

Part-Guided Relational Transformers for Fine-Grained Visual Recognition

Yifan Zhao¹, Jia Li¹, Senior Member, IEEE, Xiaowu Chen², Senior Member, IEEE, and Yonghong Tian³, Senior Member, IEEE

Abstract—Fine-grained visual recognition is to classify objects with visually similar appearances into subcategories, which has made great progress with the development of deep CNNs. However, handling subtle differences between different subcategories still remains a challenge. In this paper, we propose to solve this issue in one unified framework from two aspects, *i.e.*, constructing feature-level interrelationships, and capturing part-level discriminative features. This framework, namely PART-guided Relational Transformers (PART), is proposed to learn the discriminative part features with an automatic part discovery module, and to explore the intrinsic correlations with a feature transformation module by adapting the Transformer models from the field of natural language processing. The part discovery module efficiently discovers the discriminative regions which are highly-corresponded to the gradient descent procedure. Then the second feature transformation module builds correlations within the global embedding and multiple part embedding, enhancing spatial interactions among semantic pixels. Moreover, our proposed approach does not rely on additional part branches in the inference time and reaches state-of-the-art performance on 3 widely-used fine-grained object recognition benchmarks. Experimental results and explainable visualizations demonstrate the effectiveness of our proposed approach.

Index Terms—Fine-grained visual recognition, transformers, part discovery, relationship.

I. INTRODUCTION

FINE-GRAINED visual recognition aims to classify and distinguish the subtle differences among sub-categories with similar appearances, which has made great progress in recognizing an increasing number of categories. Benefiting from the development of Convolutional Neural Networks (CNNs), performance on recognizing common object categories, *i.e.*, birds [1], [2], cars [3], and aircrafts [4],

Manuscript received March 4, 2021; revised August 4, 2021 and September 20, 2021; accepted October 24, 2021. Date of publication November 15, 2021; date of current version November 19, 2021. This work was supported by the National Natural Science Foundation of China under Grant 61922006 and Grant 61825101. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Soma Biswas. (Corresponding author: Jia Li.)

Yifan Zhao is with the State Key Laboratory of Virtual Reality Technology and Systems, School of Computer Science and Engineering, Beihang University, Beijing 100191, China (e-mail: zhaoyf@buaa.edu.cn).

Jia Li and Xiaowu Chen are with the State Key Laboratory of Virtual Reality Technology and Systems, School of Computer Science and Engineering, Beihang University, Beijing 100191, China, and also with the Peng Cheng Laboratory, Shenzhen 518055, China (e-mail: jiali@buaa.edu.cn).

Yonghong Tian is with the Department of Computer Science and Technology, Peking University, Beijing 100871, China, and also with the Peng Cheng Laboratory, Shenzhen 518055, China.

Digital Object Identifier 10.1109/TIP.2021.3126490

has increased steadily in the last decade. However, further explorations on discriminative features and representative feature embedding still face great challenges, which also limit the performance improvement on fine-grained recognition tasks.

Existing methods in solving this challenge can be roughly divided into two predominant groups, considering the different learning manners of fine-grained features. The first group tackles the fine-grained classification problem by generating rich feature representations [5]–[9] or applying auxiliary constraints [10], [11]. Leading by the bilinear pooling operation [5], second-order matrices across different channels are widely adopted by introducing compact homogeneous representations [12], hierarchical organizations [13] and other dimensional reduction operations [14]–[16]. Second-order pooling operations describe the fine-grained object features with rich pair-wise correlations between network channels, which are regarded as high-dimensional descriptors for discovering fine-grained features. Beyond the exploitation of second-order matrices, trilinear attention across semantic channels [8], [9] are proposed to build global attention as well as maintaining the feature shapes. However, there still remain two major problems in this group of methods: 1) limited by the natural flaws of CNNs, the constructed correlations are still conducted locally and restricted in channel dimensions, and the long-term spatial relationships are still unperceived; 2) methods of the second group focus on the constraints on feature learning, but neglect the discriminative part features, which are necessary for distinguishing near-duplicated samples.

To tackle the second problem, methods of the other group propose to localize distinct parts for feature enhancement. Pioneer methods tend to obtain object parts by using part detectors [17], [18] or segmentation parsers [19], [20]. Whilst promising results have been achieved by introducing additional network branches, annotating part segmentation masks or bounding boxes is labor-consuming. Toward this issue, recent approaches [21]–[23] resort to attention mechanisms for exploring auto-emerged parts during the backpropagation process. For example, Fu *et al.* [21] proposed a zoom-in strategy to discover the most informative region in a progressive manner. Nevertheless, auto-discovered object parts are usually unstable, which would lead to overfitting on local regions. Hence a contextual understanding of global relationships is needed in solving fine-grained classification problems.

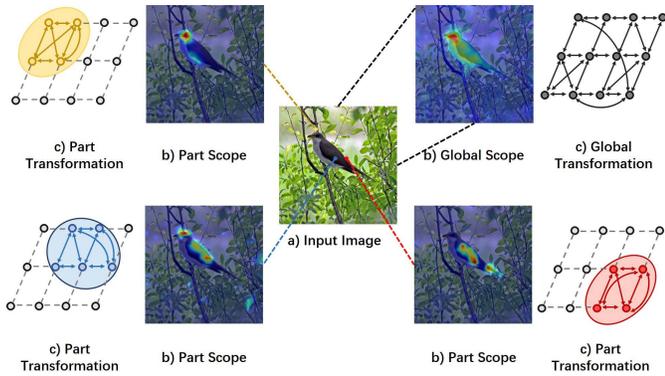


Fig. 1. Motivation of the proposed approach. a): input image. b): one global feature scope and three automatic part discovery scopes. c): relation transformation module to construct global relation embedding and local relational embedding. Our framework first discovers the discriminative local regions and then constructs contextual interactions with relation transformations.

To efficiently solve the deficiencies of these two groups of approaches, we propose to learn the discriminative part features and intrinsic correlations in one unified framework, namely PART-guided Relational Transformers (PART). To be specific, our proposed PART is composed of two essential modules, *i.e.*, a part discovery module to investigate discriminative local regions, and a relational transformation module to construct local and contextual correlations. Inspired by the success in building long-term dependencies in Natural Language Processing (NLP), we make an attempt to learn the relationships among high-level features by using Transformers [24]. In this paper, we embed the Transformer model [24] into the natural design of representative CNN architectures, incorporating self-related global relationships of network features with the stack of transformation layers. The high-level feature maps after CNN encoders are decomposed as several individual vectors by their spatial dimensions. After that, each spatial region has the potential to perceive the contextual information of any other region and builds semantic correlations with each other, which help the networks understand the holistic objects rather than restricted in local patterns. This novel network architecture greatly broadens the receptive field of conventional CNN architectures, while simultaneously keeps their spatial distributions by a positional-aware encoding.

As shown in Fig. 1, our proposed approach first investigates the discriminative regions by generated Class Activation Maps (CAMs) [25], exploring several part scopes and one global scope in Fig. 1 b). These part regions are located by an automatic part discovery module, which maximizes the diversity of different semantic parts while finding discriminative regions simultaneously. With these scopes for object and part discovery, we further propose to learn the local relational embedding and global contextual embedding in Fig. 1 c) with masked multi-head attentions. Hence not only the global relationships are explored in our approach, but also the locally-connected relations. With these two insightful modules, our approach can well embed the long-term dependencies and generate robust part localizations for the accurate classification of fine-grained species. It is worth noting that our proposed PART approach

does not rely on the additional part branches in the inference time, which introduces less computation cost. Moreover, our proposed approach does not use additional data annotations or multi-crop testing operations, and achieves state-of-the-art performances on three public fine-grained recognition benchmarks, *i.e.*, CUB-200-2011 [1], Stanford-Cars [3], and FGVC-Aircrafts [4]. Further experimental results and visualized explanations reveal the effectiveness of the two proposed modules.

In summary, the contribution of this paper is three-fold:

- 1) We propose a novel PART-guided Relational Transformers (PART) framework for fine-grained recognition tasks, which broadens the long-term dependencies of conventional CNNs with a part-aware embedding. To the best of our knowledge, it is the first attempt to handle fine-grained visual recognition tasks with joint CNN and Transformer architectures.
- 2) We propose a new part discovery module to mine the discriminative regional features, to robustly handle differential part proposals by class activation maps.
- 3) We propose a part-guided relational transformation module to learn intrinsic high-order correlations and conduct comprehensive experiments to verify the effectiveness of our proposed approach, surpassing state-of-the-art models by a margin on 3 widely-used benchmarks.

The remainder of this paper is organized as follows: Section II briefly reviews related works of fine-grained recognition tasks and Section III describes how to discover distinct part features and explore the part-guided contextual relationship in our proposed approach. Qualitative and quantitative experiments are reported in Section IV. Section V finally concludes this paper.

