# **Re-Attention Transformer for Weakly Supervised Object Localization**

Hui Su<sup>1</sup> suhui@zhejianglab.com Yue Ye<sup>1</sup> yeyue@zhejianglab.com Zhiwei Chen<sup>1</sup> chenzw@zhejianglab.com Mingli Song<sup>2</sup> brooksong@zju.edu.cn Lechao Cheng<sup>1\*</sup> chenglc@zhejianglab.com

<sup>1</sup> Zhejiang Lab, Hangzhou, China

<sup>2</sup>Zhejiang University, Hangzhou, China

#### Abstract

Weakly supervised object localization is a challenging task which aims to localize objects with coarse annotations such as image categories. Existing deep network approaches are mainly based on class activation map, which focuses on highlighting discriminative local region while ignoring the full object. In addition, the emerging transformer-based techniques constantly put a lot of emphasis on the backdrop that impedes the ability to identify complete objects. To address these issues, we present a reattention mechanism termed token refinement transformer (TRT) that captures the objectlevel semantics to guide the localization well. Specifically, TRT introduces a novel module named token priority scoring module (TPSM) to suppress the effects of background noise while focusing on the target object. Then, we incorporate the class activation map as the semantically aware input to restrain the attention map to the target object. Extensive experiments on two benchmarks showcase the superiority of our proposed method against existing methods with image category annotations. Source code is available in https://github.com/su-hui-zz/ReAttentionTransformer.

# **1** Introduction

Recently, there emerges a surge of interests in localizing activation map with weakly supervision, for example with only image-level annotations, arising in different downstream applications. Weakly supervised learning has received a lot of attention since it minimizes the demand for fine-grained annotations so as to involve less human labeling effort. This work attempts to localize objects with only image category supervision, also termed weakly supervised object localization (WSOL).

The dramatic progress of deep neural networks (DNNs) also provide tremendous benefits with impressive gains in the task of WSOL. The cornerstone work is the Class Activation



# Figure 1: The comparisons between our proposed token refinement transformer (TRT) and existing approaches.

Mapping(CAM)[] that enables classification neural networks to highlight the part of interest by computing a weighted sum of the last convolutional feature maps. However, CAM focuses only on the most discriminative features instead of the entire object areas due to its locality characteristic [10, 12, 124]. To alleviate this issue, much efforts have been made, such as adversarial erasing [22, 12, 10], multi-task joint training [13], novel networks [1, 52] and divergent activation [3]. Most of these CAM-based approaches first compute local features and proceed to acquire adjacent object pieces to embrace the whole target. However, the underlying limitations of CAM still exist despite the fact that these approaches do encourage a better activation region. Another popular alternative is to leverage the long-range dependencies from vision transformer [1] to find the object location. But this kind of approaches always suffer from irrelevant background for the reasons that canonical transformer-based methods generate patch tokens by splitting an image into a series of ordered patches and compute global relations in each layer. Therefore Gao et al. propose TS-CAM[II] by integrating both transformer and CAM to mitigate those drawbacks. As illustrated in Figure 1, CAM accentuates the most discriminative local region while transformer distracts attention from target to the background area. The results in penultimate column showcase that current ensemble approach can not deal with those issues even along with promising performance.

In this paper, we propose a re-attention strategy based on token refinement transformer (TRT) to grasp objects of interest more precisely. Specifically, TRT introduces a novel module named token priority scoring module (TPSM) to suppress the effects of background noise in transformer. We proceed to incorporate the class activation map as the semantically aware guidance to restrain the attention map to the target object. The TPSM attempts to re-score the important regions of strong response obtained by attention map in transformer to minimize the impact of jumbled background as much as possible. In addition, we also reveal that adaptive thresholding based on sampling over cumulative distribution function is superior to the regular thresholding or topK strategy. The contributions are summarized as follows:

• we propose a re-attention mechanism termed token refinement transformer (TRT) which highlights the precise object of interest.

- we propose an adaptive thresholding strategy based on sampling over cumulative importance that improves the performance significantly in the task of WSOL.
- The experimental results show convincing results of both qualitative and quantitative compared to existing approaches on ILSVRC[5] and CUB-200-2011[5].

