

ViTOL: Vision Transformer for Weakly Supervised Object Localization

Saurav Gupta

Sourav Lakhotia

Abhay Rawat

Rahul Tallamraju

Mercedes-Benz Research and Development India

{saurav.gupta, sourav.lakhotia, abhay.rawat, rahul.tallamraju}@mercedes-benz.com

Abstract

Weakly supervised object localization (WSOL) aims at predicting object locations in an image using only image-level category labels. Common challenges that image classification models encounter when localizing objects are, (a) they tend to look at the most discriminative features in an image that confines the localization map to a very small region, (b) the localization maps are class agnostic, and the models highlight objects of multiple classes in the same image and, (c) the localization performance is affected by background noise. To alleviate the above challenges we introduce the following simple changes through our proposed method ViTOL. We leverage the vision-based transformer for self-attention and introduce a patch-based attention dropout layer (p-ADL) to increase the coverage of the localization map and a gradient attention rollout mechanism to generate class-dependent attention maps. We conduct extensive quantitative, qualitative and ablation experiments on the ImageNet-1K and CUB datasets. We achieve state-of-the-art MaxBoxAcc-V2 localization scores of 70.47% and 73.17% on the two datasets respectively. Code is available on <https://github.com/Saurav-31/ViTOL>.

1. Introduction

Weakly supervised object localization (WSOL) refers to a class of problems that aim to localize objects in an image using image-level annotations as weak supervision. Typically used image-level annotations are classification labels [35], image captions and action annotations [26], object size estimates [27] and context-aware guidance [16]. In this paper, we use image classification labels as weak supervision to localize objects in images. Object location annotations are never explicitly used during training.

We detail below common challenges observed in WSOL with classification labels. Figure 1 showcases common WSOL challenges and output of ViTOL for three query images. (a) *Highlighting Only Discriminative Regions*: Fully supervised image classification does not accurately localize objects in an image as observed in the activation maps

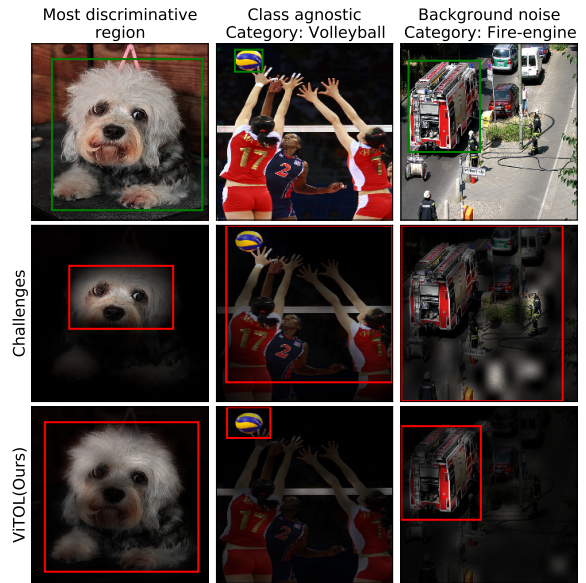


Figure 1. **Common Challenges in WSOL (row2) and ViTOL output (row3)** are visualized for three **query images (row1)**. The GT box is shown in green and the predicted box is shown in red.

(AM) overlaid on Figure 1 (column1, row2). We observe that the entire dog is not highlighted in the AM. In contrast, only the regions which the model considers as most-discriminative for classification are highlighted. This negatively affects the model’s ability to localize the complete object of interest. (b) *Class Agnostic Activation Map*: Figure 1 (column2, row2) demonstrates that objects belonging to multiple classes are highlighted in the same AM. This limits the model’s flexibility in localizing objects belonging to a specific class. (c) *Background Noise*: Accurately localizing an object of interest in a cluttered environment with many objects in the scene is challenging and results in a noisy AM as observed in Figure 1 (column3, row2).