II. RELATED WORK

Early fine-grained visual recognition researches [26]–[28] focused on tackling limited labeling data with hand-crafted conventional features. For example, Yao *et al.* [26] proposed to use dense sampling strategies and random forest for mining discriminative features. Moreover, methods [27] using these hand-crafted features show its benefits considering the independence of human labeling or encoded dictionary features. Recent ideas in solving fine-grained visual recognition tasks focus on the exploitation of deep features, which show both robustness and performance boost for handling discriminative features. Considering the different constraints of representation learning, we mainly consider two families of methods in this paper, *i.e.*, part awareness learning and regularized feature representation.

A. Part Awareness Learning

Handling subtle differences using part-level local features has been widely studied in the past decades. Pioneer works [17], [20], [29]–[32] tend to discover the part-level features with manual labels and amplify these local representations when obtaining the final features. For example, Zhang *et al.* [29] proposed a part-based R-CNN network,

using part detectors for building pose-normalized representations. Krause *et al.* [33] proposed to detect the object part with an unsupervised co-segmentation and alignment manner, enhancing the final classification abilities. Huang *et al.* [17] integrated a part detection network in the recognition network and developed two streams to encode the part-level and object-level cues simultaneously. In addition, Wei *et al.* [31] proposed to learn accurate object part masks using an additional FCN and then aggregated these masked features with the global ones. Although promising results have been achieved by using strongly supervised annotations, labeling segmentation masks and accurate object parts are still time- and labor-consuming.

On the other hand, deep neural networks have the natural ability in discovering object parts [34] with only weakly supervised labels. Object parts would emerge automatically in the feature maps during the gradient learning process. Benefiting from these observations, recent works [21]–[23], [35]–[39] proposed to learn the part-aware features with auto-discovered part masks. Leading by this motivation, Simon *et al.* [36] proposed a local region discovery model using deep neural activation maps, constructing part-level attentions for final representation. Lam *et al.* [37] proposed to search object parts using heuristic function and successor functions. Wei *et al.* [40] proposed a flood-fill method which aggregated these features maps as holistic object activations. However, this operation only considers the extraction of object features, which does not reflect the part localizations. Peng *et al.* [41] proposed to select image patches as candidates to select the informative regions and extract the feature of these regions for complementation. Wang *et al.* [39] proposed to regularize the learned feature maps in a sparse representation manner and adopted auxiliary part branches for feature representation regularization. Moreover, Ge *et al.* [18] proposed a bottom-up architecture to obtain instance detection and segmentation via weakly supervised object labels. Combining these part proposals, a context encoding LSTM is further proposed to fuse the sequential part information for image classification. However, these works only considered the localization of part regions, while the contextual relationship of these local regions is less explored. In this paper, we propose to discover the discriminative parts and investigate their spatial relations in one unified framework.

B. Regularized Feature Representation

The other family for solving the fine-grained visual categorization task is to regularize the feature learning process or formulating rich feature relationships. Inspired by the bilinear pooling operation [5], high-order relationship matrices [7], [8], [13], [42], [43] has been widely used for feature representations. For example, Yu *et al.* [13] proposed to learn the heterogeneous feature interaction from different levels of feature layers, building compact cross-layer relationships. However, simply adopting the second-order feature would lead to a dimensional disaster for optimization, thus subsequent works proposed to use low-rank presentations [14], feature factorization [15] or Grassmann constraints [16] for building compact homogeneous features. These works [14]–[16] greatly

reduce the computation cost of high order matrices but still show comparable results. Besides the efforts on the second-order relationships, the recent network tends to explore the third-order relationships of feature channels. Leading by the non-local [44] design in building global relationships of different channels, several works [8], [9] have been proposed to explore the channel-wise relationship by attention mechanisms. Gao *et al.* [9] proposed to learn the channel-wise relationship by image-level relations and cross-image relations by a contrastive constraint.

Thinking the feature representation problem from another perspective, other works tend to explore the cross-layer relationships [6], [10], [45]–[48] or auxiliary information [49]. Sun *et al.* [10] proposed a multi-attention multi-constraint (MAMC) to regularize the same class focusing on the same attention region. Based on this MAMC module, Luo *et al.* [47] proposed to learn relationships between different images and different network layers. In addition, Chen *et al.* [50] proposed a destruction and reconstruction framework to learn the local feature transformation rules for obtaining robust feature representations. Moreover, an attentive pair relation network [51] is proposed to fine the group-wise relationship inter- and intra-classes. Beyond these designs on network architectures, adopting pair-wise confusions [52] and maximizing the entropy [11] also show effectiveness in improving the final representation. Different from the aforementioned methods, in this work, we propose to find the feature embedding by two steps, *i.e.*, discovering local parts and transforming for relational embedding, which incorporates the advantages of part discovery but also constructs reliable feature regularization.

C. Vision Transformers

Transformer architectures [24], [53] have achieved great progress in the field of natural language processing. Benefiting from its ability for long-term relationship construction, several works [54], [55] proposed to introduce this architecture based on the conventional CNN encoders. For example, Carion *et al.* [54] proposed to localize different object regions with multiple attention regions in transformer layers. Beyond these works, recent ideas [56]–[59] proposed to directly encode the image patches, *e.g.*, 16×16 as the basic pattern to understand images. As the representative work, ViT [56] built a strong pretrained transformer backbone, which greatly surpasses the performance of conventional CNNs. However, this encoding manner also introduces huge computation costs. Several works proposed to introduce the shifted windows [57] and local encoding [58] to build hierarchical encoders, which exploit the advantages of local understandings in CNNs and also reduce computation costs. Based on these successful backbone encoders, utilizing vision transformers has also shown great benefits in semantic segmentation [60] and fine-grained visual classification [61]. Different from the TransFG [61] with ViT [56] encoders, our model builds relationships between high-level CNN features. In this manner, the local patterns with detailed features can be well represented and the contextual relationships are also constructed simultaneously.

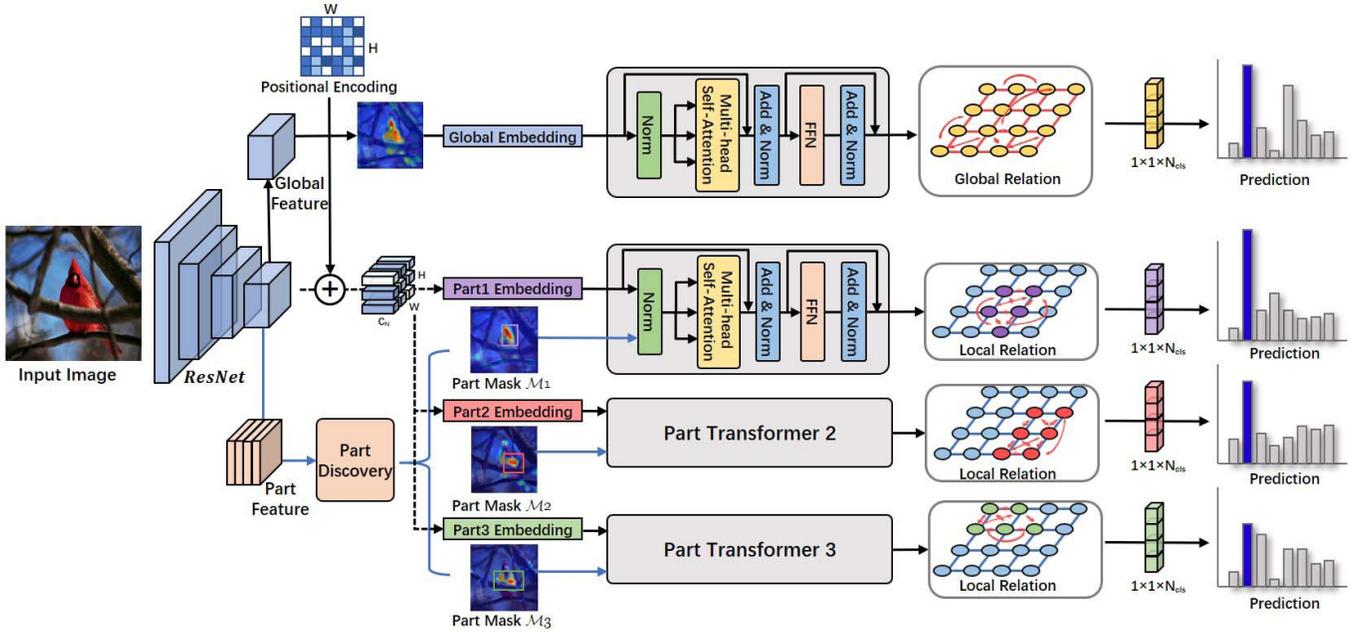


Fig. 2. Overall pipeline of our proposed PART-guided Relational Transformers (PART) framework, which is composed of a part discovery module and a relational transformation module. The part discovery module automatically generates object parts from class activation maps during the forward propagation process. And the relational transformation module takes advantages of object parts and constructs a global interaction and S ($S=3$ for example) local interactions. During the inference stage, only the global branch is maintained, introducing less computation cost.

