This ICCV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Cross-X Learning for Fine-Grained Visual Categorization

Wei Luo^{1,2} Xitong Yang² Xianjie Mo¹ Yuheng Lu^{2,5} Larry S. Davis² Jun Li³ Jian Yang⁴ Ser-Nam Lim⁵ ¹South China Agricultural University ²University of Maryland, College Park

³MIT ⁴Nanjing University of Science and Technology ⁵Facebook AI

{cswluo,yangxitongbob,cedricmo.cs,junl.mldl,sernam}@gmail.com

{ylu,lsd}@umiacs.umd.edu csjyang@njust.edu.cn

Abstract

Recognizing objects from subcategories with very subtle differences remains a challenging task due to the large intra-class and small inter-class variation. Recent work tackles this problem in a weakly-supervised manner: object parts are first detected and the corresponding part-specific features are extracted for fine-grained classification. However, these methods typically treat the part-specific features of each image in isolation while neglecting their relationships between different images. In this paper, we propose Cross-X learning, a simple yet effective approach that exploits the relationships between different images and between different network layers for robust multi-scale feature learning. Our approach involves two novel components: (i) a cross-category cross-semantic regularizer that guides the extracted features to represent semantic parts and, (ii) a cross-layer regularizer that improves the robustness of multi-scale features by matching the prediction distribution across multiple layers. Our approach can be easily trained end-to-end and is scalable to large datasets like NABirds. We empirically analyze the contributions of different components of our approach and demonstrate its robustness, effectiveness and state-of-the-art performance on five benchmark datasets. Code is available at https: //github.com/cswluo/CrossX.

1. Introduction

Fine-grained visual categorization (FGVC) aims at classifying objects from very similar categories, *e.g.* subcategories of birds [30, 10], dogs [14] and cars [16]. It has long been considered as a challenging task due to the large intra-class and small inter-class variation, as well as the deficiency of annotated data. Benefiting from the progress of deep neural networks [17, 27, 29, 9], the recognition performance of FGVC has improved steadily in recent years and the community has more recently focused on weakly-supervised FGVC that obviates the need of laborintensive part-based annotation. There are two main approaches to weakly-supervised FGVC, namely, exploiting relationships between fine-grained labels to regularize feature learning [31, 35] and localizing discriminative parts for part-specific feature extraction [6, 37]. Compared to labelrelationship based methods, the localization-based methods have the advantages of extracting fine-grained features from local regions where subtle differences between subcategories usually exist.

Early work on localization-based methods typically adopts a multi-stage learning framework: part detectors are first obtained by training on DCNN features [36] or exploiting the hidden representations in DCNNs [34, 26, 40], and then used to extract part-specific features for fine-grained classification. More recent work merges these two stages into an end-to-end learning framework that utilizes the final objective to optimize both part localization and fine-grained classification at the same time [6, 42, 37, 32]. These methods localize semantic parts independently on each image while neglecting the relationships between the part-specific features from different images. [28] explores the relationships between object parts by proposing a soft attentionbased model. The model first generates attention region features of each input image via multiple excitation modules and then guides the attention features to have semantic meaning by adopting a metric learning framework. However, the improvement from their model is limited as optimizing such a metric learning loss is challenging and involves a non-trivial sample selection procedure [33].

We propose Cross-X learning, a simple but effective approach that leverages the relationships between different images and between different network layers for robust finegrained recognition. Similar to [28], our approach first generates attention region features via multiple excitation modules, but it further involves two novel components: a *cross*category *cross*-semantic regularizer (C^3S) and a *cross*-layer regularizer (CL). C^3S is introduced to guide the attention features from different excitation modules to represent different semantic parts. Ideally, the attention features for the same semantic parts, although coming from different images with different class labels, should be more correlated than those for different semantic parts (see Fig. 2). Therefore, C^3S regulates the feature learning by maximizing the correlation between attention features extracted by the same excitation module while decorrelating those extracted by different excitation modules. Compared to the metric learning loss, C^3S can be naturally integrated into the model and easily optimized without any sampling procedure. Meanwhile, we exploit the relationships between different network layers for robust multi-scale feature learning. We first adapt FPN [20] to generate merged features. The merged features enable our model to discover local discriminative structures with both fine spatial resolution and rich highlevel semantic information. To further improve the robustness of the multi-scale features, we introduce a cross-layer regularizer (CL) that matches the prediction distribution of the mid-level features to that of the high-level features by minimizing their KL-divergence. Experimental results on five benchmark datasets show that our approach outperforms or achieves comparable performance to the state-ofthe-art methods. Moreover, our approach is easy to train and is scalable to large-scale datasets as it does not involve multi-stage or multi-crop mechanisms. We make the following contributions:

- We propose a Cross-X learning approach for finegrained feature learning. Cross-X learning explores relationships between features from different images and different network layers to learn semantic part features.
- We address the issue of robust multi-scale feature learning through cross-layer regularization, which matches prediction distributions across layers, thus increasing the robustness of features in different layers.

The remainder of the paper is organized as follows: Section 2 briefly reviews related work to our approach. Our approach is studied and detailed in Section 3. The model ablation studies and experimental results are analyzed and presented in Section 4. We conclude our work in Section 5.