To alleviate the above challenges, we leverage vision transformer architectures [10] to localize objects in an image. Post-hoc work [1, 7] has interpreted and visualized

the capabilities of self-attention layers in transformers to understand objects of interest in a given image. In our work, we introduce below *simple* modifications to the vision transformer. These modifications together with the vision transformer backbone enable a performance improvement of 6.57% (MaxBox Accuracy-V2) and 5.24% (Top-1 Localization) on ImageNet for WSOL over the current state-of-the-art.

patch-based Attention Dropout Layer (p-ADL): Attention dropout mechanism [9] is applied to a sequence of patch embeddings in ViTOL using our proposed p-ADL layer. Two components of the p-ADL layer are patch importance map and patch drop mask. During the course of training, these components are utilized to highlight informative and drop discriminative patches respectively to balance the classification v/s localization performance of the model. For a more detailed explanation, refer to Section 3.1.

Grad Attention Rollout (GAR): Attention rollout [1] generates an attention map by recursively multiplying the attention weight matrices across layers. We introduce a weighted attention rollout mechanism using the gradients computed at each attention map. We refer to this mechanism as Grad Attention Rollout. This post-hoc method in combination with the p-ADL enables the model to quantify the positive and negative contributions in the attention map for each class. This guides the model to generate class-dependent attention maps and a negative clamping operation in GAR further suppresses the effect of background noise in the attention map. We observe that this method is effective in improving upon existing state-of-the-art for WSOL (refer Section 4). In Section 4, we also showcase results for an alternative post-hoc approach called Layer Relevance Propagation [7].

Contributions: In summary, we introduce three simple yet effective modifications to current WSOL methods that enable us to achieve state-of-the-art performance on object localization benchmarks. 1) we adopt a self-attention mechanism using vision transformer in the WSOL task, 2) introduce the p-ADL layer (Figure 2 (C)) to improve the localization ability of the vision transformer model significantly, and 3) leverage grad attention rollout to generate class-dependent activation maps.

We conduct extensive ablation experiments, qualitative and quantitative comparisons against state-of-the-art approaches on the object localization and classification performance of our model. Our best model has a MaxBoxAccV2 accuracy of 70.47% on *ImageNet-1K* [25] and 73.17% on *Caltech-UCSD Birds-200-2011 (CUB)* [30] which are 6.57% and 6.67% better than the current state-of-the-art. Qualitatively, we observe that our approach generates high-quality localization maps that cover the entire object and are fairly robust to clutter in the scene (Figure 4, 5, 6).

2. Related Work

In a weakly supervised learning setup, a major challenge while training classification models is that the classification loss pushes the network to look at the most discriminative features in an image and fails to localize the entire object. To address this issue, various methods [4, 32, 35] have introduced architectural changes in CNNs. More recently, the transformer architecture was introduced by Vaswani *et al.* [29] for machine translation and is widely used since then in many computer vision tasks. Vision transformers (ViT) [10] divide the image into a sequence of image patches and uses an extra class token to learn image representation and perform classification. A more powerful and data efficient extension of transformers was introduced as DeiT [17, 28]. Naseer *et al.* [23] shows that the transformers showcase properties of handling long-range dependencies using their inherent self-attention mechanism. They preserve global information consistently across all layers (Fig.7 in [10]). Moreover, they exhibit properties of robustness to occlusions and preserve object structure which makes them suitable for the WSOL task. Therefore, we utilize the DeiT-Small-Variant (DeiT-S) and DeiT-Base-Variant (DeiT-B) transformer architecture in our work and introduce simple architectural changes to improve localization capabilities.

