓	anno.	-	87.5	93.9
	Mask-CNN [45]	✓	anno.	-	85.7	-
Weakly	iSQRT-COV [24]		anno.	91.4	88.7	93.3
	Parts Model [12]		anno.	-	90.4	-
	TASN [63]		anno.	-	89.1	93.8
	DCL [6]		anno.	93.0	87.8	94.5
Semi	Cui <i>et al.</i> [8]		anno.+web	-	89.7	-
	Xu <i>et al.</i> [48]		anno.+web	-	84.6	-
	Niu <i>et al.</i> [32]		anno.+web	-	76.5	-
	Cui <i>et al.</i> [7]		anno.+iNat	90.7	89.3	93.5
Webly	VGG-16 [38]		web	68.4	66.3	61.6
	ResNet-50 [15]		web	60.4	64.4	60.6
	B-CNN [27]		web	64.3	66.6	67.4
	Decoupling [31]		web	75.9	70.6	75.0
	Co-teaching [14]		web	72.8	73.9	73.1
	CRSSC		web	76.5	77.4	76.6

Table 2: The ACA (%) comparison by using different backbones.

Backbone	Method	Aircraft	CUB200	Cars
VGG-16	Standard	68.4	66.3	61.6
	CRSSC	78.3	77.8	81.6
ResNet-18	Standard	53.7	59.3	55.6
	CRSSC	75.6	76.8	84.2
ResNet-50	Standard	60.8	64.4	60.6
	CRSSC	82.5	81.3	87.7

Table 3: The ACA (%) comparison by using different steps.

Model	ACA (%)
ResNet-18	59.3
+ Def. 3.1	69.7
+ Def. 3.1 + Def. 3.2	69.9
+ Def. 3.1 + Def. 3.2 + Def. 3.3	73.4
+ Def. 3.1 + Def. 3.2 + Def. 3.3 + LSR (CRSSC)	75.6
+ Fine-tuned CRSSC	76.8

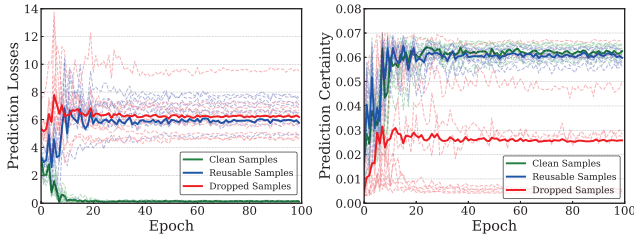


Figure 4: The prediction loss (left) and certainty (right) of clean, reusable, and dropped samples as training progresses. The value on the individual image is plotted in the dotted line and the average value of each group is plotted in the solid line.

samples in training process, we select 30 instances in total, 10 images for each group (clean, reused, and dropped) and plot their prediction losses as well as prediction certainties. The experimental results are shown in Fig. 4.

By observing Fig. 4, we can find that as the network training forwards, the losses of clean samples decrease sharply while their prediction certainties increase steadily. Regarding reusable samples, although some of them have a fairly higher loss than clean ones, their prediction certainties increase remarkably as training progresses. The explanation is that reusable samples are either hard

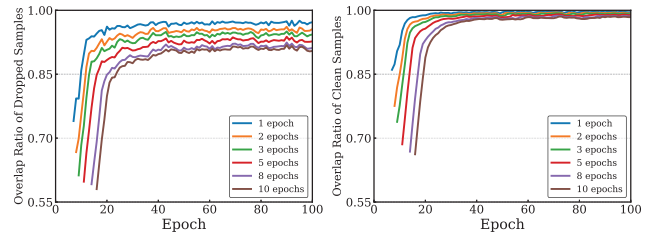


Figure 5: The sample selection overlap of dropped samples (left) and clean ones (right) between each epoch and its previous 1, 2, 3, 5, 8, 10 epochs. The comparison starts after the warm-up epochs, i.e., the blue line starts from the 7th epoch.

or mislabeled instances, thus tend to produce confident predictions consistently. The high loss and low prediction certainty of dropped samples demonstrate that our CRSSC can successfully identify and drop these irrelevant samples.