2. Related Work

Fine-grained categorization: Benefiting from the development of DCNNs, *e.g.* AlexNet [17], VGGNet [27], InceptionNet [29], ResNet [9], the study of FGVC has been gradually shifted from strongly-supervised [1, 19, 39] to weakly-supervised [6, 32, 28] in recent years. In the weakly-supervised configuration, to induce models to learn features from the mostly discriminative regions, creating

structural relationship between labels through either intermediate concepts [35, 31] or shared attributes [44, 22], often accompanied by data augmentation [3], has been proposed. Multi-task learning is typically used to make the learning feasible [35, 31, 41]. Another line of research localizes semantic parts first and then learns feature from the localized parts in a multi-stage learning framework [34, 26, 40]. Recently, this line of research combines part localization and feature learning in an end-to-end framework [6, 42, 18, 32]. Exploring relationships between objects in different images for part feature learning has also been investigated but with limited performance [28], due to the non-trivial sample selection involved in optimizing the loss function. Our approach is a step towards improving the efficiency and effectiveness of robustly exploring relationships between different images. We explore correlations between objects from different images in regularization learning and learn robust multi-scale features.

Multi-scale features: Exploiting multi-scale features improves the performance of many visual tasks. Among them, a number of methods make predictions by combining results inferred from multiple individual layers [21, 2], several other approaches first combine multiple layer features and then make a prediction [24, 8, 15]. These approaches marry low-level features' spatial resolution with high-level features' semantic properties. More recent studies have constructed high-resolution multi-scale semantic features by building feature pyramids in DCNNs through lateral connections of bottom-up and top-down feature maps [20]. Nonlinear and progressive connecting structures are studied in [38] to enhance the exploitation of multi-scale features. Multi-scale features have also studied using multigranularity labels [35, 31]. These approaches learn multiscale features by training networks with different granularity of labels. Our work also involves the utilization of multiscale features but exploits the interactions between features at different scales by matching prediction distributions of different layer feature maps.

3. Approach

Cross-X learning involves two main components: 1) A *cross*-category *cross*-semantic regularizer (C^3S) that learns semantic part features by leveraging the correlations between different images (Sec. 3.2). 2) A *cross*-layer regularizer (CL) that learns robust features by matching prediction distributions between different layers (Sec. 3.3). An overview of our approach is depicted in Fig. 1.

3.1. Preliminaries

We begin by briefly reviewing the one-squeeze multiexcitation (OSME) block [28] that learns multiple attention region features for each input image. Let $\mathbf{U} = [u_1, \cdots, u_C] \in \mathbb{R}^{W \times H \times C}$ denote the output feature map of



Figure 1. Overview of our approach. Our network outputs multiple feature maps by employing the OSME block. Two OSME blocks, each with two excitations, are depicted in the last two stages to illustrate our approach. Feature maps from stage L - 1 (blue) and L (red) are combined to generate the merged feature maps (orange). Top-left corner is a zoomed in display of the merging process of the merged features maps. Feature maps are then aggregated to obtain the corresponding pooling features through GAP or GMP. The pooled features from the same stage are mutually constrained by the C^3S regularizer and are simultaneously concatenated to feed into a fully-connected layer to generate logits. The logits are constrained through the CL regularizer after conversion into class probabilities and are combined for classification. Best viewed in color.

a residual block τ . In order to generate multiple attentionspecific feature maps, the OSME block extends the original residual block by performing one-squeeze and multipleexcitation operations.

Formally, OSME first performs global average pooling to squeeze U and produce a channel-wise descriptor $\mathbf{z} = [z_1, \dots, z_C] \in \mathbb{R}^C$. Then a gating mechanism is independently employed on \mathbf{z} for each excitation module, $p = 1, \dots, P$, to output:

$$\mathbf{m}^p = \sigma(\mathbf{W}_2^p \delta(\mathbf{W}_1^p \mathbf{z})) = [m_1^p, \cdots, m_C^p] \in \mathbb{R}^C, \quad (1)$$

where σ and δ refer to the Sigmoid and ReLU functions. Finally, the attention-specific features U_p are generated by re-weighting the channels of the original feature maps U:

$$\mathbf{U}_p = [m_1^p \mathbf{u}_1, \cdots, m_2^p \mathbf{u}_C] \in \mathbb{R}^{W \times H \times C}.$$
 (2)

Although OSME can generate attention-specific features, guiding these features to have semantic meanings is challenging. [28] tackles this by optimizing a metric learning loss which pulls features from the same excitation closer and pushes features from different excitations away. However, optimizing such a loss still poses a challenge and involves a non-trivial sample selection procedure [33].

3.2. Cross-Category Cross-Semantic Regularizer

Instead of optimizing a metric loss as in [28], we propose to learn semantic features by exploring the correlations between feature maps from different images and different excitation modules. Ideally, we want the extracted features from the same excitation module to have the same semantic meaning, even though they come from different images with different class labels. And the extracted features from different excitation modules should have different semantic meanings, even though they come from the same image (see Fig. 2 for an illustration). To achieve this goal, we introduce the cross-category cross-semantic regularizer (C^3S) that maximizes the correlation of features from the same excitation module while minimizes the correlation of features from different excitation modules.