III. APPROACH

A. Problem Formulation

Our proposed PART approach for fine-grained visual recognition is illustrated in Fig. 2. Our first key idea in exploring fine-grained recognition is to build contextual dependencies by the conjunction of successful Transformers with conventional Convolutional Neural Networks (CNNs). The second motif is to regularize the network embedding be aware of the local discriminative regions, which is desired for distinguishing visually similar samples. Beyond these designs, one global relation transformation and three local relation transformations are proposed to learn the crucial attention features by building regionally dense connections during the gradient descent process.

Given an input image \mathcal{I} , the general procedure in optimizing fine-grained recognition task is to optimize the target \mathbb{E} between prediction and groundtruth label \mathbf{y} :

$$\theta^* = \arg \min_{\theta} \mathbb{E}(\Phi(\mathcal{I}; \theta)^\top \mathbf{w}_c, \mathbf{y}), \quad (1)$$

where $\Phi(\cdot; \theta)$ represents feature extractor Φ with learnable parameters θ . \mathbf{w}_c represents the final classification projection vectors. Different from the general classification optimizations, our proposed network first generates two different feature embedding, *i.e.*, $\mathbf{X}^p, \mathbf{X}^g \in \mathbb{R}^{W \times H \times C}$. Thus we could localize the object part proposals \mathbf{P} with part discovery module \mathcal{M}_{part} , *i.e.*, $\mathcal{P} = \mathcal{M}_{part}(\mathbf{X}^p)$, and construct high-order correlations with Transformers \mathcal{G}, \mathcal{T} for global and local branches.

Beyond these foundations, we thus make a general assumption that the object part proposals \mathcal{P} share the same groundtruth labels \mathbf{y} with the global image \mathcal{I} . Hence the

overall optimization in our framework can be presented as:

$$\theta^* = \arg \min_{\theta} \mathbb{E}((\Phi(\mathcal{I}; \theta_b) \otimes \mathcal{G}(\mathbf{X}^g; \theta_g))^\top \mathbf{w}_g, \mathbf{y}) + \sum_{\mathbf{p} \in \mathcal{P}} \lambda_p \mathbb{E}((\Phi(\mathcal{I}; \theta_b) \otimes \mathcal{T}^p(\mathbf{p}, \mathbf{X}^p; \theta_p))^\top \mathbf{w}_l, \mathbf{y}), \quad (2)$$

where \otimes indicates that the latter term extract features from the former one. $\mathbf{w}^{(g,l)}$ represent the global and local classification projections and λ is the balanced weight. Hence our network introduces the relational transformation constraints and part localization constraints in fine-grained classification tasks. We will elaborate the details of these two modules in Section III-B and Section III-C.

B. Part Discovery Module

Excavating discriminative features by localizing object parts has been widely explored by previous works [17], [20], [29]–[32] with bounding box annotations or segmentation mask annotations. In our proposed PART framework, we propose to explore the object parts in a weakly supervised manner, which takes advantage of the class activation maps generated during the gradient descent process. Recent research [34] shows that object parts emerge automatically during the learning of classification tasks, *i.e.*, during the training process, network features can be automatically regularized for excavating discriminative features and distinguishing near-duplicated objects.

Based on these findings, several recent ideas, *e.g.*, RA-CNN [21], MA-CNN [46], propose to constrain these activation regions into part groups by their similarities. Different from these works, our part discovery module aims to provide

discriminative region priors \mathcal{P} for the transformer module. With these priors, the transformer module is proposed to mine the contextual object clues within each region. To achieve this, our proposed part discovery module follows two meaningful rules: 1) the highly activated features maps are attached with high priorities during the back-propagation process; 2) different activated parts should have the minimum overlap for discriminative part finding. These rules ensure CAMs are not restricted to local regions and enhance the generalization ability of deep models. Besides that, our part discovery module can be trained end-to-end in one unified scheme and does not cost additional computation costs for inference.

Given the extracted backbone features $\mathbf{X}^p \in \mathbb{R}^{W \times H \times C}$, we define the class activation scores \mathbf{s}_c of the c th channel:

$$\mathbf{v}_c = \frac{1}{W \times H} \sum_{i=1}^H \sum_{j=1}^W \mathbf{X}_{i,j,c}^p, \quad (3)$$

$$\mathbf{s}_c = \frac{\mathbf{v}_c}{\sum_{k=1}^C (\mathbf{v}_k + \epsilon)},$$

where ϵ is set to keep nonzero of the denominator. With this score as the ranking indicator, we thus rearrange the $\mathbf{X}^* = \text{ActivationSort}(\mathbf{X}, \mathbf{s})$ and build a part proposal stack to select the most important one from the stack top.

Starting from this part proposal stack, we first sample a random rate for each part $\eta_c \sim N(\mu, \sigma)$, enhancing robustness for threshold selection. For each selected feature $\mathbf{x}_c^* \in \mathbf{X}$, the part proposals \mathbf{p}^* are localized with the ROIcrop operation:

$$\mathbf{p}_{i,j,c}^* = \begin{cases} 1 & \frac{\mathbf{x}_{i,j,c}^* - \min(\mathbf{x}_c^*)}{\max(\mathbf{x}_c^*) - \min(\mathbf{x}_c^*)} \geq \eta_c, \\ 0 & \frac{\mathbf{x}_{i,j,c}^* - \min(\mathbf{x}_c^*)}{\max(\mathbf{x}_c^*) - \min(\mathbf{x}_c^*)} < \eta_c, \end{cases} \quad (4)$$

The final bounding box regions can be automatically obtained by calculating the range of coordinates in the x -axis and y -axis. However, as an unsupervised part discovery method, different feature maps usually localize similar part regions, which may restrict the learning of discriminative features. Inspired by the Non-Maximum Suppression operation [62] in edge detection, we remove the redundant part proposals by calculating the Intersection over Union (IoU) of the current bounding box with the existing boxes in the part selection set \mathcal{P} . By repeating this process, all the legal part proposals will be automatically added into the final selection set until $|\mathcal{P}| \geq N$. Specially, we further apply an iteration terminator if all this part selection module could not find appropriate object parts. The detailed process is elaborated in Algorithm 1. With this proposed part selection module, our approach has the potential to find discriminative regional features automatically.

C. Relation Embedding With Transformers

As mentioned in Section III-A, one crucial problem in fine-grained recognition is to build contextual dependencies for different semantic parts. However, limited by the design of 2D Convolutional Neural Networks, each feature unit has a restricted receptive field which is fixed by the convolutional

Algorithm 1 Part Discovery Algorithm

Input: Part feature maps $\mathbf{X} \in \mathbb{R}^{W \times H \times C}$, Hyperparameters: Number of parts N , Selection range R , IoU threshold th
Output: Part Selection Set: $\mathcal{P} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N\}$

- 1: Initialize selection set $\mathcal{P} = \phi$
- 2: **while** $|\mathcal{P}| < N$ **do**
- 3: Calculate part activation scores \mathbf{s} of \mathbf{X} in Eqn. (3)
- 4: Rearrange \mathbf{X} by the ranking of \mathbf{s} :
 $\mathbf{X}^* = \text{ActivationSort}(\mathbf{X}, \mathbf{s})$
- 5: **for** $\mathbf{x}_i^* \in \mathbf{X}^*$ **do**
- 6: Sample Random Activation Rate: $\eta_i \sim N(\mu, \sigma)$
- 7: $\tilde{\mathbf{x}}_i = \text{Norm}(\mathbf{x}_i^*)$
- 8: $\mathbf{p}_i^* = \text{ROIcrop}(\tilde{\mathbf{x}}_i \geq \eta_i)$:
- 9: **for** $\mathbf{p}_j \in \mathcal{P}$ **do**
- 10: Calculate bounding box IoU:
 $\text{BIoU} = \text{BboxIoU}(\mathbf{p}_j, \mathbf{p}_i^*)$
- 11: **if** $\text{BIoU} > th$ **then**
- 12: **break**
- 13: Add new legal proposals: $\mathcal{P} = \mathcal{P} \cup \{\mathbf{p}_i^*\}$
- 14: iter = iter+1
- 15: **if** iter > maxiter and $|\mathcal{P}| < N$ **then**
- 16: Select top N parts from \mathbf{X}^* and repeat Step 6 ~ 8
- 17: **return** Part Selection Set \mathcal{P}

kernels. To solve this problem, we propose to adopt the Transformer [24], which has been established as a state-of-the-art transduction model to take advantage of long-term correlation in semantic sentence understanding. Motivated by its successes in natural language processing, we make an attempt to introduce it into the fine-grained recognition task, which is presented as the transformation operation in Eqn. (2).

1) Fine-Grained Recognition With Visual Transformers:

The conventional transformers tend to take the input of a whole language sentence and building correlations of different semantic words, which are now used to replace the Recurrent Neural Networks (RNNs) in many NLP tasks to build long-term dependencies. Besides its superior performance in relation to modeling, embedding this architecture into existing CNN backbones, *e.g.*, ResNet, VGGNet, is thus essential. In our PART framework, CNN backbones and transformers play different roles in fine-grained visual understanding. Solely adopting transformer models may result in a larger receptive field in relational modeling, but leads to a high computation burden for network optimization. In addition, keeping the main body of CNNs can also benefit from the pretrained common knowledge, *e.g.*, ImageNet [63]. Hence in our approach, we propose to make a conjunction of conventional CNNs and the successful Transformer model. The detailed architecture is illustrated in Fig. 3.