# 2 Related Work

### 2.1 CAM-based Attention Map Generation

The milestone work of the class activation map (CAM [53]) aims to localize the discriminative regions for a specific category. In this study, we grouped CAM techniques into two categories: activation-based approaches [28, 50, 40, 51] and gradient-based approaches [5, 13, 22, 29, 13]. Typical activation-based works dispose of the reliance on the gradient as they only require access to class activation map and the channel-wise confidence. For example, Wang et al. [1] propose Score-CAM that computes the importance weights with the channel-wise increase of confidence performed on the activation map. Further, an extended version named SS-CAM [III] provides better post-hoc explanations about the centralized target object in the image using smoothing to obtain a better understanding. In addition, other popular works, such as integrated score CAM (IS-CAM) [23], attempt to alleviate the issue of saturation and false confidence and also achieve significant results. Apart from activation-based approaches, researchers often employ gradient information to enhance the salience of the image. Gradient-based methods highlight the important regions based on the back-propagated gradients of the specific class. However, the naive Grad-CAM [1] often fails to localize objects in an image when multiple occurrences of the same category exist. Chattopadhyay et al. propose Grad-CAM++ [] to overcome above issue by introducing pixel-wise weighting of the gradients of the output with respect to a particular region in the final feature map. Moreover, the work of smooth Grad-CAM++ both smoothing and Grad-CAM++ [1] to enable either visualizing a layer, subset of feature maps, or subset of neurons within a feature map. Many additional techniques, including XGrad-CAM [1], LayerCAM [2] and others, have also contributed to the field's progress and delivered favorable results.

### 2.2 Transformer-based Attention Map Generation

 solid numerical results have been achieved, the highlight regions still bring issues derived from both CAM and Transformer.

### 2.3 Weakly Supervised Object Localization

# 3 Method

In this section, we first revisit the preliminaries for vision transformer [1]. Then we discuss the details of the proposed token refinement transformer (TRT) for WSOL.



Figure 2: Token Refinement Transformer (TRT) framework. TRT consists of two branches, Token Priority Scoring Module (TPSM) and CAM, respectively. TPSM attempts to generate context-aware features  $\mathbf{M}_T$  that contribute most to the target class. CAM is introduced to obtain discriminative features  $\mathbf{M}_C$ . We finally get the attention map  $\mathbf{M}$  by performing element-wise multiplication as  $\mathbf{M} = \mathbf{M}_C \odot \mathbf{M}_T$ .

### 3.1 Preliminary

Let us consider  $I \in \mathbb{R}^{W \times H \times 3}$  as an input image, we split the image I based on the patch size P, resulting in N ( $N = [\frac{W}{P}] \times [\frac{H}{P}]$ ) non-overlap flattened patch blocks  $\mathbf{x}_p \in \mathbb{R}^{N \times (3 \times P^2)}$ . Each patch block  $\mathbf{x}_p^n (n \in \{1, ..., N\})$  is linearly projected into D dimensional patch embedding before being fed into transformer block. As part of the embeddings, an extra learnable class token  $\mathbf{z}_0^{cl_s} \in \mathbb{R}^{1 \times D}$  is introduced. In addition, we incorporate a position embedding  $\mathbf{E}_{pos} \in \mathbb{R}^{(N+1) \times D}$  to form the whole patch embedding for transformer as follows:

$$\mathbf{Z}_0 = [\mathbf{z}_0^{cls}; E(\mathbf{x}_p^1); E(\mathbf{x}_p^2); ...; E(\mathbf{x}_p^N)] + \mathbf{E}_{pos}$$
(1)

Where E(.) indicates patch embedding projection and  $\mathbf{Z}_0 \in \mathbb{R}^{(N+1) \times D}$  is the input for the first transformer block.

### 3.2 Overview

We define  $\mathbf{Z}_l \in \mathbb{R}^{(N+1) \times D}$   $(l \in \{1, ..., L\})$  as the output feature embedding for *l*-th transformer block. As illustrated in Figure 2, the output of the penultimate transformer  $\mathbf{Z}_{L-1}$  is fed into two branch, one of which is Token Priority Scoring Module (TPSM) that aims to re-attention patch tokens with the proposed adaptive thresholding strategy while the other branch compute standard class activation map.