Formally, we first perform global average pooling (GAP) on \mathbf{U}_p to obtain the corresponding pooled features $\mathbf{f}_p \in \mathbb{R}^C$, followed by ℓ_2 normalization ($\mathbf{f}_p \leftarrow \mathbf{f}_p / \|\mathbf{f}_p\|$). Then the correlations between all pairs of excitation modules p and p' form a matrix S:

$$S_{p,p'} = \frac{1}{N^2} \sum \mathbf{F}_p{}^T \mathbf{F}_{p'}, \qquad (3)$$

where T is the transpose operator, N is the batch size and $\mathbf{F}_p = [\mathbf{f}_{p,1}, \cdots, \mathbf{f}_{p,N}] \in \mathbb{R}^{C \times N}$ is a matrix storing the pooled features from excitation module p for all samples in the batch.

The C^3S regularization loss is then constructed from two parts: 1) maximizing the diagonal of S to maximize the correlation within the same excitation module and, 2) penalizing the norm of S to minimize the correlation between different excitation modules:

$$\mathcal{L}_{C^{3}S}(S) = \frac{1}{2} \left(\|S\|_{F}^{2} - 2\|diag(S)\|_{2}^{2} \right), \qquad (4)$$

where $\|\cdot\|$ is the Frobenius norm, and the $diag(\cdot)$ operator extracts the main diagonal of a matrix into a vector. Compared to the triplet based metric learning loss, C^3S loss can be naturally integrated into the OSME block and is easily optimized without any sampling procedure.



Figure 2. An illustration of the C^3S learning. Take the center image as an example, C^3S encourages the excitation modules, \mathbf{U}_1 and \mathbf{U}_2 , to be activated on different semantic parts by exploiting relationships between features from different images (orange dash box) and features from different excitation modules (blue shaded box). Best viewed in color.

3.3. Cross-Layer Regularizer

Exploiting semantic features from different layers of CNNs has been shown to be beneficial to many vision tasks [24, 8, 15, 21, 2]. A simple extension of this idea to fine-grained recognition is to combine the prediction outputs of different layers for the final prediction. However, we observe in our experiments that this simple strategy usually leads to inferior performance (see Sec. 4.3). We hypothesize that the problem is due to two reasons: 1) midlevel features are more sensitive to input changes [5] which makes them less robust for fine-grained recognition where the intra-class variation is large, 2) relationships between the predictions of features are not exploited. To alleviate these problems, we adapt the feature pyramid network (FPN) [20] to integrate features from different layers and propose a novel cross-layer regularizer (CL) that learns robust features by matching the prediction distribution between different layers.

Formally, let $\mathbf{U}^L = {\{\mathbf{U}_p^L\}}_{p=1}^P, \mathbf{U}^{L-1} = {\{\mathbf{U}_p^{L-1}\}}_{p=1}^P$ be the feature maps at stage L and L-1 (here a stage refers to a group of layers that produce feature maps with the same size [9]). We generate the merged feature maps \mathbf{U}_p^G in a similar way to FPN [20] but with two differences. First, the dimensionality reduction of \mathbf{U}_p^L is performed before upsampling. Second, batch normalization (BN) [13] is used after the anti-aliasing operation on the merged feature maps. The procedure can be summarized as:

$$\mathbf{U}_{p}^{G} = \mathbf{BN} \left(\mathbf{K}_{2} * \left(\mathbf{U}_{p}^{L-1} + \text{Bilinear}(\mathbf{K}_{1} * \mathbf{U}_{p}^{L}) \right) \right), \quad (5)$$

where * is convolutional operation, Bilinear(·) denotes bilinear interpolation, \mathbf{K}_1 , \mathbf{K}_2 are 1×1 and 3×3 filters, respectively. \mathbf{U}^G integrates the property of fine spatial resolution in the mid-level layers and the rich high-level semantic in the top-level layers.

To further exploit the relationships between the predictions of features, we propose the CL regularizer that matches the prediction distribution between different layers. Let $\mathbf{Pr}^{L} = \sigma(f(U^{L}))$ and $\mathbf{Pr}^{L-1} = \sigma(f(U^{L-1}))$ be the prediction outputs of stage L and L - 1, where $\sigma(\cdot)$ is the softmax function and $f(\cdot)$ denotes the output layer. The CL regularizer encourages \mathbf{Pr}^{L-1} to match \mathbf{Pr}^{L} by minimizing their KL-divergence:

$$\mathcal{L}_{CL}(\mathbf{Pr}^{L}, \mathbf{Pr}^{L-1}) = \mathrm{KL}(\mathbf{Pr}^{L} \mid\mid \mathbf{Pr}^{L-1})$$
$$= \frac{1}{N} \sum_{n=1}^{N} \sum_{k=1}^{K} p_{nk}^{L} \log \frac{p_{nk}^{L}}{p_{nk}^{L-1}}, \qquad (6)$$

where K is the number of classes. A similar regularizer can be added to constrain the feature maps U^L and U^G as well. The CL regularizer can be viewed as knowledge distillation [7] that uses "soft targets" from U^L with rich structure information to guide the feature learning of U^{L-1} and U^G .