2) *Positional Encoding*: Another drawback in attention-based modeling methods is that structural information would be forgotten, especially when transforming CNN feature maps $\mathbf{X} \in \mathbb{R}^{W \times H \times C}$ to multiple vectors. When dense connections are constructed, their original structural information could be easily lost. Inspired by [24], [54] in keeping position

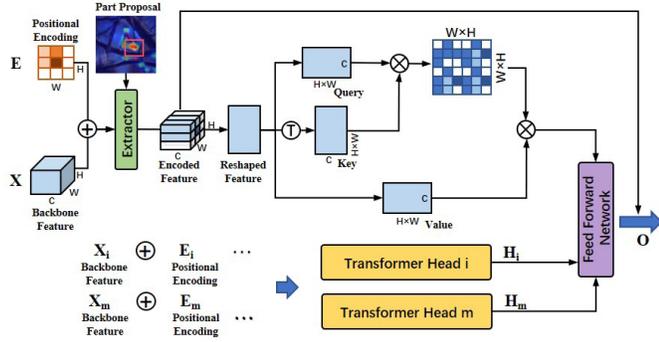


Fig. 3. Illustration of the part-guided relational transformation module. When input backbone features and positional encoding, this module builds local transformation embedding with efficient self-attention mechanisms with a multi-head ($\mathbf{H}_1 \dots \mathbf{H}_m$) encoder manner.

information, here we adopt the positional encoding \mathbf{E} into the original backbone features \mathbf{X} . With the positional encoded features regardless of part proposals in Fig. 3, the key idea is to build global relation matrices $\mathbf{A} \in \mathbb{R}^{W \times H \times W \times H}$, indicating the dense dependencies of each pixel unit to all the other pixels. Let i, j be the positions of query and key values, this spatial correlation can be formulated as:

$$\mathbf{A}_{i,j} = (\mathbf{X}_{i,:} + \mathbf{E}_{i,:}) \mathbf{W}_{qry} \mathbf{W}_{key}^T (\mathbf{X}_{j,:} + \mathbf{E}_{j,:})^T, \quad (5)$$

where $\mathbf{W}_{\{qry,key\}}$ denote the learnable weight of query and key projections. Here we adopt the relative positional encoding of Transformer-XL [64] to build different positional encoding for different pixels. We build the 1D learning for relative encoding and 2D positional learning for absolute encoding. By expanding Eqn. (5) and using learnable \mathbf{u}, \mathbf{v} , this relative encoding correlations \mathbf{A} can be presented as:

$$\mathbf{A}_{i,j} = \mathbf{X}_{i,:}^T \mathbf{W}_{qry}^T \mathbf{W}_{key} \mathbf{X}_{j,:} + \mathbf{X}_{i,:}^T \mathbf{W}_{qry}^T \widehat{\mathbf{W}}_{key} \mathbf{R}_{i-j} + \mathbf{u}^T \mathbf{W}_{key} \mathbf{X}_{j,:} + \mathbf{v}^T \widehat{\mathbf{W}}_{key} \mathbf{R}_{i-j}, \quad (6)$$

where key weights are split into \mathbf{W}_{key} and relative position $\widehat{\mathbf{W}}_{key}$ respectively. \mathbf{R} denotes the sinusoid encoding matrix in [24] without learnable weights.

3) *Multi-Head Attention With Transformers*: The relational matrices $\mathbf{A}_{i,j}$ are multiplied with the original feature \mathbf{X} , resulting in the enhanced feature of one head \mathbf{H} :

$$\mathbf{H} = \frac{\exp(\mathbf{A} / \sqrt{C})}{\sum_{j=1}^{W \times H} \exp(\mathbf{A}_{:,j} / \sqrt{C})} \mathbf{X} \mathbf{W}_{val}. \quad (7)$$

As illustrated in Fig. 3, each transformer unit is composed of multiple transformer heads, *i.e.*, $\mathcal{H} = \{\mathbf{H}_1, \dots, \mathbf{H}_m\}$ and m denotes the number of attention heads. Employing multiple attention heads encourages different projections of query, key and value triplets to construct relational embedding. Then these attention heads are concatenated and fused with the feed-forward network, and the final output \mathbf{O} can be formally presented as:

$$\mathbf{O} = \text{Concat}(\mathbf{H}_1, \dots, \mathbf{H}_m) \mathbf{W}^o + \mathbf{b}^o, \quad (8)$$

where $\mathbf{W}^o \in \mathbb{R}^{m C_h \times C}$ denotes the learnable weight for linear transformations, and C_h denotes the dimension of one attention head.

4) *Relations With CNNs*: Based on this designment of relational embedding transformers, networks have the potential to learn the long-term dependencies between any two pixels. Notably, this attentive transformer layer and convolutional layers process 2D features in similar ways. In other words, the transformer layers are attached with more relational embedding learning matrices compared to the convolutional layers and in extreme cases, the transformers can be degenerated to simulate CNNs. The main theorem from [65] is as follows:

Theorem 1: A transformer layer with m heads and with hidden dimension C_h and output dimension C_{out} can simulate convolutional neural networks with $\min(C_h, C_{out})$ channels.

Detailed proof can refer to [65], which is abstracted in Appendix . Theorem 1 shows that the transformer modules have the potential to represent existing CNN layers but have the strong potential to explore cross-pixel relations. It means that the CNNs are “restricted” transformers that have fixed attentional embedding or limited receptive fields.

5) *Part-Guided Relation Embedding*: As aforementioned in Eqn. (2), our proposed optimization framework aims to construct global relations as well as discovering discriminative part relations simultaneously. Benefiting from the strong relation embedding ability of transformer models, we propose to learn the part-level relations by the auto-discovered part in Section III-B. We first adopt the learned part masks $\mathbf{p}_i \in \mathcal{P}$ as attentive masks before feeding into the transformer layers. Thus the learned architecture of each attention head can be presented as:

$$\mathbf{H}^{\mathbf{p}_i} = \text{Attention}((\mathbf{E}(\mathbf{p}_i) + \mathbf{X}_{i,:} \odot \mathbf{p}_i)), \quad (9)$$

where \odot denotes the Hadamard product for mask region selection. We thus concatenate the multiple heads of each part following Eqn. (8), forming output $\mathbf{O}^{\mathbf{p}_i}$ of each part \mathbf{p}_i . After passing the mask selection operation, a dense connection within each local part is constructed, which discovers the spatial relations while neglects the unnecessary ones.

Moreover, the part-guided transformer units can be stacked in sequence, passing each enhanced feature $\mathbf{O}^{\mathbf{p}_i}$ into the next layer. This operation enlarges the reception field of each pixel and constructs strong correlations by mixing connections hops:

$$\mathbf{O}_t^{\mathbf{p}_i} = \mathcal{T}_{t-1}^{\mathbf{p}_i}(\mathbf{O}_{t-1}^{\mathbf{p}_i}, \mathbf{E}_t + \mathbf{X}, \mathbf{p}_i), \quad t = 1 \dots S, \quad (10)$$

where S denotes the number of stacked layers of transformer model \mathcal{T} . Similarly, the global relation embedding transformer \mathcal{G} is defined in similar manners.

In our training scheme, we conduct a collaborative learning scheme of the global and N part transformation branches in Fig. 2. Each branch is supervised with an individual softmax cross-entropy loss for minimizing target \mathbb{E} in Eqn. (2). Thus our overall learning is composed of one global loss for contextual relation learning and N part losses for local feature discovery. Note that our proposed PART approach resort to the N auxiliary part branches for feature regularization in the training stage, whilst not relying on these parts in the testing stage. Moreover, simply aggregating these features would also lead to overfitting issues and inferior results. Compared to the ResNet backbone, we only introduce limited learnable matrices for relational embedding and conduct an end-to-end

training with the conjunction of part discovery module and relational embedding module. Our framework greatly alleviates the computation burden and retains the discriminative part simultaneously.

IV. EXPERIMENTS

In this section, we first introduce the experimental settings with dataset statistics in Section IV-A and elaborate the implementation details and network architectures in Section IV-B. Subsequently, the comparison with state-of-the-art methods is exhibited in Section IV-C. The detailed performance analyses, as well as the explainable visualizations, are presented in Section IV-D.

A. Experimental Settings

To evaluate the effectiveness of our proposed PART approach, we conduct experiments on three public popular benchmarks, namely Caltech-UCSD Birds (CUB-200-2011) [1], Stanford-Cars [3], and FGVC-Aircrafts [4]. CUB-200-2011 dataset [1] contains 11,788 images of 200 wild bird species, which is the most widely used benchmark. Stanford-Cars dataset [3] includes 16,185 images of 196 car subcategories. FGVC-Aircraft dataset [4] contains 10,000 images of 100 classes of fine-grained aircrafts. We follow the standard training/testing splits as in the original works [1], [3], [4]. For evaluations, we use the top-1 evaluation criterion following previous works and only adopt the classification label for supervised training without any additional part annotations.