In training stage, we evaluate the inconsistency between the output and ground truth in both two branches with commonly used cross entropy loss. Let  $\mathbf{z}_{L-1}^c \in \mathbb{R}^{1 \times D}$  and  $\mathbf{z}_{L-1}^p \in \mathbb{R}^{N \times D}$  be the output class token and patch token at the penultimate layer, respectively. For CAM branch, we reshape  $\mathbf{z}_{L-1}^p \in \mathbb{R}^{\sqrt{N} \times \sqrt{N} \times D}$ , which serves as the effective input for the subsequent convolution layer. Further, the output feature is globally average pooled followed by a softmax layer to get classification prediction  $\mathbf{p}^c$ . For TPSM branch,  $\mathbf{Z}_{L-1}(=[\mathbf{z}_{L-1}^c;\mathbf{z}_{L-1}^p])$  is subsequent processed with the re-attention module to obtain classification probability  $\mathbf{p}^t$ . The loss function is thus defined as:

$$L_{ce} = -\sum_{k=1}^{K} y_k (\log p_k^c + \log p_k^t)$$
(2)

where  $L_{ce}$  denotes cross entropy loss function. *K* is the number of category.  $y_k$  is the binary indicator (0 or 1) if class label *k* is the correct classification for the observation.  $p_k^c$  and  $p_k^t$  denote the output probability for class *k* in CAM and TPSM branch, respectively.

During testing, we first forward  $\mathbf{z}_{L-1}^p$  into TPSM to obtain context-aware feature maps  $\mathbf{M}_T \in \mathbb{R}^{\sqrt{N} \times \sqrt{N}}$  by performing a re-attention operation on patch tokens. More details will be discussed in Section 3.3. For the counterpart, we apply standard CAM to generate class specific activation maps  $\mathbf{M}_C \in \mathbb{R}^{K \times \sqrt{N} \times \sqrt{N}}$ . Consequently, attention maps can be obtained by:

$$\mathbf{M} = \mathbf{M}_T \odot \mathbf{M}_C \tag{3}$$

where  $\odot$  denotes an element-wise multiplication operation. The proposed context-aware feature map  $\mathbf{M}_T \in \mathbb{R}^{\sqrt{N} \times \sqrt{N}}$  is class-agnostic so we incorporate the class-aware activation map  $\mathbf{M}_C \in \mathbb{R}^{K \times \sqrt{N} \times \sqrt{N}}$  as the semantically guidance to restrain the attention map to the target class. We further resize the maps  $\mathbf{M}$  to the same size of the original images by bilinear interpolation. Specifically, we separate the foreground from the background using a defined

threshold as described in  $[\square]$ . Then we look for the tight bounding boxes that enclose the most linked region in the foreground pixels. Finally with the grid search approach, the thresholds for obtaining bounding boxes are updated to the optimal values.

### 3.3 Token Priority Scoring Module

Our token priority scoring module(TPSM) consists of three components as indicated in Figure 2. Firstly, we generate a preliminary attention map by exploiting long-range dependencies of class token and patch tokens over transformer blocks. Then an adaptive thresholding strategy is introduced to screen out patch tokens with high response in preliminary attention map. Finally, we perform re-attention operation on the selected tokens to capture more effective global relationships. We will detail the descriptions in following three subsections.

#### 3.3.1 Token Preliminary Attention

Multi-head self-attention(MHSA) [ $\square$ ] is widely used in transformer to model the long range dependency. We first compute the preliminary self-attention for *l*-th transformer block as:

$$\mathbf{A}_{l} = softmax(\frac{\mathbf{Q}_{l}\mathbf{K}_{l}^{\mathrm{T}}}{\sqrt{D}}) \tag{4}$$

where  $\mathbf{Q}_l$  and  $\mathbf{K}_l$  are the query and key representations projected from the previously output  $\mathbf{Z}_{l-1}$ . *D* indicates dimension of patch embeddings. **T** is a transpose operator.