3.4. Optimization

L

Given the feature maps \mathbf{U}^L , \mathbf{U}^{L-1} and \mathbf{U}^G , our final prediction can be obtained by combining their prediction outputs:

$$\mathbf{Pr} = \sigma \left(f(\mathbf{U}^L) + f(\mathbf{U}^{L-1}) + f(\mathbf{U}^G) \right).$$
(7)

Putting this all together, the full objective function of Cross-X learning is:

$$\mathcal{L} = \mathcal{L}_{data} + \gamma \mathcal{L}_{C^3 S} + \lambda \mathcal{L}_{CL}, \qquad (8)$$

$$\mathcal{L}_{data} = -\frac{1}{N} \sum_{n=1}^{N} \sum_{k=1}^{K} c_{nk} \log p_{nk},$$
(9)

$$\mathcal{L}_{C^{3}S} = \gamma_{1}\mathcal{L}_{C^{3}S}(S^{L}) + \gamma_{2}\mathcal{L}_{C^{3}S}(S^{L-1}) + \gamma_{3}\mathcal{L}_{C^{3}S}(S^{G}),$$
(10)
$$\mathcal{L}_{CL} = \lambda_{1}\mathcal{L}_{CL}(\mathbf{Pr}^{L}, \mathbf{Pr}^{L-1}) + \lambda_{2}\mathcal{L}_{CL}(\mathbf{Pr}^{L}, \mathbf{Pr}^{G}),$$
(11)

where L_{data} is the classification loss, γ and λ are hyperparameters that balance the contribution of different costs. Our model can be trained end-to-end using stochastic gradient descent (SGD) and does *not* require other optimization tricks such as multiple crops [37], data augmentation [3], model ensemble [42], and separate initialization [32]

4. Experiments

4.1. Datasets and Baselines

Datasets: We conduct experiments on five fine-grained visual categorization datasets, including NABirds[10],

Datasets	#category	#training	#testing
NABirds[10]	555	23,929	24,633
CUB-Birds [30]	200	5,994	5,794
Stanford Cars [16]	196	8,144	8,041
Stanford Dogs [14]	120	12,000	8,580
FGVC-Aircraft [25]	100	6,667	3,333

Table 1. The statistics of fine-grained datasets in this paper

Caltech-UCSD Birds (CUB-Birds) [30], Stanford Cars [16], Stanford Dogs [14] and FGVC-Aircraft [25]. Note that NABirds is a recently released dataset with much larger scale and many more fine-grained categories. The detailed statistics such as category numbers and data splits are summarized in Tab. 1. We report top-1 accuracy in this study.

Baselines: We compare our approach with various stateof-the-art methods using weakly-supervised learning for fine-grained recognition. For fair comparison, we mainly compare to the results with ResNet-50 as their backbone network and include the best results of VGG based methods for completeness in the following, unless otherwise stated. In addition, an ablation study of Cross-X learning is analyzed based on the SENet backbone [11], since OSME is a direct extension of the SE block. Moreover, we also report results of our approach on the ResNet-50 backbone [9]. All the baselines are listed as follows:

- FCAN [23]: fully convolutional attention network that adaptively selects multiple task-driven visual attentions by reinforcement learning.
- **RA-CNN** [6]: recurrent attention convolutional neural network that localizes discriminative areas and extracts features from coarse to fine scale.
- **DT-RAM** [18]: recurrent visual attention model that selects a sequence of regions through a dynamic continue/stop gating mechanism.
- MA-CNN [42]: multi-attention convolutional neural network that generates multiple parts from spatially-correlated channels via multi-task learning.
- **DFB-CNN** [32]: discriminative filter bank approach that learns a bank of convolutional filters that capture class-specific discriminative patches.
- NTS-Net [37]: navigator-teacher-scrutinizer network finds consistent informative regions through multiagent cooperation.
- MAMC-CNN [28]: multi-attention multi-class constraint approach that learns soft attention masks by regularizing features from different images.
- MaxEnt-CNN [4]: maximum entropy approach provides a training routine to maximize the entropy of the output probability distributions for FGVC.

4.2. Implementation Details

We develop our model in PyTorch, on top of the implementation of SENet/ResNet-50. Specifically, we place the OSME block after conv5_3 and conv4_6 in SENet/ResNet-50. The size of the output feature maps of the two blocks are $14 \times 14 \times 2048$ and $28 \times 28 \times 1024$, respectively. Therefore, the channel sizes of U^L , U^{L-1} and U^G are 4096, 2048 and 2048 when P = 2. We initialize most of our network using the weights pretrained on ImageNet and initialize the newly introduced layers (OSME blocks, FPN blocks) from scratch. No part or bounding box annotations are used during training.