5.2 Overlap of Identified Samples in Epochs

To further investigate the robustness of our sample selection strategy, we explore the overlap ratio of selected samples between adjacent epochs. To this end, we record the sample selection overlap

Table 4: The ACA (%) results of combining CRSSC with Co-teaching in same and different backbones on CUB200-2011.

Dataset	Net 1	ACA 1	Net 2	ACA 2
CUB200-2011	ResNet-18	77.4	ResNet-18	78.7
	ResNet-18	78.8	ResNet-50	81.9
	ResNet-50	81.9	ResNet-50	81.0

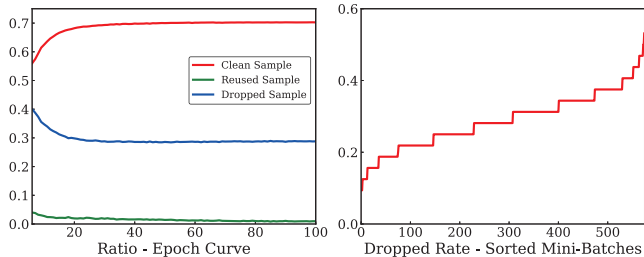


Figure 6: The variations of clean, reusable, and dropped samples during the training processes (left). The drop rate among sorted mini-batches in one epoch (right).

between each epoch and its previous 1, 2, 3, 5, 8, 10 epochs. Let C^i represent the selected dropped (clean) sample set in the i -th epoch, Fig. 5 presents the sample selection overlap of dropped (clean) samples among the current i -th epoch and its previous t epochs $\mathcal{F}_t^i = C^i \cap C^{i-1} \cap \dots \cap C^{i-t+1} \cap C^{i-t}$. From Fig. 5, we can observe that both dropped and clean samples grow steadily and finally converge to a high level, which firmly proves the consistency and robustness of our sample selection strategy.

5.3 Influence of Different Backbones

It is well known that the choice of CNN architectures has a critical impact on object recognition performance. To investigate the influence of different backbones, we conduct experiments by using different basic networks VGG-16 [38], ResNet-18, and ResNet-50 [15]. The experimental results are shown in Table 2.

From Table 2, we can have the following observations: 1) with a much deeper backbone network like ResNet-50, our CRSSC can yield significantly better performance than ResNet-18 and VGG-16. 2) When training a basic network directly with noisy web images, the basic network with higher capacity may produce a worse result. However, by adopting our CRSSC, we can make full use of the learning capacity of basic networks via properly selecting reusable samples and correcting their labels. Compared with the standard network, the improvement of performance demonstrates the superiority of our proposed approach.

5.4 Influence of Different Steps

In this subsection, we investigate the influence of various steps on a basic network like ResNet-18. We first add the Def. 3.1 on the ResNet-18 network to construct a baseline. We then add the Def. 3.1 and 3.2 to construct another baseline. For the third baseline, we add all the Def. 3.1, 3.2, and 3.3 to the ResNet-18. For the fourth baseline, we add the label smoothing technique to complete our CRSSC method. Finally, we present a fine-tuned CRSSC model as the last baseline. The experimental results on the CUB200-2011 dataset

Table 5: The ACA (%) results of our CRSSC training on the combined datasets (web+labeled).

Backbone	FGVC Aircraft	CUB200	Stanford Cars
VGG-16	88.4	85.7	92.4
ResNet-18	88.4	86.8	92.4
ResNet-50	93.4	87.7	94.0

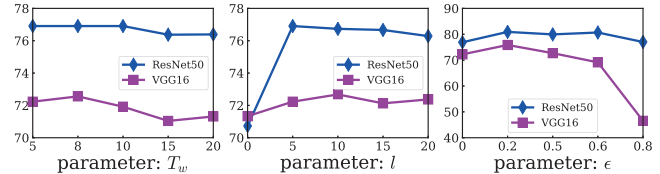


Figure 7: The parameter sensitivities of warm-up epochs T_w (left), history list length l (middle), and the LSR smoothing level ϵ (right) by using different backbones VGG-16 and ResNet-50.

are summarized in Table 3. By observing Table 3, we can find that the fine-tuned CRSSC framework obtains the best performance.