We later investigate the characteristic of  $\mathbf{A}_l \in \mathbb{R}^{(1+N)\times(1+N)}$ . It is found that attention vector in the first row, which records dependency of class token to patch tokens, is driven to highlight object regions when Eq.2 is optimized. We aggregate these attention vectors over transformer blocks as  $\mathbf{m} = \sum_{l=1}^{L-1} \mathbf{A}_l[0, 1:]$ .  $\mathbf{m} \in \mathbb{R}^{1\times N}$  is then reshaped back to  $\mathbf{M}_m \in \mathbb{R}^{\sqrt{N} \times \sqrt{N}}$  as a preliminary attention map as illustrated in Figure 2.

#### 3.3.2 Token Selection Strategy

Token preliminary attention considers the cumulative dependency between patch tokens and class tokens based on the attention map of multi-head self-attention. As shown in Figure 2, we are suggested to suppress the irreverent background response to highlight the objective regions. Intuitively, we may pick up the top k largest responses or responses that exceed a fixed threshold  $\tau$ . However we experimentally find that these two basic strategies do not work well in practice. Thus we propose an adaptive thresholding strategy based on sampling over cumulative importance that boosts the performance significantly.

We first calculate the cumulative distribution function F of **m** and define strictly monotone transformation  $\mathbb{T}: U \sim [0, 1] \mapsto \mathbb{R}$  as the inverse function, thus

$$F(x) = \mathbf{P}_r(\mathbf{m} < x) = \mathbf{P}_r(\mathbb{T}(U) < x) = \mathbf{P}_r(U < \mathbb{T}^{-1}(x)) = \mathbb{T}^{-1}(x)$$
(5)

 $\mathbf{P}_r$  is the probability function. *F* is considered as the inverse function of  $\mathbb{T}$ , or  $\mathbb{T}(u) = F^{-1}(u), u \sim [0, 1]$ . Specifically, we first sort the values in **m** from high to low and calculate the cumulative attention. Then the adaptive threshold  $\tau'$  is obtained based on the inverse transform sampling given the contribution *u* over cumulative attention. More theoretical details about adaptive threshold generation can be referred to "Inverse Transform Sampling." *u* controls the proportion of selected token attention in token preliminary attention. During

the training and inference stages,  $\tau'$  varies adaptively for different images and can achieve good results for objects with different scales. Therefore, we are capable of generating adaptive thresh  $\tau'$  from  $F^{-1}(u)$ . We denote  $\mathbf{b} = [\mathbf{m} > \tau']$  as the binary mask for the existence of selected patch tokens.

#### 3.3.3 Token Re-Attention

Let  $\mathbf{I}_N$  be the identity matrix with size *N*. To achieve the goal of drawing more attention to class-specific objectiveness instead of background, we generate the selection matrix  $\mathbf{B} \in \mathbb{R}^{N \times N}$  for token re-attention as:

$$\mathbf{B} = \mathbf{J} \otimes \mathbf{b} + \mathbf{J} \otimes (\mathbf{J}^{\mathrm{T}} - \mathbf{b}) \odot \mathbf{I}_{N}$$
(6)

 $\mathbf{J} \in \mathbb{R}^{N \times 1}$  is a matrix where every element is equal to one.  $\otimes$  means tensor product. **B** is a binary matrix where each entry  $\mathbf{B}_{i,j}$  means the *j*-th token will contribute to the update of the *i*-th token. We replace self-attention modules in transformer block by masked self-attention modules as follows:

$$\mathbf{S} = \frac{\mathbf{Q}_{L-1}\mathbf{K}_{L-1}^{\mathrm{T}}}{\sqrt{D}} \tag{7}$$

$$\mathbf{A}_{ij}^{r} = \frac{\exp(\mathbf{S}_{ij}) * \mathbf{B}_{ij}}{\sum_{k=1}^{N} \exp(\mathbf{S}_{ik}) * \mathbf{B}_{ik}}$$
(8)