B. Implementation Details

For the implementation of baseline and our proposed PART framework, we adopt ResNet-50 and ResNet-101 [67] networks pretrained on ImageNet [63] as our backbone for fair comparisons. We remove the last bottleneck in stage 4 to directly get the feature of 512 dimensions. We use the SGD optimizer with an initial learning rate of $8e - 4$ annealed by 0.1 for every 60 epochs. We adopt the commonly used data augmentation techniques, *i.e.*, random cropping and erasing, left-right flipping, and color jittering for robust feature representations. Our model is relatively lightweight and is trained end-to-end on two NVIDIA 2080Ti GPUs for acceleration. We use the group-wise sampler with a batchsize of 16, sampling 4 samples for each class. For hyper-parameter fine-tuning, we randomly select 10% of the training set as validation. The balanced weights λ_p for multiple part loss functions are empirically set as 0.1. The channel dimension for feeding into relation embedding module is $C = 512$ and the part branch number is set as 4, the IoU threshold is $th = 0.6$. The stacked layer number S is set as 3 for the global branch and 1 for the local branches. The training and testing protocol follow the state-of-the-art works [47], [50], [51] using random cropping of 448×448 in training and only use one-crop during inference, which does not rely on the multi-crop features like previous works [21], [23], [46]. Notably, in the inference stage, we only use the global feature for accuracy reports, and the local branches can be simply omitted for acceleration.

TABLE I
COMPARISON WITH STATE-OF-THE-ART RESULTS ON CUB-200-2011 DATASET. †: USING ADDITIONAL ANNOTATIONS

Type	Method	Backbone	Accuracy
Part Based	PA-CNN† [33]	VGG-19	84.3%
	FCAN† [66]	ResNet50	84.7%
	RA-CNN [21]	VGG-19	85.3%
	MA-CNN [46]	VGG-19	86.5%
	Interpret [20]	ResNet101	87.3%
	NTS-Net [23]	ResNet50	87.5%
	DF-GMM [39]	ResNet50	88.8%
Feature Based	Bilinear [5]	VGG-16	84.0%
	MAMC [10]	ResNet101	86.5%
	MaxEnt [11]	DenseNet-161	86.5%
	DFL-CNN [6]	VGG-16	86.7%
	PC [52]	DenseNet-161	86.9%
	HBP [13]	VGG-16	87.1%
	DFL-CNN [6]	ResNet50	87.4%
	Cross-X [47]	ResNet50	87.7%
	DCL [50]	ResNet50	87.8%
	TASN [8]	ResNet50	87.9%
	ACNet [48]	ResNet50	88.1%
	API-Net [51]	ResNet101	88.6%
PMG [43]	ResNet50	89.6%	
Joint	PART (Ours)	ResNet50	89.6%
	PART (Ours)	ResNet101	90.1%

C. Comparison With State of the Arts

In this subsection, we conduct experimental comparisons with state-of-the-art models on three representative benchmarks, *i.e.*, CUB-200-2011 [1], Stanford-Cars [3], and FGVC-Aircrafts [4].

1) *Comparison on CUB-200-2011 Dataset:* CUB-200-2011 is the most widely-used dataset, consisting of 200 bird species with visually similar appearances. The classification accuracy can be found in Tab. I. Following the aforementioned group ways in Section I, we divide the state-of-the-art models into two groups, *i.e.*, feature regularized learning and part-based learning. It can be found that the earlier works, *e.g.*, [33], [66] relies on the accurate part localization information, while achieving only 84.7% base performance. Hence with the development of auto-discovered part localization methods [23], [46], the classification accuracy has been increased by over 3% without any additional annotations.

On the other hand, feature-learning-based algorithms also show advantages in extracting informative knowledge and forming efficient representations. Hence our proposed PART approach is a joint learning algorithm with collaborative feature relation learning and part feature discovery. Combining these merits, our proposed method generates state-of-the-art results of 89.6% accuracy, which demonstrates the effectiveness of our proposed learning framework. Following API-Net [51], we extend our model with the deeper ResNet101 backbones, which achieves 90.1% and is 1.5% higher than [51].

TABLE II
COMPARISON WITH STATE-OF-THE-ART RESULTS ON FGVC-AIRCRAFT DATASET

Type	Method	Backbone	Accuracy
Part Based	RA-CNN [21]	VGG-19	88.2%
	MA-CNN [46]	VGG-19	89.9%
	NTSNet [23]	ResNet50	91.4%
Feature Based	Bilinear [5]	VGG-16	84.1%
	PC [52]	DenseNet-161	89.2%
	MaxEnt [11]	DenseNet-161	89.8%
	HBP [13]	VGG-16	90.3%
	ACNet [48]	VGG-16	91.5%
	DFL-CNN [6]	ResNet50	92.0%
	ACNet [48]	ResNet50	92.4%
	CIN [9]	ResNet50	92.6%
	Cross-X [47]	ResNet50	92.7%
	CIN [9]	ResNet101	92.8%
	DCL [50]	ResNet50	93.0%
	API-Net [51]	ResNet50	93.0%
	PMG [43]	ResNet50	93.4%
	API-Net [51]	ResNet101	93.4%
Joint	PART (Ours)	ResNet50	94.4%
	PART (Ours)	ResNet101	94.6%

2) *Comparison on FGVC-Aircraft Dataset:* Similar to the benchmarking on CUB-200-2011 dataset, here we make comparisons on sub-categories of aircrafts in Tab. II. Recent approaches, *e.g.*, DCL [50] and Cross-X [47], achieve performance of 93.0% and 92.7%, which is much higher than the previous work [5]. To make fair comparisons with these works, we adopt the lightweight ResNet-50 backbone to achieve the state-of-the-art performance of 94.4%, while previous works rely on deep backbones for feature extraction, *e.g.*, [9] with ResNet-101 and [11], [52] with DenseNet-161. In addition, other works [21], [23], [46] adopt the multi-crop inference to obtain multi-scale features, introducing additional computation burdens. Different from these mentioned issues, our proposed PART approach presents robust features with ResNet-50 backbone and one-crop inference.

3) *Comparison on Stanford-Cars Dataset:* Tab. III exhibits the results on the Stanford-Cars dataset. Stanford-Cars is a much easier dataset, in which previous methods achieve preliminary results of over 92.5% accuracy. It can be found that methods on this dataset perform very similar performance, *e.g.*, TASN [8] and HBP [13] of 93.7% performance. While obtaining multi-level feature representations [47] and using separate layer initialization strategies [6], high performance as 94.5% can be achieved. Surprisingly, even on this dataset, our method can provide a steady improvement compared to state-of-the-art results, reaching 95.1% performance. With deep backbones *i.e.*, ResNet101, the performance shows a slight improvement to 95.3% in top-1 accuracy.

D. Performance Analysis

In this subsection, we first conduct ablation studies of our proposed different modules, and study the manner of positional

TABLE III
COMPARISON WITH STATE-OF-THE-ART RESULTS ON STANFORD-CARS DATASET. †: ADDITIONAL BOUNDING BOX OR SEGMENTATION ANNOTATION

Type	Method	Backbone	Accuracy
Part Based	PA-CNN† [33]	VGG-19	92.8%
	MA-CNN [46]	VGG-19	92.8%
	NTSNet [23]	ResNet50	93.9%
Feature Based	Bilinear [5]	VGG-16	91.3%
	Grassmann [16]	VGG-16	92.8%
	PC [52]	DenseNet-161	92.9%
	MaxEnt [11]	DenseNet-161	93.0%
	MAMC [10]	ResNet101	93.0%
	HBP [13]	VGG-16	93.7%
	TASN [8]	ResNet50	93.7%
	DFL-CNN [6]	ResNet50	93.8%
	CIN [9]	ResNet50	94.1%
	CIN [9]	ResNet101	94.5%
	Cross-X [47]	ResNet50	94.5%
	DCL [50]	ResNet50	94.5%
	API-Net [51]	ResNet50	94.8%
	API-Net [51]	ResNet101	95.1%
PMG [43]	ResNet50	95.1%	
Joint	PART (Ours)	ResNet50	95.1%
	PART (Ours)	ResNet101	95.3%

TABLE IV
ABLATION STUDIES OF OUR DIFFERENT COMPONENTS ON THREE BENCHMARKS. $\mathcal{M}_{Relation}$, \mathcal{M}_{Part} DENOTES THE PROPOSED RELATION TRANSFORMATION MODULE AND PART-DISCOVERY MODULE RESPECTIVELY. QUATERNION (x, y_1, y_2, y_3) INDICATES THE HEAD NUMBER OF GLOBAL TRANSFORMER IS x AND HEADS OF THREE PARTS ARE y_1, y_2 AND y_3

$\mathcal{M}_{Relation}$	\mathcal{M}_{Part}	CUB-200-2011	FGVC-Aircraft	Stanford-Cars
N/A	N/A	85.4%	90.7%	93.4%
(2,0,0,0)	N/A	88.2%	93.5%	94.7%
(2,1,1,1)	×3	89.2%	94.2%	95.1%
(4,1,1,1)	×3	89.3%	94.3%	95.1%
(4,1,1,1)	×4	89.6%	94.4%	95.1%

encoding in relation to transformation modules. Then we make a detailed analysis of the part effects in our framework and finally make explanations using visualized comparisons.