In WSOL, some pioneering work has been proposed to generate localization maps. CAM [35] uses a global average pooling layer [19] to fuse feature maps and generate a localization map. This has a limitation of focusing only on the small discriminative part of the object. Wonho Bae *et al.* [4] improved CAM by using thresholded average pooling, negative weight clamping and percentile based thresholding techniques. TS-CAM [11] introduces a visual transformer based variant of CAM. They explicitly couple semantic agnostic maps from the vision transformer with semantic-aware maps from CAM to effectively localize objects in the image. In contrast, we introduce a grad attention rollout mechanism (see Sec. 3.2) in addition to architectural changes to the vision transformer which implicitly enable semantic aware object localization. HaS [18], Cut-Mix [31] drop some random parts in the input image to enforce the network to learn the importance of object sub-parts. Despite hiding the image region randomly, it does not necessarily erase the most discriminative parts. In contrast, our approach drops the most-discriminative regions in the attention maps to enforce localization of the entire object (see Sec. 3.1). Some CNN based recent work include ACoL [32], MEIL [21], SPG [33], GC-Net [20] and I2C [34]. We empirically compare the localization performance against each of the above approaches in Sec. 4.

Attention Dropout Mechanism: ADL [9] applies attention based dropout mechanism to feature maps in deeper layers in CNNs. In our approach, we introduce a patch-

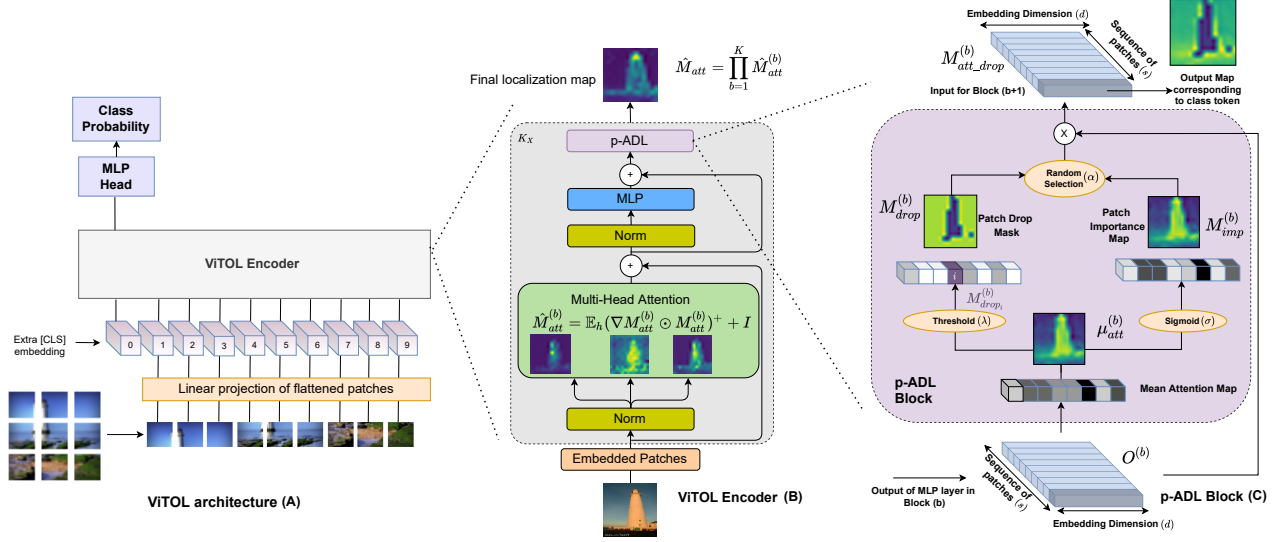


Figure 2. Model overview: (A) ViTOL architecture (B) ViTOL Encoder which consists of an additional p-ADL block. (C) p-ADL Block and its components. We visualize attention maps in B, p-ADL components in C and also highlight critical symbols and GAR equations.

based counterpart to ADL called p-ADL (see Section 3.1) which acts as a regularizer and enhances the localization capability of the model. Unlike ADL which operates on feature maps, p-ADL operates on the patch embeddings including the class token embedding. Owing to the global attention property across all the layers in a transformer, we apply p-ADL at every transformer encoder block.