Our network is trained using SGD on a single NVIDIA P6000 GPU with momentum 0.9 and a mini-batch size of 32. The initial learning rate is set to be 0.01 except for the experiments on Stanford Dogs where 0.001 is used. We train the network for 30 epochs and decay the learning rate by 0.1 every 15 epochs. For datasets that do not provide a validation set, we randomly take 10% out of the training samples from each category for validation. Input images are cropped to 448×448 and flipped horizontally with a probability of 0.5. We report our results on a single scale of 448×448 from a single model. More details can be found in the supplementary material.

4.3. Ablation Studies

Effectiveness of C^3S and CL: The effectiveness of our regularization is studied in Fig. 3. We find the performance of our base network (OSME, putting the OSME block after conv5 3 in SENet-50.) is lower than that of the SENet-50 on almost all datasets (SE vs. OSME), this indicates the training difficulties when employing the OSME block for multiple outputs. As we expected, C^3S can effectively regularize the learning of our network to force excitations in the OSME block to be activated on different semantic parts, thus resulting in better features for classification (C3S vs. OSME). In addition, we find combining mid-level (stage L-1) and high-level (stage L) features without a constraint between them results in a performance drop (C3S vs. C3S+GMP). However, CL can effectively increase the robustness of the mid-level feature and thus boosting the performance (C3S+GMP vs. C3S+GMP+CL).

Benefits from the merged feature maps: Employing the merged feature maps can bring systematic performance improvement on all datasets, whether CL is used or not (C3S+GxP vs. C3S+GxP+FP and C3S+GxP+CL vs. C3S+GxP+FP+CL in Fig. 3–4). This validates the effectiveness of our proposal that extra semantic features can be introduced to improve the FGVC performance and, the correctness of our operation that generates the merged feature maps. An interesting observation is that the performance of C3S+GxP+FP is systematically lower than that of C3S+GxP+FP+CL in Fig. 3–4; this signifies that increasing



Figure 3. Ablation performance on 5 benchmark datasets with GMP employed on U_p^{L-1} . The legend only shows the added block/regularizer names with the default ResNet-50 backbone omitted, *e.g.* SE means SENet-50. Best viewed in color.

	NABirds	CUB-Birds	Cars	Dogs	Aircraft
GMP-	81.7	85.2	93.0	83.0	91.1
GAP-	76.3	84.7	90.4	87.3	89.4
GMP+	80.9	84.2	91.9	86.7	90.7
GAP+	81.7	84.7	93.8	87.3	91.3

Table 2. Performance of our approach on five benchmark datasets with GAP and GMP alternatively employed on \mathbf{U}_p^{L-1} . The top group compares results from the approach with CL but without merged feature maps. The bottom group shows results from the approach with merged feature maps but without CL.

the robustness of the newly introduced merged feature maps is also necessary and it further demonstrates that CL has the capability to improve the robustness of mid-level features.

GMP vs. GAP: As indicated in Fig. 1, GAP and global max pooling (GMP) can be alternatively adopted to pool feature maps. However, we only switch the pooling method from GAP to GMP in \mathbf{U}^{L-1} , since we initially thought the discriminative structure of FGVC is local and subtle, thus GMP should have advantage over GAP to capture the these structures, and provide a better feature representation. This is verified on almost all datasets (the top group of Tab. 2). The results indicate CL can collaborate well with GMP to provide robust mid-level features. However, when the network is enhanced by the merged feature maps, which use GAP, but without CL, the results show different behaviour (the bottom group of Tab. 2). GAP+, where GAP is employed on \mathbf{U}^{L-1} , achieves the best performance on Cars, Dogs, and Aircraft but fails to surpass the performance of GMP-, where GMP is employed on \mathbf{U}^{L-1} , on Birds. This phenomenon indicates that GMP is necessary for ascertaining local and subtle structures in categories with fine-andrich texture. The difference caused by employing GMP or GAP on \mathbf{U}^{L-1} can also be observed in Fig. 5 (b) (see Sec. 4.5). Therefore, we report the final results on Cars, Dogs, and Aircraft with GAP employed on \mathbf{U}^{L-1} while on Birds with GMP employed on $U^{\hat{L}-1}$ in Sec. 4.4



Figure 4. Ablation performance on 5 benchmark datasets with GAP employed on U_p^{L-1} . C3S, CL and FP represent C^3S , CL and merged feature maps, respectively. Best viewed in color.

Approach	1-Stage	Sep. Init.	Accuracy
AlexNet-fc6 [10]		Х	35.0
PN-CNN [10]	×	×	74.0
MaxEnt-CNN [4]	,		60.2
(ResNet-50)	V		09.2
SENet-50 [11]	\checkmark	×	82.1
ResNet-50 [9]		×	82.2
MaxEnt-CNN [4]	,	×	82.0
(DenseNet-161)	V		85.0
Cross-X (SENet)		×	86.4
Cross-X (ResNet)	\checkmark	×	86.2

Table 3. Performance on NABirds. The result of PN-CNN is implemented with part annotations based on AlexNet. 1-Stage means the network is trained end-to-end after initialization. Sep. Init. denotes separate initialization.