1) *Ablation Studies on Different Components:* To evaluate the effectiveness of our proposed part discovery module and relation transformation module, we conduct detailed ablation studies on three representative benchmarks [1], [3], [4]. In Tab. IV, the first line shows the baseline performance on these three benchmarks, using the same data augmentation and training hyper-parameters as our final model. It can be found that our baseline model surpasses many earlier works and our proposed modules, *i.e.*, relation transformation module $\mathcal{M}_{Relation}$ and part discovery module \mathcal{M}_{Part} can improve the recognition ability steadily, *e.g.*, from 85.4% to 89.6% (4 parts) on CUB-200-2011 dataset. In the second row, we only adopt the relation transformation module with stacked manners, and set the head number of multi-head attention as 2.

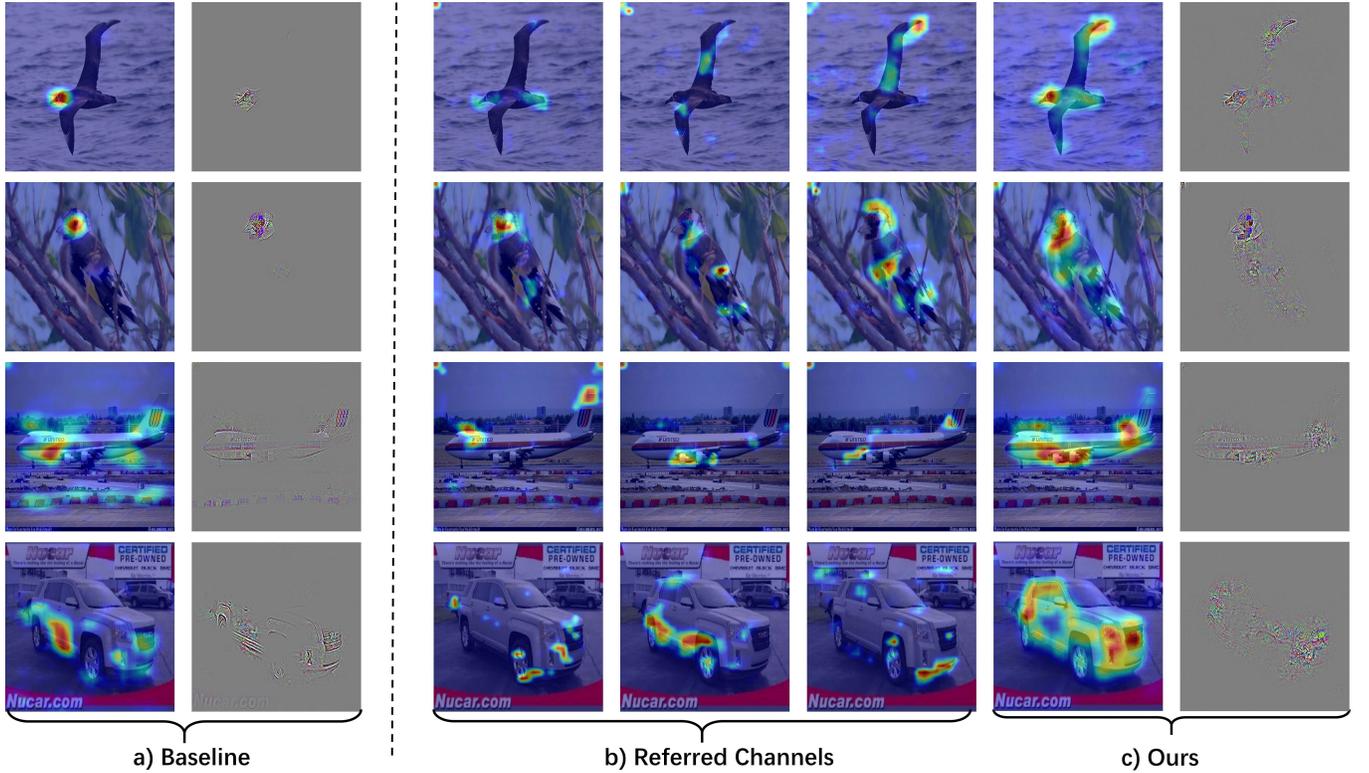


Fig. 4. Class activation visualizations of baseline (ResNet50) and our proposed PART. Left figures in a) and c) are the weighted summed class activation map of all channels. right figures in a) and c) are the guided gradient of baseline and our model generated by Grad-CAMs [68]. Images in b) are referred channels generated by our proposed PART.

TABLE V

PERFORMANCE ANALYSIS OF FEATURE EXTRACTION AND FEATURE ENCODING MANNER ON CUB-200-2011 BENCHMARK. $\mathcal{M}_{Relation}$. \mathcal{M}_{Part} DENOTES THE PROPOSED RELATION TRANSFORMATION MODULE AND PART-DISCOVERY MODULE RESPECTIVELY

Feature Extraction	Encoding Manner	Accuracy
Baseline (w/o group sampler)	N/A	84.2%
Baseline	N/A	85.4%
$\mathcal{M}_{Relation}$	N/A	87.5%
$\mathcal{M}_{Relation}$	Absolute Encoding	88.1%
$\mathcal{M}_{Relation}$	Learnable Encoding	88.2%
$\mathcal{M}_{Relation} + \mathcal{M}_{Part}$	Learnable Encoding	89.2%

It can be found that the performance has increased notably by incorporating the part-level features. This verifies that the effectiveness of our proposed module, even on the high-performance baseline. Moreover, when setting the number of attention heads as 4, performance on these three datasets can be slightly improved.

2) *Effects of Feature Encoding*: We first present the performance of the baseline method in the first two rows on the CUB benchmark dataset. The first row indicates that without the group sampler in Section IV-B of a class-wise balanced manner, baseline model would result in a performance drop of over 1%. As a natural drawback in building contextual dense relations, the spatial structure information is lost during

TABLE VI

HYPER PARAMETER EXPERIMENTS ON PART DISCOVERY MODULE ON CUB-200-2011 BENCHMARK. $|\mathcal{X}^*|$ AND $|\mathcal{P}|$ DENOTES THE NUMBER OF PARTS IN THE PART STACK AND FINAL SELECTED PARTS RESPECTIVELY. \mathcal{H} INDICATES THE HEAD NUMBER OF TRANSFORMERS

$ \mathcal{X}^* $	N/A	64	64	64	32	64	64	64
$ \mathcal{P} $	N/A	1	2	3	3	3	4	5
\mathcal{H}	2	2	2	2	2	4	4	4
Acc.	88.2%	88.8%	89.0%	89.2%	89.1%	89.3%	89.6%	88.8%

the vectorization process. The result (with 2 attention heads) without positional information can be found in the third-row of Tab. V. By incorporating the absolute positional encoding operation (sinusoid encoding function in [24]), the recognition accuracy has been increased from 87.5% to 88.1%. In this paper, we adopt relative encoding which uses learned encoding parameters to encode the unique positional information for each pixel. The result can be found in the fifth row, which also has the potential to simulate any convolutional layers as in Theorem 1. And finally our full model with the part discovery module achieves the best performance in the last row.

3) *Effects of Part Discovery*: To evaluate the effectiveness of the part discovery module, here we exhibit different hyper-parameters in the part discovery module in Tab. VI. In our proposed algorithm, the part selection is mainly affected by two hyper-parameters, *i.e.*, the capacity of the part stack $|\mathcal{X}^*|$

and the final selected part number $|\mathcal{P}|$. The selected part number also affects the network architecture, deciding the number of regularized part-level branches. When applying no auxiliary part branches, in Tab. VI, our proposed approach with only transformer modules achieves the performance of 88.2%. With the increase of selected part number $|\mathcal{P}|$, auxiliary supervision is also automatically added into the overall framework. It can be found that our approach with more parts ($|\mathcal{P}| = 4$) reaches a high performance of 89.6%. This indicates the effectiveness of using auxiliary part branches as regularization. However, with the part numbers $|\mathcal{P}|$ increasing to be 5, the final performance decreases to 88.8%. Discovering more parts would incorporate meaningless background regions into computation and also restrain the learning of global features, which jointly lead to an inferior response on class activation maps. Moreover, when squeezing the capacity of the part stack $|\mathbf{X}^*|$, the part discovery module would also lead to similar performance, which verifies the robustness of our proposed method.