Class Dependent Attention: Attention maps also need to be class-dependent to correctly identify the object of interest in an image. For CNNs, class dependent behavior was introduced using gradient-based methods such as Grad-CAM [26], Grad-CAM++ [6]. Furthermore, attribution propagation methods [3, 5] propagate relevance for the predicted class based on Deep Taylor Decomposition (DTD) framework [22]. Class-specific behavior is also introduced using Contrastive-LRP [12] and Softmax-Gradient LRP [15].

For transformers, Abnar et al. [1] proposed attention rollout method, which recursively multiply softmax activations corresponding to the class token in each attention layer. Our proposed method uses the gradients with respect to target class to generate final localization maps. Chefer et al. [7] proposed an alternate method of assigning local relevance based on the DTD principle to generate class dependent attention maps. We evaluate attention rollout, gradient-based rollout and layer relevance propagation methods empirically and qualitatively in our proposed approach.

3. Method

In this section, we discuss the details of the proposed method, ViTOL, for WSOL. ViTOL mainly introduces (a) architectural changes and (b) localization map generation

methods. In *architectural changes*, we discuss changes to vision transformer architecture by introducing a patch-based Attention Dropout Layer (p-ADL) as shown in Figure 2. The p-ADL layer acts as a regularizer and assists the network to focus on both discriminative as well as non-discriminative regions of an object. This results in coverage of the entire object region, thereby greatly enhancing the localization ability. In *map generation methods*, we discuss grad attention rollout mechanism (GAR). We compare this against attention rollout (AR) mechanism [1] and layer relevance propagation (LRP) [7] for generating class dependent attention maps. We discuss below both these changes in detail and also explain the procedure for generating high-quality localization maps.

3.1. Architectural Changes

Vision Transformer: We leverage self-attention mechanism of vision transformers and employ DeiT-B and DeiT-S variants to our task. Self-attention module [29] in transformers is defined below.

$$M_{att}^{(b)} = \text{softmax}(Q^{(b)} \cdot K^{(b)T}) \quad (1)$$

$$Z^{(b)} = M_{att}^{(b)} \cdot V^{(b)} \quad (2)$$

$M_{att}^{(b)} \in \mathbb{R}^{(h \times s \times s)}$ is the attention matrix for encoder block b . $Z^{(b)} \in \mathbb{R}^{(h \times s \times d_h)}$ is the output of the self-attention layer in block b . Here, h, s, d_h, d denotes the number of heads in each attention block, number of input sequence of patches, dimension of each head and embedding dimension of the input respectively. Generally, $d = h \cdot d_h$ is used. Here, notations follow [7].

patch-based Attention Dropout Layer (p-ADL): The p-ADL block is illustrated in Figure 2 (C). Largely inspired

by the success of the attention-based dropout layer [9] on CNN-based architectures, we propose the *p-ADL* layer for transformers. This layer has two main components, patch importance map and the patch drop mask. Both these components operate on the mean attention map, which is computed by taking the mean over the embedding dimension.¹ The *patch drop mask* is created by dropping the most activated mean patch embedding based on a *drop threshold* (λ) parameter. The *patch importance map* is calculated by normalizing the mean attention map using a sigmoid activation, which denotes the importance of patches in an image.

Either the patch importance map or the patch drop mask is randomly chosen during training based on an *embedding drop rate* (α) parameter. Patch drop mask facilitates the network to improve its localization capability. This is because the most discriminative region of the attention map is dropped, forcing the model to focus on the remaining region of the desired object (as shown in Fig. 3). Patch importance map aids the classification accuracy. This is because, over the course of training, the patch drop mask might drop informative regions of the object of interest. Patch importance map counters this behavior by ensuring necessary information about the object of interest is retained (see Figure 2 (C)). Some ablations regarding p-ADL layer are studied in Section 4.4.2. Note that this layer does not introduce any new trainable parameters. We formulate the following equations for the p-ADL layer.