4.4. Comparison with State-of-the-Art

Results on NABirds: Most previous methods do not report results on this dataset because of the computational complexity of the multi-crop, multi-scale, and multi-stage optimization. Due to the simplicity of our approach, it scales well to big datasets. Tab. 3 compares results from methods that are all optimized on single-crop inputs. Our re-implementation of SENet/ResNet-50 is better than the more sophisticated posed-normalized PN-CNN [1] and the maximum entropy regularized MaxEnt-ResNet-50. The MaxEnt-CNN improves its performance to 83.0% by employing the DenseNet-161 architecture [12]. This shows the benefits brought by more advanced network architectures. However, our Cross-X learning can further outperform it by 3.2% with a relatively simple ResNet-50 backbone, which signifies the effectiveness of our approach.

Results on CUB-Birds: The classification results for CUB birds are presented in Tab. 4. Compared to previous methods, our approach achieves the state-of-the-art performance in a much easier experimental setting, in which only one feedforward operation on a single scale input is needed, without any specialized initialization. Notice that DFB-

Approach	1-Stage	Sep. Init.	Accuracy
FCAN [23]		Х	84.3
RA-CNN [6]	×	\checkmark	85.3
DT-RAM [18]	×	×	86.0
MA-CNN [42]	\checkmark	\checkmark	86.5
NTS-Net [37]	\checkmark	×	87.5
MaxEnt-CNN [4]		×	80.4
SENet-50 [11]	\checkmark	×	83.0
ResNet-50 [9]	\checkmark	×	84.5
Kernel-Pooling [3]	\checkmark	\checkmark	84.7
MAMC-CNN [28]	\checkmark	×	86.2
DFB-CNN [32]	\checkmark	\checkmark	87.4
Cross-X (SENet)		×	87.5
Cross-X (ResNet)		×	87.7

Table 4. Performance on CUB-Birds. RA-CNN and MA-CNN are based on VGGNet. Multi-crop operations are employed in the first group while not in others.

Approach	1-Stage	Sep. Init.	Accuracy
RA-CNN [6]	×		92.5
MA-CNN [42]		\checkmark	92.8
FCAN [23]		×	93.1
DT-RAM [18]	×	×	93.1
NTS-Net [37]		×	93.9
SENet-50 [11]		×	91.6
Kernel-Pooling [3]	\checkmark	\checkmark	92.4
ResNet-50 [9]	\checkmark	×	92.9
MAMC-CNN [28]		×	93.0
DFB-CNN [32]		\checkmark	93.8
MaxEnt-CNN [4]	\checkmark	×	93.9
Cross-X (SENet)	\checkmark	Х	94.5
Cross-X (ResNet)		×	94.6

Table 5. Performance on Stanford Cars. Kernel-Pooling, RA-CNN, and MA-CNN are based on VGGNet. Multi-crop training and testing are employed in the first group.

CNN [32] needs a separate layer initialization to prevent the model learning from degeneration and NTS-Net [37] conducts feature combinations from multiple crops. MA-CNN [42] obtains comparable results based on VGGNet with part localization pretraining and multi-crop inputs. MaxEnt-CNN [4] can achieve 86.5% when implemented with DenseNet-161, MAMC-CNN [28] improves to 86.5% when using ResNet-101 and Kernel-Pooling reaches 86.2% when combined with VGGNet as reported in their work; however, still clearly lower than ours.

Results on Stanford Cars: Tab. 5 shows the results on Stanford Cars. Our Cross-X learning also achieves the state-of-the-art performance on this dataset, even though DBF-CNN [32] and NTS-Net [37] employ separate layer initialization and multi-scale crops, respectively. Kernel-Pooling attains a better result when coupled with VGGNet compared to that of ResNet-50 (91.9%), thus we report the

Approach	1-Stage	Sep. Init.	Accuracy
RA-CNN [6]	×		87.3
FCAN [23]	\checkmark	×	88.9
MaxEnt-CNN [4]		×	73.6
MAMC-CNN [28]	\checkmark	×	84.8
SENet-50 [11]		×	87.1
ResNet-50 [9]	\checkmark	×	88.1
Cross-X (SENet)		×	88.2
Cross-X (ResNet)	\checkmark	×	88.9

Table 6. Performance on Stanford Dogs. The first group uses multi-crop operations while the others are not.

Approach	1-Stage	Sep. Init.	Accuracy
MA-CNN [42]			89.9
NTS-Net [37]	\checkmark	×	91.4
Kernel-Pooling [3]		\checkmark	86.9
MaxEnt-CNN [4]	\checkmark	×	89.8
ResNet-50 [9]		×	90.3
SENet-50 [11]		×	90.6
DFB-CNN [32]	\checkmark	\checkmark	92.0
Cross-X (SENet)		×	92.7
Cross-X (ResNet)	\checkmark	×	92.6

Table 7. Performance on FGVC-Aircraft. The first group uses multi-crop operations. Kernel-Pooling and DFB-CNN are based on VGGNet. MaxEnt-CNN is implemented with DenseNet-161.

VGGNet-driven results in Tab. 5. Compared to MAMC-CNN [28], which learns multiple feature maps by embedding the OSME block in a metric learning framework, our Cross-X learning outperforms it by 1.6%. The improvement indicates the effectiveness of our proposal to learn semantic part features by exploring the correlations between excitation modules and to extract robust features by bridging the relationship between features in different layers.