During the inference phase, we only rely on the global branches and drop the redundant part features. We also conduct experiments of direct concatenating these branches rather than using them as regularization. The performance of condition $|\mathcal{P}| = 1$ in the first row of Tab. VI drops to 88.5%. Assuming the same classification task of 200 classes, the overall parameters of ResNet50 is 25.6M. Comparing to the full model with all 4 parts of 32.9M, our final model only requires 28.7M parameters for learning, which is implementation-friendly.

4) *Visualized Explanations*: We further investigate the class activation maps using state-of-the-art Grad-CAMs [68]. As shown in Fig. 4, we present the summation of all channels of baseline model (ResNet50) in Fig. 4 a), and our proposed PART in c). Interestingly, the baseline model can easily fall into local attentions to distinguish the visually similar samples, which leads to inferior generalization capabilities. While our proposed PART can localize the full object, resulting in a robust recognition ability. Moreover, we visualize the referred channels of our model in Fig. 4 b). It can be found that different channels focus on different semantic parts, forming the final robust representations. Although these “parts” do not strictly meet the definition of natural semantics, while still responding to meaningful object regions.

V. CONCLUSION

In this paper, we propose a novel PART-guided Relational Transformers (PART) framework for fine-grained recognition tasks. To solve the deficiencies in building long-term relationships in conventional CNNs, we make the first attempt to introduce the Transformer architecture into fine-grained visual recognition tasks. Beyond these insights, we further present a new part discovery module to automatically mine discriminative regions by utilizing the class activation maps during the training process. With these generated part regions, we propose a part-guided transformation module to learn high-order spatial relationships among semantic pixels. Our full model is composed of several local regional transformers and one global transformer, enhancing the discriminative regions while maintaining the contextual one. Experimental results

verify the effectiveness of our proposed modules and our proposed PART reaches new state-of-the-art on 3 widely-used fine-grained recognition benchmarks.

APPENDIX

PROOF OF THE THEOREM 1

This Appendix provides a brief proof of theorem 1, which is derived from previous research [65] and not the main contribution of our paper.

Theorem 2: A transformer layer with m heads and with hidden dimension C_h and output dimension C_{out} can simulate convolutional neural networks with $\min(C_h, C_{out})$ channels.

Proof: A typical transformer layer \mathcal{T} is composed of multi-head attention (MHA) layers with m heads and several feed forward layers (multi-layer perceptron). Considering Eqn. (7) and Eqn. (8), it can be formally presented as:

$$\begin{aligned} \mathcal{T} &= ((\text{MHA}(\mathbf{X})\mathbf{X}\mathbf{W}_{val} + \mathbf{b}^o)\mathbf{W}^f + \mathbf{b}^f) \\ &= (\text{Concat}_{h=1}^m (\sum_k \varphi(\mathbf{A}^{(h)})_k \mathbf{X}_k) \mathbf{W}_{val}^h + \mathbf{b}^o) \mathbf{W}^f + \mathbf{b}^f, \end{aligned} \quad (11)$$

where φ denotes the softmax operation and $\mathbf{W}^f, \mathbf{b}^f$ are learnable weights for feed-forward layers. In this equation, the learnable weights \mathbf{W}_{val}^h and \mathbf{W}^f can be formulated into one unified projection of \mathbf{W}_{proj}^h . In this manner, Eqn. (11) can be further reformulated as:

$$\mathcal{T} = \text{Concat}_{h=1}^m (\sum_k \varphi(\mathbf{A}^{(h)})_k \mathbf{X}_k) \mathbf{W}_{proj}^h + \mathbf{b}, \quad (12)$$

Considering a convolutional layer $\text{Conv}(\mathbf{X}) = \mathbf{X}_k \mathbf{W} + \mathbf{b}$, the transformer layer would degenerated to a convolutional layer, when the attention value $\varphi(\mathbf{A}^{(h)})_k$ equals to 1, when the positional biases $\Delta_{\{1,2\}}$ of convolutional kernel centers equals to the attention biases $k - q$. Considering the relative encoding in the main manuscript:

$$\begin{aligned} \mathbf{A}_{k,q} &= \mathbf{X}_{k,:}^\top \mathbf{W}_{qry}^\top \mathbf{W}_{key} \mathbf{X}_{q,:} + \mathbf{X}_{k,:}^\top \mathbf{W}_{qry}^\top \widehat{\mathbf{W}}_{key} \mathbf{R}_{k-q} \\ &\quad + \mathbf{u}^\top \mathbf{W}_{key}^\top \mathbf{X}_{q,:} + \mathbf{v}^\top \widehat{\mathbf{W}}_{key} \mathbf{R}_{k-q}, \end{aligned} \quad (13)$$

As proved by Cordonnier *et al.* [65], there exists an attention matrix $\varphi(\mathbf{A}^{(h)})_k$ with relative encoding to be 1 if the bias in convolutional kernels equals to $k - q$, otherwise $\varphi(\mathbf{A}^{(h)})_k = 0$. To construct this, we set $\mathbf{W}_{key} = \mathbf{W}_{qry} = 0$, thus Eqn. (13) only has the last term of $\mathbf{v}^\top \widehat{\mathbf{W}}_{key} \mathbf{R}_{k-q}$.

In the above expression, we set $\widehat{\mathbf{W}}_{key} = \mathbf{I}$ and assume $(\mathbf{A}_q^{(h)})_k = \mathbf{v}^\top \mathbf{R}_{k-q} = -\alpha \|k - q - \Delta\|^2 + \alpha c$, the softmax operation $\varphi(\mathbf{A}_q^{(h)})_k$ yields 1. α and c denotes the coefficient and constant respectively. Note that $\Delta_{\{1,2\}} \in \mathbb{Z}^2$ denote the biases of x-axis and y-axis direction.

With these values fixed, the limitation of softmax value $\varphi(\mathbf{A}^{(h)})_k$ could reach 1. On the other conditions ($k - q \neq \Delta$), its limitation would be zero. Hence theorem 1 can be concluded. ■

REFERENCES

- [1] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, “The caltech-UCSD birds-200–2011 dataset,” California Inst. Technol., CA, USA, Tech. Rep. CNS-TR-2011-001, 2011.