5.5 Combining with Co-teaching

Our proposed method is flexible with regard to combining with other techniques since it only involves sample selection and relabeling. Here, we combine our CRSSC with a state-of-the-art method Co-teaching [14] for further performance improvement. Following the setup in Co-teaching, we maintain two networks simultaneously. In each mini-batch, each network constructs its own \mathcal{B}_{C_E} and $\mathcal{B}_{\mathcal{R}}$ and subsequently feeds them into its peer network for further updating. Different from Co-teaching which only trains model with clean, easy samples, the combined method additionally leverages hard and mislabeled samples to promote the network optimization. Table 4 demonstrates the ACA results of combining CRSSC with Co-teaching in same and different backbones on CUB200-2011 dataset. Compared with the naive CRSSC (presented in Table 2), we can observe that great improvement has been achieved in Table 4.

5.6 Combining Web and Labeled Data

One of the roadblocks that limit the performance of fine-grained visual classification is the lack of enough labeled training data. The widely-used FGVC benchmark datasets (e.g., FGVC-Aircraft, CUB200-2011, and Stanford Cars) all suffer from limited training data, which severely prevented the FGVC task from being sufficiently benefited from the high learning capability of deep CNN. Therefore, employing web images as a supplement to existing fine-grained datasets also attracts considerable attention in recent years. Following the semi-supervised manner, we leverage collected web images as data augmentation to the labeled training data for training deep FGVC model. The experimental results of our CRSSC training on the combined data are shown in Table 5.

5.7 Trend of Samples in Training

We present the ratio variations of identified clean, reusable, and dropped samples during the training processes in Fig. 6 (left). From Fig. 6 (left), we can notice that the ratio of clean samples increases

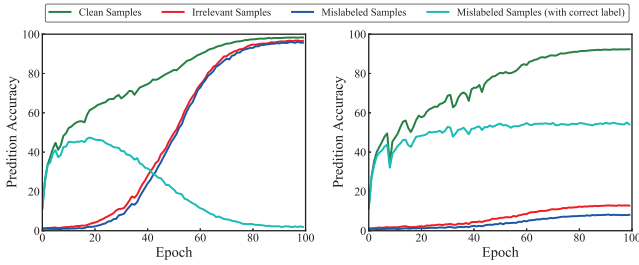


Figure 8: The prediction accuracy (%) of the baseline model (left) and our CRSSC model (right) during the process of training. The “Mislabeled Samples” curve (blue) represents the prediction accuracy with regard to corrupted labels while the “Mislabeled Samples (with correct label)” curve (cyan) denotes the prediction accuracy with regard to ground-truth labels.

steadily until convergence while training progresses. As the training continues, the previous hard examples will gradually become “easy” for our model. On the contrary, as the network training proceeds, the certainty-based criterion will gradually reduce the mistaken dropping, thus leading to a firm decrease in the ratio of dropped samples until convergences to the ground-truth noise rate. Additionally, since hard examples get fewer with the growth of network capability, the ratio of reusable samples also decreases, until only mislabeled examples are left. The final convergence of three groups demonstrates the stability and robustness of our sample selection strategy. Fig. 6 (right) shows the drop rate among sorted mini-batches in one epoch. From Fig. 6 (right), we can find the imbalance of dropped sample ratios across each mini-batch, which proves the necessity of avoiding using a predefined drop rate.

5.8 Parameters in Proposed Approach

For the parameters analysis, we concern three parameters, including 1) the number of warm-up epochs T_w , 2) the length of history list l , and 3) the LSR smoothing level ϵ . Fig. 7 gives the results on CUB200-2011 dataset.

From Fig. 7 (left), we can observe that CRSSC is relatively stable when varying T_w by fixing other two parameters. Both cases achieve the best performance when T_w is selected in [5, 8] and we select $T_w = 5$ as the default option. The length of history list affects the precision of label correction. Intuitively, a higher value of l may benefit the label correction. However, the relabeling could also be misled due to poor predictions of early epochs. From Fig. 7 (middle), we notice that the best performance can be obtained when $l \in [5, 10]$, we select $l = 5$ as the default option. It should be noted that when $l = 0$, this means that we use the prediction label of the current epoch instead of that with the highest accumulated probability in the previous m epochs.