Results on Stanford Dogs: Classification results are presented in Tab. 6. Surprisingly, the performance of our reimplementation of SENet/ResNet-50 surpasses many previous methods. Even though they can improve their performance by employing more advanced architectures, *e.g.*, MAMC-CNN [28] with ResNet-101 (85.2%) and MaxEnt-CNN [4] with DenseNet-161 (83.6%) as reported in their papers, still falling behind us. However, our Cross-X learning can beat ResNet-50 a bit and achieve the state-of-theart performance by combining with SENet-50 and ResNet-50, respectively. FCAN [23] also achieves the best performance, but it is more complicated than our approach and needs multi-scale multi-crops for model training and testing. In contrast, Cross-X learning is simple and effective.

Results on FGVC-Aircraft: Tab. 7 reports the average class-prediction accuracy. Our approach obtains the best result among methods reporting results on this dataset, even compared to those based on more advanced network



Figure 5. Superimposed display of activation maps (b) \mathbf{U}_p^{L-1} , (c) \mathbf{U}_p^L and (d) \mathbf{U}_p^G from CUB-Birds, Stanford Cars, and FGVC-Aircraft. The first column (a) shows original images and the last two columns (e) are combined activation maps from corresponding columns of \mathbf{U}_p^{L-1} , \mathbf{U}_p^L and \mathbf{U}_p^G . Each of (b)~(e) shows activations of two excitation modules in the corresponding layers. Best viewed in color.

architectures. As the main difference of the categories in this dataset results from the changes of aircraft structures, this result implies that our Cross-X learning is applicable to classification problems with relatively large inter-class structural variation. Notice that the performance of Kernel-Pooling [3], MaxEnt [4] and DFB-CNN [32] methods drop to 83.9%, 85.7% and 91.7% respectively when supported by ResNet-50 instead of VGGNet.

4.5. Visualization

Fig. 5 displays the resized activation maps [43] of 6 images from 3 datasets (see the supplementary material for more displays). Activation maps from the same layer complement each other — they concentrate on different regions of the objects. In addition, we find the activations in corresponding columns of (b)~(d) cover the same object parts at different scales. Compared to the activation maps (c) \mathbf{U}^L , the highly-activated area in (b) \mathbf{U}^{L-1} and (d) \mathbf{U}^G have respectively a relatively small scale and a highlighted center. The activation maps of \mathbf{U}^G can further be seen as the enhanced activation maps of \mathbf{U}^L from that of \mathbf{U}^{L-1} , *e.g.* head of birds, wings of planes. This is consistent with the design of the fine spatial-resolution and rich high-level semantic feature in FPN [20]. The difference caused by employing GMP or GAP on \mathbf{U}^{L-1} can also be observed in (b) where GMP leads to consistent activation in a single region (the first two rows) while GAP results in scattered activations in multiple regions (the last 4 rows). We further present the combined activation maps in (e) to demonstrate the refined final maps taken as input in our approach for classification.

5. Conclusion

We propose Cross-X learning to learn robust fine-grained feature by exploiting relationships between features from different images and different network layers. Our approach leverages the fact that features for the same semantic parts, although coming from different images with different class labels, should be more correlated than those for different semantic parts. Experiments evaluated on five benchmark datasets, ranging from 100 to 555 categories, validate the effectiveness of our approach. Ablation studies further demonstrate the role of every component of Cross-X.

Acknowledgement This work was supported by NSFC (No.61702197), NSFGD (No.2017A030310261), and Facebook AI.

References

- Steve Branson, Grant Van Horn, Serge Belongie, and Pietro Perona. Bird species categorization using pose normalized deep convolutional nets. In *arXiv preprint arXiv:1406.2952*, 2014.
- [2] Zhaowei Cai, Quanfu Fan, Rogerio S. Feris, and Nuno Vasconcelos. A unified multi-scale deep convolutional neural network for fast object detection. In *ECCV*, 2016.
- [3] Yin Cui, Feng Zhou, Jiang Wang, Xiao Liu, Yuanqing Lin, and Serge Belongie. Kernel pooling for convolutional neural networks. In *CVPR*, 2017.
- [4] Abhimanyu Dubey, Otkrist Gupta, Ramesh Raskar, and Nikhil Naik. Maximum entropy fine-grained classification. In *NIPS*, 2018.
- [5] David Eigen, Jason Rolfe, Rob Fergus, and Yann LeCun. Understanding deep architectures using a recursive convolutional network. In *ICLR*, 2014.
- [6] Jianlong Fu, Heliang Zheng, and Tao Mei. Recurrent attention convolutional neural network for fine-grained image recognition. In *CVPR*, 2017.
- [7] Hinton Geoffrey, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *arXiv preprint arXiv:1503.02531*, 2015.
- [8] Bharath Hariharan, Pablo Arbelez, Ross Girshick, and Jitendra Malik. Hypercolumns for object segmentation and finegrained localization. In CVPR, 2015.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In CVPR, 2016.
- [10] Grant Van Horn, Steve Branson, Ryan Farrell, Scott Haber, Jessie Barry, Panos Ipeirotis, Pietro Perona, and Serge Belongie. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In CVPR, 2015.
- [11] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In CVPR, 2017.
- [12] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, 2017.
- [13] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- [14] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Fei-Fei Li. Novel dataset for fine-grained image categorization. In *First Workshop on Fine-Grained Visual Categorization (FGVC) at CVPR*, 2011.
- [15] Tao Kong, Anbang Yao, Yurong Chen, and Fuchun Sun. Hypernet: Towards accurate region proposal generation and joint object detection. In *CVPR*, 2016.
- [16] Jonathan Krause, Michael Stark, Jia Deng, and Fei-Fei Li. 3d object representations for fine-grained categorization. In 4th IEEE Workshop on 3D Representation and Recognition at ICCV, 2013.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.