- [2] G. Van Horn *et al.*, "Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 595–604.
- [3] J. Krause, M. Stark, J. Deng, and L. Fei-Fei, "3D object representations for fine-grained categorization," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops*, Jun. 2013, pp. 554–561.
- [4] S. Maji, E. Rahtu, J. Kannala, M. Blaschko, and A. Vedaldi, "Fine-grained visual classification of aircraft," 2013, *arXiv:1306.5151*.
- [5] T.-Y. Lin, A. RoyChowdhury, and S. Maji, "Bilinear CNN models for fine-grained visual recognition," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1449–1457.
- [6] Y. Wang, V. I. Morariu, and L. S. Davis, "Learning a discriminative filter bank within a CNN for fine-grained recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 4148–4157.
- [7] L. Zhang, S. Huang, W. Liu, and D. Tao, "Learning a mixture of granularity-specific experts for fine-grained categorization," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 8331–8340.
- [8] H. Zheng, J. Fu, Z.-J. Zha, and J. Luo, "Looking for the devil in the details: Learning trilinear attention sampling network for fine-grained image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 5012–5021.
- [9] Y. Gao, X. Han, X. Wang, W. Huang, and M. Scott, "Channel interaction networks for fine-grained image categorization," in *Proc. AAAI Conf. Artif. Intell. (AAAI)*, 2020, pp. 10818–10825.
- [10] M. Sun, Y. Yuan, F. Zhou, and E. Ding, "Multi-attention multi-class constraint for fine-grained image recognition," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 805–821.
- [11] A. Dubey, O. Gupta, R. Raskar, and N. Naik, "Maximum-entropy fine grained classification," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2018, pp. 637–647.
- [12] Y. Gao, O. Beijbom, N. Zhang, and T. Darrell, "Compact bilinear pooling," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 317–326.
- [13] C. Yu, X. Zhao, Q. Zheng, P. Zhang, and X. You, "Hierarchical bilinear pooling for fine-grained visual recognition," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 574–589.
- [14] S. Kong and C. Fowlkes, "Low-rank bilinear pooling for fine-grained classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 365–374.
- [15] Y. Li, N. Wang, J. Liu, and X. Hou, "Factorized bilinear models for image recognition," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2079–2087.
- [16] X. Wei, Y. Zhang, Y. Gong, J. Zhang, and N. Zheng, "Grassmann pooling as compact homogeneous bilinear pooling for fine-grained visual classification," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Sep. 2018, pp. 355–370.
- [17] S. Huang, Z. Xu, D. Tao, and Y. Zhang, "Part-stacked CNN for fine-grained visual categorization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 1173–1182.
- [18] W. Ge, X. Lin, and Y. Yu, "Weakly supervised complementary parts models for fine-grained image classification from the bottom up," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 3034–3043.
- [19] M. M. Kalayeh, E. Basaran, M. Gokmen, M. E. Kamasak, and M. Shah, "Human semantic parsing for person re-identification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 1062–1071.
- [20] Z. Huang and Y. Li, "Interpretable and accurate fine-grained recognition via region grouping," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 8662–8672.
- [21] J. Fu, H. Zheng, and T. Mei, "Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4438–4446.
- [22] A. Recasens, P. Kellnhofer, S. Stent, W. Matusik, and A. Torralba, "Learning to zoom: A saliency-based sampling layer for neural networks," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Sep. 2018, pp. 51–66.
- [23] Z. Yang, T. Luo, D. Wang, Z. Hu, J. Gao, and L. Wang, "Learning to navigate for fine-grained classification," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Sep. 2018, pp. 420–435.
- [24] A. Vaswani *et al.*, "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5998–6008.
- [25] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2921–2929.
- [26] B. Yao, A. Khosla, and L. Fei-Fei, "Combining randomization and discrimination for fine-grained image categorization," in *Proc. CVPR*, Jun. 2011, pp. 1577–1584.
- [27] B. Yao, G. Bradschi, and L. Fei-Fei, "A codebook-free and annotation-free approach for fine-grained image categorization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2012, pp. 3466–3473.
- [28] S. Lazebnik, C. Schmid, and J. Ponce, "Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories," in *Proc. IEEE Comput. SoC. Conf. Comput. Vis. Pattern Recognit.*, vol. 2, Jun. 2006, pp. 2169–2178.
- [29] N. Zhang, J. Donahue, R. Girshick, and T. Darrell, "Part-based R-CNNs for fine-grained category detection," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*. Cham, Switzerland: Springer, 2014, pp. 834–849.
- [30] X. He and Y. Peng, "Weakly supervised learning of part selection model with spatial constraints for fine-grained image classification," in *Proc. 31st AAAI Conf. Artif. Intell.*, 2017, pp. 4075–4081.
- [31] X.-S. Wei, C.-W. Xie, J. Wu, and C. Shen, "Mask-CNN: Localizing parts and selecting descriptors for fine-grained bird species categorization," *Pattern Recognit.*, vol. 76, pp. 704–714, Apr. 2018.
- [32] B. He, J. Li, Y. Zhao, and Y. Tian, "Part-regularized near-duplicate vehicle re-identification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 3997–4005.
- [33] J. Krause, H. Jin, J. Yang, and L. Fei-Fei, "Fine-grained recognition without part annotations," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 5546–5555.
- [34] G. G. Abel, D. Modolo, and V. Ferrari, "Do semantic parts emerge in convolutional neural networks?" *Int. J. Comput. Vis.*, vol. 126, no. 5, pp. 476–494, May 2018.
- [35] T. Xiao, Y. Xu, K. Yang, J. Zhang, Y. Peng, and Z. Zhang, "The application of two-level attention models in deep convolutional neural network for fine-grained image classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 842–850.
- [36] M. Simon and E. Rodner, "Neural activation constellations: Unsupervised part model discovery with convolutional networks," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1143–1151.
- [37] M. Lam, B. Mahasseni, and S. Todorovic, "Fine-grained recognition as HSnet search for informative image parts," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2520–2529.
- [38] Y. Ding, Y. Zhou, Y. Zhu, Q. Ye, and J. Jiao, "Selective sparse sampling for fine-grained image recognition," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6599–6608.
- [39] Z. Wang, S. Wang, S. Yang, H. Li, J. Li, and Z. Li, "Weakly supervised fine-grained image classification via gaussian mixture model oriented discriminative learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 9749–9758.
- [40] X.-S. Wei, J.-H. Luo, J. Wu, and Z.-H. Zhou, "Selective convolutional descriptor aggregation for fine-grained image retrieval," *IEEE Trans. Image Process.*, vol. 26, no. 6, pp. 2868–2881, Jun. 2017.
- [41] Y. Peng, X. He, and J. Zhao, "Object-part attention model for fine-grained image classification," *IEEE Trans. Image Process.*, vol. 27, no. 3, pp. 1487–1500, Mar. 2017.
- [42] Y. Zhao, K. Yan, F. Huang, and J. Li, "Graph-based high-order relation discovery for fine-grained recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2021, pp. 15079–15088.
- [43] R. Du *et al.*, "Fine-grained visual classification via progressive multi-granularity training of jigsaw patches," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*. Cham, Switzerland: Springer, 2020, pp. 153–168.
- [44] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7794–7803.
- [45] D. Chang *et al.*, "The devil is in the channels: Mutual-channel loss for fine-grained image classification," *IEEE Trans. Image Process.*, vol. 29, pp. 4683–4695, 2020.
- [46] H. Zheng, J. Fu, T. Mei, and J. Luo, "Learning multi-attention convolutional neural network for fine-grained image recognition," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 5209–5217.
- [47] W. Luo *et al.*, "Cross-X learning for fine-grained visual categorization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 8242–8251.
- [48] R. Ji *et al.*, "Attention convolutional binary neural tree for fine-grained visual categorization," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 10468–10477.
- [49] O. M. Aodha, E. Cole, and P. Perona, "Presence-only geographical priors for fine-grained image classification," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 9595–9605.

- [50] Y. Chen, Y. Bai, W. Zhang, and T. Mei, "Destruction and construction learning for fine-grained image recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 5157–5166.
- [51] P. Zhuang, Y. Wang, and Y. Qiao, "Learning attentive pairwise interaction for fine-grained classification," in *Proc. 34th AAAI Conf. Artif. Intell. (AAAI)*, 2020, pp. 13130–13137.
- [52] A. Dubey, O. Gupta, P. Guo, R. Raskar, R. Farrell, and N. Naik, "Pairwise confusion for fine-grained visual classification," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Sep. 2018, pp. 70–86.
- [53] K. M. Choromanski *et al.*, "Rethinking attention with performers," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2020.
- [54] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*. Cham, Switzerland: Springer, 2020, pp. 213–229.
- [55] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, "Deformable DETR: Deformable transformers for end-to-end object detection," 2020, *arXiv:2010.04159*.
- [56] A. Dosovitskiy *et al.*, "An image is worth 16×16 words: Transformers for image recognition at scale," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2020.
- [57] Z. Liu *et al.*, "Swin transformer: Hierarchical vision transformer using shifted Windows," 2021, *arXiv:2103.14030*.
- [58] L. Yuan *et al.*, "Tokens-to-token ViT: Training vision transformers from scratch on ImageNet," 2021, *arXiv:2101.11986*.
- [59] A. Srinivas, T.-Y. Lin, N. Parmar, J. Shlens, P. Abbeel, and A. Vaswani, "Bottleneck transformers for visual recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 16519–16529.
- [60] E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Alvarez, and P. Luo, "SegFormer: Simple and efficient design for semantic segmentation with transformers," 2021, *arXiv:2105.15203*.
- [61] J. He *et al.*, "TransFG: A transformer architecture for fine-grained recognition," 2021, *arXiv:2103.07976*.
- [62] A. Rosenfeld and M. Thurston, "Edge and curve detection for visual scene analysis," *IEEE Trans. Comput.*, vol. C-20, no. 5, pp. 562–569, May 1971.
- [63] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2009, pp. 248–255.
- [64] Z. Dai, Z. Yang, Y. Yang, J. G. Carbonell, Q. Le, and R. Salakhutdinov, "Transformer-XL: Attentive language models beyond a fixed-length context," in *Proc. 57th Annu. Meeting Assoc. Comput. Linguistics*, 2019, pp. 2978–2988.
- [65] J.-B. Cordonnier, A. Loukas, and M. Jaggi, "On the relationship between self-attention and convolutional layers," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2020.
- [66] X. Liu, T. Xia, J. Wang, Y. Yang, F. Zhou, and Y. Lin, "Fully convolutional attention networks for fine-grained recognition," 2016, *arXiv:1603.06765*.
- [67] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [68] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 618–626.



Yifan Zhao received the B.E. degree from the Harbin Institute of Technology in July 2016. He is currently pursuing the Ph.D. degree with the State Key Laboratory of Virtual Reality Technology and System, School of Computer Science and Engineering, Beihang University. His research interests include computer vision and image understanding.



Jia Li (Senior Member, IEEE) received the B.E. degree from Tsinghua University in 2005 and the Ph.D. degree from the Institute of Computing Technology, Chinese Academy of Sciences, in 2011. He is currently a Full Professor with the School of Computer Science and Engineering, Beihang University. His research interests include computer vision and multimedia big data, especially the understanding and generation of visual contents.



Xiaowu Chen (Senior Member, IEEE) received the Ph.D. degree in computer science from Beihang University, Beijing, China, in 2001. His research interests include virtual reality, augmented reality, computer graphics, and computer vision.



Yonghong Tian (Senior Member, IEEE) received the Ph.D. degree from the Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China, in 2005. He is currently a Full Professor with the Department of Computer Science and Technology, Peking University, Beijing. He has authored or coauthored more than 160 technical articles in refereed journals and conferences. He has owned more than 57 Chinese and U.S. patents. His research interests include machine learning, computer vision, and multimedia big data.