$$\mu_{att}^{(b)} = \mathbb{E}_d(O^{(b)}) \quad (3)$$

$$M_{imp}^{(b)} = \sigma(\mu_{att}^{(b)}) \quad (4)$$

$$M_{drop_i}^{(b)} = \begin{cases} 0 & \mu_{att_i}^{(b)} \geq (\lambda \cdot \max(\mu_{att}^{(b)})) \\ 1 & \mu_{att_i}^{(b)} < (\lambda \cdot \max(\mu_{att}^{(b)})) \end{cases} \quad (5)$$

$$M_{att.drop}^{(b)} = \begin{cases} O^{(b)} \cdot M_{drop}^{(b)} & p_{random} < \alpha \\ O^{(b)} \cdot M_{imp}^{(b)} & p_{random} \geq \alpha \end{cases} \quad (6)$$

where $O^{(b)} \in \mathbb{R}^{(s \times d)}$ denotes the output patch embedding matrix at encoder block b , \mathbb{E}_d denotes the mean along embedding dimension, $\mu_{att}^{(b)}, M_{imp}^{(b)}, M_{drop}^{(b)} \in \mathbb{R}^{(s)}$ denotes the mean patch embedding, patch importance map, patch drop mask for encoder block b , $M_{drop_i}^{(b)}$ denotes the i^{th} index of $M_{drop}^{(b)}$, s is the number of patches in input sequence, $M_{att.drop}^{(b)} \in \mathbb{R}^{(s \times d)}$ is the output of the p-ADL layer for encoder block b . σ denotes the sigmoid function. $p_{random} \in [0, 1]$ can be generated using any random number generator function.

Patch drop mask intuition: In Fig. 3, we visualize the output of patch drop mask corresponding to the class token, after every encoder block, for a query image. Initially, a

very small patch of crane was activated which is dropped in Block 1 by the patch drop mask. Dropping that information forces the model to explore other parts of crane, to help classify it correctly. This happens after each encoder block (Block 2-5), thus driving the model to explore more object regions. However, dropping every informative patch progressively after each block might lead to embeddings which no longer have information about the object (Block 5). This loss of object information may lead to object mis-classification. To avoid such circumstances, patch importance map is also chosen randomly which highlights the activated patches and retains object specific information. The visualization for patch importance and patch drop mask for an example image is observed in Figure 2 (C).

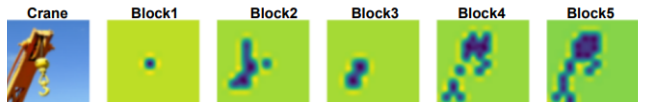


Figure 3. **Patch Drop Mask progression.** Patch drop mask for set of encoder blocks are shown here for a sample crane image using our model with p-ADL layer. Dark color denotes dropped regions.

3.2. Generation of localization maps

Grad Attention Rollout (GAR): Attention rollout [1] is an intuitive mechanism in which self-attention maps are averaged over each attention head in the encoder block and recursively multiplied over all transformer layers to produce a localization map. Rollout adds an Identity matrix I to the attention matrix at each layer to account for residual connections [13] in the network. However, this method assumes that attentions are linearly combined. As stated by Chefer *et al.* [7], this overlooks the fact that GELU [14] activation is used in all intermediate layers. An ill-effect of this ignorance is that it fails to distinguish between positive and negative contributions to the final attention map. In our method, we compute a weighted attention rollout using gradient maps for each respective attention map and then use negative attention clamping to consider only positive values for each attention head. The GAR method is as given below.

$$\hat{M}_{att}^{(b)} = I + \mathbb{E}_h(\nabla M_{att}^{(b)} \odot M_{att}^{(b)})^+ \quad (7)$$

$$\hat{M}_{att} = \prod_{b=1}^K \hat{M}_{att}^{(b)} \quad (8)$$

Here \odot is the Hadamard product, and \mathbb{E}_h is the mean over head dimension, $\nabla M_{att}^{(b)} \in \mathbb{R}^{(h \times s \times s)}$ is the gradient matrix corresponding to each attention matrix $M_{att}^{(b)}$, K is the total number of encoder blocks. $\hat{M}_{att} \in \mathbb{R}^{(s \times s)}$ is the final rollout attention map. Superscript $+$ in Eq. 7 refers to negative weight clamping operation.