Patch tokens  $\mathbf{z}_{L-1}^{p}$  are fed to mask transformer block followed by a fully connected layer and a mask softmax layer, resulting in importance weights  $\lambda$ . For training stage, the fusion embedding is generated by computing a weighted sum of importance weights  $\lambda$  with patch tokens  $\mathbf{z}_{L-1}^{p}$ . We further concatenate class embedding and fusion embedding to feed into final transformer block to yield classification loss. For inference stage, we retrieve the weights from the original relation **m** for pruned tokens. Hence, the re-attention vector is then defined as:

$$r = \frac{\sum_{k=1}^{N} \mathbf{m}_k * \mathbf{b}_k}{\sum_{k=1}^{N} \lambda_k}$$
(9)

$$\mathbf{m}' = \mathbf{m} \odot (\mathbf{J}^{\mathbf{T}} - \mathbf{b}) + \lambda * r$$
(10)

We further reshape  $\mathbf{m}' \in \mathbb{R}^N$  to yield context-aware feature map  $\mathbf{M}_{\mathbf{T}} \in \mathbb{R}^{\sqrt{N} \times \sqrt{N}}$ .

## **4** Experiments

### 4.1 Experimental Settings

Our method is implemented under PyTorch with four 32GB V100 GPUs. Adam is used for optimization, and the weight decay is set to  $5 \times 10^{-4}$ . The basic backbone used is the Deit-Base[52] pre-trained on ImageNet-1K[53]. We evaluate our proposed framework on two commonly used benchmark datasets (ILSVRC[53] and CUB-200-2011[53]) in WSOL. For CUB-200-2011[53], each input image is resized to 256 and randomly cropped to 224. The batch size is 128, and the adaptive uniform threshold is 0.65. We train the backbone

Methods	Backbone	Loc.Acc		
		Top-1	Top-5	Gt-Known
CAM[	VGG16	44.2	52.2	56.0
SPG[1]	VGG16	48.9	57.2	58.9
SLT-Net[	VGG16	67.8	-	87.6
CAM[53]	InceptionV3	41.1	50.7	55.1
SPG[51]	InceptionV3	46.7	57.2	-
I2C[52]	InceptionV3	66.0	68.3	72.6
SLT-Net[	InceptionV3	66.1	-	86.5
TS-CAM[	Deit-B	75.8	84.1	86.6
TS-CAM*[	Deit-B-384	77.8	88.6	90.8
TRT(Ours)	Deit-B	76.5	88.0	91.1
TRT(Ours)	Deit-B-384	80.5	91.7	94.1

Table 1: Experimental results on CUB-200-2011 for metrics of Loc.Acc.

and TPSM branches concurrently for 30 epochs, then fix the backbone and TPSM branch parameters and train the CAM branch for additional 15 epochs. For ILSVRC[ $\square$ ], we follow [ $\square$ ] to employ AutoAugment [ $\square$ ] on the inputs. Similar to CUB-200-2011, we also train the backbone and TPSM branches in first 12 epochs and proceed to tune CAM branch for 10 epochs. *u* is set to 0.95 for ILSVRC. \* means we reproduce the results based on the provided codes unless otherwise stated. We adopt the evaluation metrics as previous works :

- Gt-Known *Loc.Acc*[**D**] is the localization accuracy with ground-truth class that provides a positive value when Intersection over Union(IOU) between the predicted box and ground truch box is larger than 50%.
- Top-1/Top-5 *Loc.Acc*[51] means Top-1/Top-5 localization accuracy that fulfills: 1) the predicted class is correct. 2). the IOU between the predicted box and ground truch box is larger than 50%.
- MaxBoxAccV2 [8] averages the localization accuracy across IOU thresholds in [30%, 50%, 70%], regardless of classification results.

Methods	Backbone	MaxBoxAccV2	
CAM[53]	VGG	63.70	
ADL[	VGG	66.30	
VITOL[	Deit-B	73.17	
TS-CAM*[	Deit-B	76.74	
TRT(Ours)	Deit-B	82.08	

Table 2: Experimental results on CUB-200-2011 for MaxBoxAccV2.