The smoothing level has an influence on the generalization ability. Compared with the case in which LSR is not leveraged (*i.e.* $\epsilon = 0$), from Fig. 7 (right), we can find that the classification accuracy increases considerably when adopting a proper level of label smoothing. When $\epsilon \in [0.2, 0.6]$, the performance is fairly robust. However, as the ϵ gets larger than 0.6, the performance starts to decrease. This probably results from the fact that too large ϵ leads to a lack of proper ground-truth guidance in the training process.

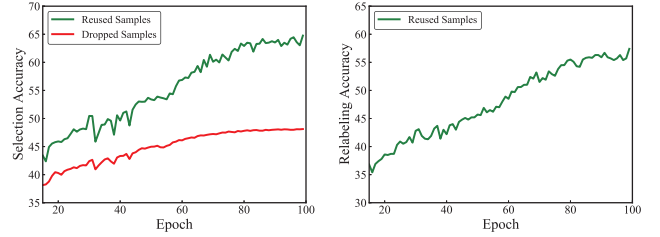


Figure 9: The samples selection accuracy (%) (left) and the sample relabeling accuracy (%) (right) during the training process.

5.9 Further Studies on Noisy-CIFAR100

To further explore the effectiveness of our approach, we follow Co-teaching [14] and generate a synthetic dataset based on CIFAR100 [21] for further study. We first regard the last 20 categories of CIFAR100 as the irrelevant categories. Then we randomly select 20% of the remaining training samples and corrupt their labels to simulate mislabeled data. We named this synthetic dataset as Noisy-CIFAR100.

Prediction Accuracy for Different Samples: As the memorization effects [1, 57] indicated, CNN tends to fit clean samples in initial epochs and will eventually fit noise data (*i.e.*, irrelevant samples, and mislabeled ones). Fig. 8 presents the prediction accuracy of baseline model ResNet-18 (left) and our CRSSC (right) model. By observing Fig. 8 (left), we can observe that, while the prediction accuracy on clean training samples grows steadily in the training process, CNN eventually fits noisy training data and degrades the classification ability of the final model. From Fig. 8 (right), we can notice that, by using our CRSSC, the over-fitting to noise data are effectively suppressed. Besides, by using our relabeling strategy, mislabeled samples can be better learned and finally contribute to boosting the model classification ability.

Samples Selection and Relabeling Accuracy: Fig. 9 (left) shows the samples selection accuracy of our approach and (right) presents the samples relabeling accuracy. From Fig. 9 (left), we can find that the samples selection accuracy (including both reused and dropped samples) grows steadily as the training proceeds. In addition, the samples relabeling accuracy also has a steady increases in training.

6 CONCLUSION

In this work, we studied the problem of training FGVC models directly with noisy web images. Accordingly, we proposed a simple yet effective approach, termed as CRSSC, which trained a deep neural network using additionally selected hard and mislabeled samples to boost the robustness of the model. Comprehensive experiments showed that our approach has achieved state-of-the-art performance, compared with existing webly supervised methods.

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (No. 61976116) and Fundamental Research Funds for the Central Universities (No. 30920021135).