- [18] Zhichao Li, Yi Yang, Xiao Liu, Feng Zhou, Shilei Wen, and Wei Xu. Dynamic computational time for visual attention. In *ICCV*, 2017.
- [19] Di Lin, Xiaoyong Shen, Cewu Lu, and Jiaya Jia. Deep lac: Deep localization, alignment and classification for finegrained recognition. In *CVPR*, 2015.
- [20] Tsung-Yi Lin, Piotr Dollr, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In CVPR, 2017.
- [21] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. Ssd: Single shot multibox detector. In ECCV, 2016.
- [22] Xiao Liu, Jiang Wang, Shilei Wen, Errui Ding, and Yuanqing Lin. Localizing by describing: Attribute-guided attention localization for fine-grained recognition. In AAAI, 2017.
- [23] Xiao Liu, Tian Xia, Jian Wang, and Yuanqing Lin. Fully convolutional attention localization networks. In *arXiv preprint arXiv:1603.06765*, 2016.
- [24] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [25] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. In *arXiv preprint arXiv:1306.5151*, 2013.
- [26] Marcel Simon and Erik Rodner. Neural activation constellations: Unsupervised part model discovery with convolutional networks. In *ICCV*, 2015.
- [27] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [28] Ming Sun, Yuchen Yuan, Feng Zhou, and Errui Ding. Multiattention multi-class constraint for fine-grained image recognition. In ECCV, 2018.
- [29] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In CVPR, 2015.
- [30] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. Technical report, California Institute of Technology, 2011.
- [31] Dequan Wang, Zhiqiang Shen, Jie Shao, Wei Zhang, Xiangyang Xue, and Zheng Zhang. Multiple granularity descriptors for fine-grained categorization. In *ICCV*, 2015.
- [32] Yaming Wang, Vlad I. Morariu, and Larry S. Davis. Learning a discriminative filter bank within a cnn for fine-grained recognition. In CVPR, 2018.
- [33] Chao-Yuan Wu, R Manmatha, Alexander J Smola, and Philipp Krahenbuhl. Sampling matters in deep embedding learning. In Proceedings of the IEEE International Conference on Computer Vision, pages 2840–2848, 2017.
- [34] Tianjun Xiao, Yichong Xu, Kuiyuan Yang, Jiaxing Zhang, Yuxin Peng, and Zheng Zhang. The application of two-level attention models in deep convolutional neural network for fine-grained image classification. In CVPR, 2015.
- [35] Saining Xie, Tianbao Yang, Xiaoyu Wang, and Yuanqing Lin. Hyper-class augmented and regularized deep learning for fine-grained image classification. In *CVPR*, 2015.

- [36] Zhe Xu, Shaoli Huang, Ya Zhang, and Dacheng Tao. Augmenting strong supervision using web data for fine-grained categorization. In *ICCV*, 2015.
- [37] Ze Yang, Tiange Luo, Dong Wang, Zhiqiang Hu, Jun Gao, and Liwei Wang. Learning to navigate for fine-grained classification. In *ECCV*, 2018.
- [38] Fisher Yu, Dequan Wang, Evan Shelhamer, and Trevor Darrell. Deep layer aggregation. In *CVPR*, 2018.
- [39] Han Zhang, Tao Xu, Mohamed Elhoseiny, Xiaolei Huang, Shaoting Zhang, Ahmed Elgammal, and Dimitris Metaxas. Spda-cnn: Unifying semantic part detection and abstraction for fine-grained recognition. In *CVPR*, 2016.
- [40] Xiaopeng Zhang, Hongkai Xiong, Wengang Zhou, Weiyao Lin, and Qi Tian. Picking deep filter responses for finegrained image recognition. In CVPR, 2016.
- [41] Xiaofan Zhang, Feng Zhou, Yuanqing Lin, and Shaoting Zhang. Embedding label structures for fine-grained feature representation. In *CVPR*, 2016.
- [42] Heliang Zheng, Jianlong Fu, Tao Mei, and Jiebo Luo. Learning multi-attention convolutional neural network for finegrained image recognition. In *ICCV*, 2017.
- [43] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016.
- [44] Feng Zhou and Yuanqing Lin. Fine-grained image classification by exploring bipartite-graph labels. In *CVPR*, 2016.