### 4.2 Comparisons to the State-of-the-arts

Table 1 and Table 2 showcase the competitive results of our proposed TRT framework on the CUB-200-2011. Overall, our approach achieves remarkable performance across all metrics

Methods	Backbone	Loc.Acc		
		Top-1	Top-5	Gt-Known
CAM[53]	VGG16	42.8	54.9	59.0
ADL[	VGG16	44.9	-	-
SLT-Net[🔼]	VGG16	51.2	62.4	67.2
CAM[53]	InceptionV3	46.3	58.2	62.7
ADL[	InceptionV3	48.7	-	-
SLT-Net[	InceptionV3	55.7	65.4	67.6
TS-CAM*[	Deit-B	47.8	60.0	64.4
LCTR*[	Deit-B	53.4	63.9	67.1
TRT(ours)	Deit-B	58.8	68.3	70.7

Table 3: Experimental results on ILSVRC for metrics of Loc.Acc.

compared to existing methods. Specifically, TRT outperforms state-of-the-art methods with a large margin in Gt-Known, yielding a localization accuracy of 91.1% against 86.6%. It can also be observed that the benefits over other approaches stand out more when higher quality images are used. As for ILSVRC, Table 3 demonstrated that TRT is superior to both existing CAM-based and transformer-based approaches. We have reproduced the results of [**G**] and [**IG**] on ILSVRC with backbones Deit-S and Deit-B based on the released codes. It shows that the results of [**G**] and [**IG**] based on Deit-B are even worse than Deit-S, although Deit-B has much more capacity. In contrast, our results on ILSVRC outperform [**G**] and [**IG**] on both Deit-S and Deit-B. And we obtain better performances based on Deit-B than those from Deit-S for our proposed approach. A reasonable explanation may be that inductive bias (the re-attention constraint, etc.) allows us to better use the model capacity in the vision transformer. Figure 3 illustrates the qualitative results. When compared to CAM-based and Transformer-based techniques, our produced attention maps show visual superiority that focuses more on the objects.



Figure 3: Visualization of localization maps on CUB-200-2011 and ILSVRC datasets. Red means ground truth and green means predicted bounding box.

95







(b) The impact of uniform distribution





## 4.3 Ablation Study

In this section, we conduct ablation studies to verify the effectiveness of our proposed approach. First, we investigate the impact for different token selection strategies as shown in figure 4 (a). The experimental results emphasize the necessity of the proposed adaptive thresholding method. Figure 4 (b) empirically suggest us the default setting about u(=0.65) on CUB-200-2011. In addition, we also did the ablations about the re-attention module. Both qualitative and quantitative results as shown in figure 5 reveal that the aforementioned re-attention block plays a critical role in framework.

# 5 Conclusion

In this work, we introduce a Re-Attention method called the token refinement transformer (TRT) that captures object-level semantics for the task of WSOL. To reduce the impact of background noise while concentrating on the target object, TRT introduces the token priority scoring module (TPSM). Then, we integrate the class activation map as the guidance to amend the produced context-aware feature maps. Extensive experiments on two benchmarks demonstrate that our proposed approach outperforms existing methods.

# Acknowledgments

This work is partially supported by the National Natural Science Foundation of China (Grant No. 62106235), by the Exploratory Research Project of Zhejiang Lab(2022PG0AN01), by the Zhejiang Provincial Natural Science Foundation of China (LQ21F020003), by Open Research Projects of Zhejiang Lab(NO. 2019KD0AD01/018).