REFERENCES

- [1] Devansh Arpit, Stanislaw Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. 2017. A closer look at memorization in deep networks. In *International Conference on Machine Learning*. 233–242.
- [2] Steve Branson, Grant Van Horn, Serge Belongie, and Pietro Perona. 2014. Bird species categorization using pose normalized deep convolutional nets. In *British Machine Vision Conference*.
- [3] Haw-Shiuan Chang, Erik Learned-Miller, and Andrew McCallum. 2017. Active bias: Training more accurate neural networks by emphasizing high variance samples. In *Advances in Neural Information Processing Systems*. 1002–1012.
- [4] Haw-Shiuan Chang, Erik Learned-Miller, and Andrew McCallum. 2017. Active bias: Training more accurate neural networks by emphasizing high variance samples. In *Advances in Neural Information Processing Systems*. 1002–1012.
- [5] Tao Chen, Jian Zhang, Guo-Sen Xie, Yazhou Yao, Xiaoshui Huang, and Zhenmin Tang. 2020. Classification Constrained Discriminator For Domain Adaptive Semantic Segmentation. In *IEEE International Conference on Multimedia and Expo*. 1–6.
- [6] Yue Chen, Yalong Bai, Wei Zhang, and Tao Mei. 2019. Destruction and construction learning for fine-grained image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*. 5157–5166.
- [7] Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge Belongie. 2018. Large scale fine-grained categorization and domain-specific transfer learning. In *IEEE Conference on Computer Vision and Pattern Recognition*. 4109–4118.
- [8] Yin Cui, Feng Zhou, Yuanqing Lin, and Serge Belongie. 2016. Fine-grained categorization and dataset bootstrapping using deep metric learning with humans in the loop. In *IEEE Conference on Computer Vision and Pattern Recognition*. 1153–1162.
- [9] Yin Cui, Feng Zhou, Jiang Wang, Xiao Liu, Yuanqing Lin, and Serge Belongie. 2017. Kernel pooling for convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*. 2921–2930.
- [10] Abhimanyu Dubey, Otkrist Gupta, Ramesh Raskar, and Nikhil Naik. 2018. Maximum-entropy fine grained classification. In *Advances in Neural Information Processing Systems*. 637–647.
- [11] Yang Gao, Oscar Beijbom, Ning Zhang, and Trevor Darrell. 2016. Compact bilinear pooling. In *IEEE Conference on Computer Vision and Pattern Recognition*. 317–326.
- [12] Weifeng Ge, Xiangru Lin, and Yizhou Yu. 2019. Weakly supervised complementary parts models for fine-grained image classification from the bottom up. In *IEEE Conference on Computer Vision and Pattern Recognition*. 3034–3043.
- [13] Jacob Goldberger and Ehud Ben-Reuven. 2017. Training deep neural-networks using a noise adaptation layer. In *International Conference on Learning Representations*.
- [14] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in Neural Information Processing Systems*. 8527–8537.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.
- [16] Xian-sheng Hua, Fumin Shen, Jian Zhang, and Zhenmin Tang. 2016. A domain robust approach for image dataset construction. In *ACM international conference on Multimedia*. 212–216.
- [17] Shaoli Huang, Zhe Xu, Dacheng Tao, and Ya Zhang. 2016. Part-stacked cnn for fine-grained visual categorization. In *IEEE Conference on Computer Vision and Pattern Recognition*. 1173–1182.
- [18] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. 2017. Mentor-net: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International Conference on Machine Learning*.
- [19] Shu Kong and Charless Fowlkes. 2017. Low-rank bilinear pooling for fine-grained classification. In *IEEE Conference on Computer Vision and Pattern Recognition*. 365–374.
- [20] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 2013. 3d object representations for fine-grained categorization. In *CVPRW*. 554–561.
- [21] Alex Krizhevsky and Geoffrey Hinton. 2009. Learning multiple layers of features from tiny images. *Technical report, University of Toronto* 1, 4 (2009), 7.
- [22] M Pawan Kumar, Benjamin Packer, and Daphne Koller. 2010. Self-paced learning for latent variable models. In *Advances in Neural Information Processing Systems*. 1189–1197.
- [23] Michael Lam, Behrooz Mahasseni, and Sinisa Todorovic. 2017. Fine-grained recognition as hsnet search for informative image parts. In *IEEE Conference on Computer Vision and Pattern Recognition*. 2520–2529.
- [24] Peihua Li, Jiangtao Xie, Qilong Wang, and Zilin Gao. 2018. Towards faster training of global covariance pooling networks by iterative matrix square root normalization. In *IEEE Conference on Computer Vision and Pattern Recognition*. 947–955.
- [25] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *IEEE International Conference on Computer Vision*. 2980–2988.