Layer Relevance Maps: Chefer *et al.* [7] examined the problems in attention rollout and proposed a method that

¹We use the class token embedding while calculating the mean. This empirically results in a better localization performance.

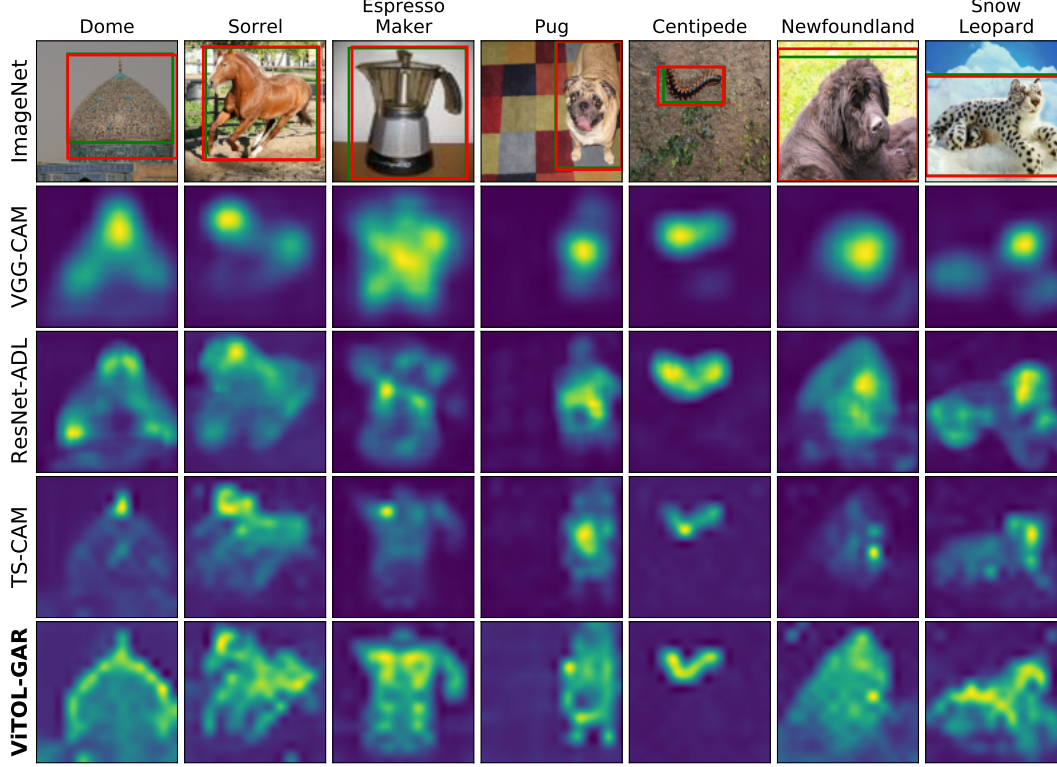


Figure 4. **Qualitative comparison of activation maps with the state-of-the-art weakly supervised object localization models.**

Row-1: Query Image from ImageNet, Row-2: VGG16-CAM, Row-3: ResNet50-ADL, Row-4: TS-CAM based on DeiT-S backbone, Row-5: ViTOL based on Gradient Attention Rollout (GAR). Class label is displayed as a column label.

computes local relevance based on the deep Taylor decomposition principle. This provides class-specific visualizations. For more details we refer the reader to [7]. The final attention map is computed as follows.

$$\tilde{M}_{att}^{(b)} = I + \mathbb{E}_h(\nabla M_{att}^{(b)} \odot R^{(n_b)})^+ \quad (9)$$

$$\tilde{M}_{att} = \prod_{b=1}^K \tilde{M}_{att}^{(b)} \quad (10)$$

where \odot is the Hadamard product, n_b denotes the softmax layer in encoder block b . $R^{(n_b)}$ denotes the relevance for n_b layer, $\tilde{M}_{att} \in \mathbb{R}^{(s \times s)}$ is the final attention map.