# References

- [1] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4981–4990, 2018.
- [2] Wonho Bae, Junhyug Noh, and Gunhee Kim. Rethinking class activation mapping for weakly supervised object localization. In *European Conference on Computer Vision*, pages 618–634. Springer, 2020.
- [3] Ricardo Cabral, Fernando De la Torre, Joao Paulo Costeira, and Alexandre Bernardino. Matrix completion for weakly-supervised multi-label image classification. *IEEE transactions on pattern analysis and machine intelligence*, 37(1):121–135, 2014.
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020.
- [5] Aditya Chattopadhay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In 2018 IEEE winter conference on applications of computer vision (WACV), pages 839–847. IEEE, 2018.
- [6] Zhiwei Chen, Changan Wang, Yabiao Wang, Guannan Jiang, Yunhang Shen, Ying Tai, Chengjie Wang, Wei Zhang, and Liujuan Cao. Lctr: On awakening the local continuity of transformer for weakly supervised object localization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 410–418, 2022.
- [7] Junsuk Choe and Hyunjung Shim. Attention-based dropout layer for weakly supervised object localization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 2219–2228, 2019.
- [8] Junsuk Choe, Seong Joon Oh, Sanghyuk Chun, Seungho Lee, Zeynep Akata, and Hyunjung Shim. Evaluation for weakly supervised object localization: Protocol, metrics, and datasets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [9] Ramazan Gokberk Cinbis, Jakob Verbeek, and Cordelia Schmid. Weakly supervised object localization with multi-fold multiple instance learning. *IEEE transactions on pattern analysis and machine intelligence*, 39(1):189–203, 2016.
- [10] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the*

*IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 113–123, 2019.

- [11] Zhigang Dai, Bolun Cai, Yugeng Lin, and Junying Chen. Up-detr: Unsupervised pretraining for object detection with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1601–1610, 2021.
- [12] Thomas Deselaers, Bogdan Alexe, and Vittorio Ferrari. Localizing objects while learning their appearance. In *European conference on computer vision*, pages 452–466. Springer, 2010.
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020.
- [14] Chaowei Fang, Dingwen Zhang, Liang Wang, Yulun Zhang, Lechao Cheng, and Junwei Han. Cross-modality high-frequency transformer for mr image super-resolution. arXiv preprint arXiv:2203.15314, 2022.
- [15] Ruigang Fu, Qingyong Hu, Xiaohu Dong, Yulan Guo, Yinghui Gao, and Biao Li. Axiom-based grad-cam: Towards accurate visualization and explanation of cnns. arXiv preprint arXiv:2008.02312, 2020.
- [16] Wei Gao, Fang Wan, Xingjia Pan, Zhiliang Peng, Qi Tian, Zhenjun Han, Bolei Zhou, and Qixiang Ye. Ts-cam: Token semantic coupled attention map for weakly supervised object localization. In *Proceedings of the IEEE/CVF International Conference* on Computer Vision, pages 2886–2895, 2021.
- [17] Benjamin Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. Levit: a vision transformer in convnet's clothing for faster inference. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12259–12269, 2021.
- [18] Guangyu Guo, Junwei Han, Fang Wan, and Dingwen Zhang. Strengthen learning tolerance for weakly supervised object localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7403–7412, 2021.
- [19] Saurav Gupta, Sourav Lakhotia, Abhay Rawat, and Rahul Tallamraju. Vitol: Vision transformer for weakly supervised object localization. *arXiv preprint arXiv:2204.06772, 2022.*
- [20] Yanbin Hao, Shuo Wang, Pei Cao, Xinjian Gao, Tong Xu, Jinmeng Wu, and Xiangnan He. Attention in attention: Modeling context correlation for efficient video classification. *IEEE Transactions on Circuits and Systems for Video Technology*, 2022.
- [21] Ju He, Jie-Neng Chen, Shuai Liu, Adam Kortylewski, Cheng Yang, Yutong Bai, Changhu Wang, and Alan Yuille. Transfg: A transformer architecture for fine-grained recognition. arXiv preprint arXiv:2103.07976, 2021.
- [22] Peng-Tao Jiang, Chang-Bin Zhang, Qibin Hou, Ming-Ming Cheng, and Yunchao Wei. Layercam: Exploring hierarchical class activation maps for localization. *IEEE Transactions on Image Processing*, 30:5875–5888, 2021.