- [26] Tsung-Yu Lin and Subhransu Maji. 2017. Improved bilinear pooling with cnns. In *British Machine Vision Conference*. 117.1–117.12.
- [27] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. 2017. Bilinear convolutional neural networks for fine-grained visual recognition. 40, 6 (2017), 1309–1322.
- [28] Jiarou Lu, Huafeng Liu, Yazhou Yao, Shuyin Tao, Zhenming Tang, and Jianfeng Lu. 2020. Hsi Road: A Hyper Spectral Image Dataset For Road Segmentation. In *IEEE International Conference on Multimedia and Expo*. 1–6.
- [29] Haonan Luo, Guosheng Lin, Zichuan Liu, Fayao Liu, Zhenmin Tang, and Yazhou Yao. 2019. Segeqa: Video segmentation based visual attention for embodied question answering. In *IEEE International Conference on Computer Vision*. 9667–9676.
- [30] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. 2013. Fine-grained visual classification of aircraft. *arXiv:1306.5151* (2013).
- [31] Eran Malach and Shai Shalev-Shwartz. 2017. Decoupling “when to update” from “how to update”. In *Advances in Neural Information Processing Systems*. 960–970.
- [32] Li Niu, Ashok Veeraraghavan, and Ashutosh Sabharwal. 2018. Webly supervised learning meets zero-shot learning: A hybrid approach for fine-grained classification. In *IEEE Conference on Computer Vision and Pattern Recognition*. 7171–7180.
- [33] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. 2017. Making deep neural networks robust to label noise: A loss correction approach. In *IEEE Conference on Computer Vision and Pattern Recognition*. 1944–1952.
- [34] Scott Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. 2014. Training deep neural networks on noisy labels with bootstrapping. In *International Conference on Learning Representations*. 1–11.
- [35] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. 2018. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning*. 4334–4343.
- [36] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. 2016. Training region-based object detectors with online hard example mining. In *IEEE Conference on Computer Vision and Pattern Recognition*. 761–769.
- [37] Xiangbo Shu, Jinhui Tang, Guojun Qi, Wei Liu, and Jian Yang. 2019. Hierarchical long short-term concurrent memory for human interaction recognition. *IEEE transactions on pattern analysis and machine intelligence* (2019).
- [38] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556* (2014).
- [39] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. 2019. SELFIE: Refurbishing unclean samples for robust deep learning. In *International Conference on Machine Learning*. 5907–5915.
- [40] Zeren Sun, Fumin Shen, Li Liu, and Limin et al. Wang. 2019. Dynamically visual disambiguation of keyword-based image search. *International Joint Conference on Artificial Intelligence* (2019), 996–1002.
- [41] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *IEEE Conference on Computer Vision and Pattern Recognition*. 2818–2826.
- [42] Jinhui Tang, Zechao Li, Hanjiang Lai, Liyan Zhang, Shuicheng Yan, et al. 2017. Personalized age progression with bi-level aging dictionary learning. *IEEE transactions on pattern analysis and machine intelligence* 40, 4 (2017), 905–917.
- [43] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. 2018. The inaturalist species classification and detection dataset. In *IEEE Conference on Computer Vision and Pattern Recognition*. 8769–8778.
- [44] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. 2011. The Caltech-UCSD Birds-200-2011 Dataset. *CNS-TR-2011-001* (2011).
- [45] Xiu-Shen Wei, Chen-Wei Xie, Jianxin Wu, and Chunhua Shen. 2018. Mask-CNN: Localizing parts and selecting descriptors for fine-grained bird species categorization. *Pattern Recognition* 76 (2018), 704–714.
- [46] Guo-Sen Xie, Li Liu, Xiaobo Jin, Fan Zhu, Zheng Zhang, Jie Qin, Yazhou Yao, and Ling Shao. 2019. Attentive region embedding network for zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*. 9384–9393.
- [47] Guo-Sen Xie, Li Liu, Fan Zhu, Fang Zhao, Zheng Zhang, Yazhou Yao, Jie Qin, and Ling Shao. 2020. Region Graph Embedding Network for Zero-Shot Learning. In *European Conference on Computer Vision*.
- [48] Zhe Xu, Shaoli Huang, Ya Zhang, and Dacheng Tao. 2016. Webly-supervised fine-grained visual categorization via deep domain adaptation. 40, 5 (2016), 1100–1113.
- [49] Jufeng Yang, Xiaoxiao Sun, Yu-Kun Lai, Liang Zheng, and Ming-Ming Cheng. 2018. Recognition from web data: A progressive filtering approach. 27, 11 (2018), 5303–5315.
- [50] Hantao Yao, Shiliang Zhang, Yongdong Zhang, Jintao Li, and Qi Tian. 2016. Coarse-to-fine description for fine-grained visual categorization. 25, 10 (2016), 4858–4872.

- [51] Yazhou Yao, Fumin Shen, Guosen Xie, Li Liu, Fan Zhu, Jian Zhang, and Heng Tao Shen. 2020. Exploiting web images for multi-output classification: From category to subcategories. *IEEE Transactions on Neural Networks and Learning Systems* 31, 7 (2020), 2348–2360.
- [52] Yazhou Yao, Fumin Shen, Jian Zhang, Li Liu, Zhenmin Tang, and Ling Shao. 2018. Extracting multiple visual senses for web learning. *IEEE Transactions on Multimedia* 21, 1 (2018), 184–196.
- [53] Yazhou Yao, Fumin Shen, Jian Zhang, Li Liu, Zhenmin Tang, and Ling Shao. 2018. Extracting privileged information for enhancing classifier learning. *IEEE Transactions on Image Processing* 28, 1 (2018), 436–450.
- [54] Yazhou Yao, Jian Zhang, Fumin Shen, Xiansheng Hua, Jingsong Xu, and Zhenmin Tang. 2017. Exploiting web images for dataset construction: A domain robust approach. *IEEE Transactions on Multimedia* 19, 8 (2017), 1771–1784.
- [55] Yazhou Yao, Jian Zhang, Fumin Shen, Li Liu, Fan Zhu, Dongxiang Zhang, and Heng Tao Shen. 2019. Towards automatic construction of diverse, high-quality image datasets. *IEEE Transactions on Knowledge and Data Engineering* 32, 6 (2019), 1199–1211.
- [56] Kun Yi and Jianxin Wu. 2019. Probabilistic end-to-end noise correction for learning with noisy labels. In *IEEE Conference on Computer Vision and Pattern Recognition*. 7017–7025.
- [57] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. 2017. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*.
- [58] Chuanyi Zhang, Yazhou Yao, Huafeng Liu, Guo-Sen Xie, Xiangbo Shu, Tianfei Zhou, Zheng Zhang, Fumin Shen, and Zhenmin Tang. 2020. Web-Supervised Network with Softly Update-Drop Training for Fine-Grained Visual Classification. In *AAAI Conference on Artificial Intelligence*. 12781–12788.
- [59] Chuanyi Zhang, Yazhou Yao, Jiachao Zhang, Jiaxin Chen, and et al. 2020. Web-Supervised Network for Fine-Grained Visual Classification. In *IEEE International Conference on Multimedia and Expo*. 1–6.
- [60] Jian Zhang, Fumin Shen, Xiansheng Hua, Jingsong Xu, and Zhenmin Tang. 2016. Automatic image dataset construction with multiple textual metadata. In *IEEE International Conference on Multimedia and Expo*. 1–6.
- [61] Jian Zhang, Fumin Shen, Xiansheng Hua, Jingsong Xu, and Zhenmin Tang. 2017. A new web-supervised method for image dataset constructions. *Neurocomputing* 236 (2017), 23–31.
- [62] Jian Zhang, Fumin Shen, Wankou Yang, Pu Huang, and Zhenmin Tang. 2018. Discovering and distinguishing multiple visual senses for polysemous words. In *AAAI Conference on Artificial Intelligence*. 523–530.
- [63] Heliang Zheng, Jianlong Fu, Zheng-Jun Zha, and Jiebo Luo. 2019. Looking for the devil in the details: Learning trilinear attention sampling network for fine-grained image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*. 5012–5021.
- [64] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Antonio Torralba, and Aude Oliva. 2016. Places: An image database for deep scene understanding. *arXiv:1610.02055* (2016).