- [23] M Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. Advances in neural information processing systems, 23, 2010.
- [24] Jungbeom Lee, Eunji Kim, and Sungroh Yoon. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4071– 4080, 2021.
- [25] Dong Li, Jia-Bin Huang, Yali Li, Shengjin Wang, and Ming-Hsuan Yang. Weakly supervised object localization with progressive domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3512–3520, 2016.
- [26] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021.
- [27] Jinjie Mai, Meng Yang, and Wenfeng Luo. Erasing integrated learning: A simple yet effective approach for weakly supervised object localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8766–8775, 2020.
- [28] Rakshit Naidu, Ankita Ghosh, Yash Maurya, Soumya Snigdha Kundu, et al. Is-cam: Integrated score-cam for axiomatic-based explanations. *arXiv preprint arXiv:2010.03023*, 2020.
- [29] Daniel Omeiza, Skyler Speakman, Celia Cintas, and Komminist Weldermariam. Smooth grad-cam++: An enhanced inference level visualization technique for deep convolutional neural network models. *arXiv preprint arXiv:1908.01224*, 2019.
- [30] Harish Guruprasad Ramaswamy et al. Ablation-cam: Visual explanations for deep convolutional network via gradient-free localization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 983–991, 2020.
- [31] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115 (3):211–252, 2015.
- [32] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [33] Yunhang Shen, Rongrong Ji, Yan Wang, Yongjian Wu, and Liujuan Cao. Cyclic guidance for weakly supervised joint detection and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 697–707, 2019.

- [34] Krishna Kumar Singh and Yong Jae Lee. Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In 2017 IEEE international conference on computer vision (ICCV), pages 3544–3553. IEEE, 2017.
- [35] Zhiqing Sun, Shengcao Cao, Yiming Yang, and Kris M Kitani. Rethinking transformerbased set prediction for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3611–3620, 2021.
- [36] Peng Tang, Xinggang Wang, Xiang Bai, and Wenyu Liu. Multiple instance detection network with online instance classifier refinement. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2843–2851, 2017.
- [37] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347– 10357. PMLR, 2021.
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [39] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- [40] Haofan Wang, Rakshit Naidu, Joy Michael, and Soumya Snigdha Kundu. Sscam: Smoothed score-cam for sharper visual feature localization. *arXiv preprint arXiv:2006.14255*, 2020.
- [41] Haofan Wang, Zifan Wang, Mengnan Du, Fan Yang, Zijian Zhang, Sirui Ding, Piotr Mardziel, and Xia Hu. Score-cam: Score-weighted visual explanations for convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer* vision and pattern recognition workshops, pages 24–25, 2020.
- [42] Jun Wang, Xiaohan Yu, and Yongsheng Gao. Feature fusion vision transformer for fine-grained visual categorization. arXiv preprint arXiv:2107.02341, 2021.
- [43] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 568–578, 2021.
- [44] Yunchao Wei, Jiashi Feng, Xiaodan Liang, Ming-Ming Cheng, Yao Zhao, and Shuicheng Yan. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1568–1576, 2017.
- [45] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. Advances in Neural Information Processing Systems, 34, 2021.
- [46] Haolan Xue, Chang Liu, Fang Wan, Jianbin Jiao, Xiangyang Ji, and Qixiang Ye. Danet: Divergent activation for weakly supervised object localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6589–6598, 2019.

14

- [47] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zi-Hang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proceedings of the IEEE/CVF International Conference* on Computer Vision, pages 558–567, 2021.
- [48] Dingwen Zhang, Junwei Han, Gong Cheng, and Ming-Hsuan Yang. Weakly supervised object localization and detection: a survey. *IEEE transactions on pattern analysis and machine intelligence*, 2021.
- [49] Hao Zhang, Lechao Cheng, Yanbin Hao, and Chong-Wah Ngo. Long-term leap attention, short-term periodic shift for video classification. *arXiv preprint arXiv:2207.05526*, 2022.
- [50] Xiaolin Zhang, Yunchao Wei, Jiashi Feng, Yi Yang, and Thomas S Huang. Adversarial complementary learning for weakly supervised object localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1325–1334, 2018.
- [51] Xiaolin Zhang, Yunchao Wei, Guoliang Kang, Yi Yang, and Thomas Huang. Selfproduced guidance for weakly-supervised object localization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 597–613, 2018.
- [52] Xiaolin Zhang, Yunchao Wei, and Yi Yang. Inter-image communication for weakly supervised localization. In *European Conference on Computer Vision*, pages 271–287. Springer, 2020.
- [53] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.