In our approach, we evaluate both the methods defined above to get a self-attention map using Eq. (8), (10). This final matrix \tilde{M}_{att} of size $(s \times s)$ represents how each patch in an image interacts with other patches. We only consider the first row of size $(1 \times s)$ corresponding to the $[CLS]$ token to compute the final attention map. We skip the first value in this row to ignore the $[CLS]$ token’s attention with itself and reshape the extracted attention embedding of size $(s - 1)$ to size $(\sqrt{s - 1} \times \sqrt{s - 1})$ which denotes the final localization map for our method.

4. Experiments

4.1. Dataset

We evaluate ViTOL on two benchmark datasets in WSOL namely ImageNet-1K [25] and Caltech-UCSD Birds-200-2011 (CUB) [30]. We use the same data splits as defined in [8]. Imagenet-1K dataset has 1000 classes with approximately 1.3M images for training and 50k images for testing. It is one of the most challenging datasets because of its large inter-class variation. Previous WSOL methods [9, 18, 31–33, 35] have seen relatively small improvements (MaxBoxAccV2) for WSOL as observed in Table 1. CUB dataset is comparatively small with 200 classes consisting of 5994 train and 5794 test images.

4.2. Evaluation Metric

We generate a binary mask from localization map and use it to generate a prediction bounding box around the foreground. If the Intersection over Union (IOU) of the predicted box and ground truth box is greater than a threshold δ then it is classified as a correct prediction. We use 4 metrics for evaluating all our methods: 1) *MaxBoxAccV2* [8]: average of the localization accuracy across IOU thresholds

	ImageNet		CUB	
Method	Architecture	MaxBoxAccV2	Architecture	MaxBoxAccV2
CAM	ResNet	63.7	VGG	63.7
HaS	Inception	63.7	ResNet	64.7
ACoL	Inception	63.7	ResNet	66.5
SPG	Inception	63.3	ResNet	60.4
ADL	ResNet	63.7	VGG	66.3
CutMix	Inception	63.9	ResNet	62.8
Ours-GAR	DeiT-B+p-ADL	69.17	DeiT-B+p-ADL	72.42
Ours-LRP	DeiT-B+p-ADL	70.47	DeiT-B+p-ADL	73.17

Table 1. **Comparison with the state-of-the-art methods (MaxBoxAccV2).** Scores for MaxBoxAccV2 for the baseline methods are borrowed from [8] on test split for both *ImageNet* and *CUB* datasets.

Method	Architecture	Top1-Loc	GT-Known	Top1-CLS
CAM	InceptionV3	46.3	62.7	70.6
ACoL	VGG16	45.8	63.0	64.5
SPG	InceptionV3	48.6	64.7	71.1
ADL	InceptionV3	48.7	-	61.2
CutMix	VGG16	43.5	-	66.4
MEIL [21]	InceptionV3	49.5	-	-
GC-Net [20]	InceptionV3	49.1	-	-
CAAM+SSAM [2]	ResNet50	52.36	67.89	-
SPA [24]	InceptionV3	52.7	68.33	-
I ² C [34]	InceptionV3	53.1	68.5	-
TS-CAM [11]	DeiT-S	53.4	67.6	74.3
Ours-GAR	DeiT-S+p-ADL	54.74	71.86	71.84
Ours-LRP	DeiT-S+p-ADL	53.62	70.48	71.84
Ours-GAR	DeiT-B+p-ADL	57.62	71.32	77.08
Ours-LRP	DeiT-B+p-ADL	58.64	72.51	77.08

Table 2. **Comparison with the state-of-the-art methods (Top1-Loc and GT-Known):** Top1-Loc and GT-Known metrics are compared with other WSOL methods on ImageNet validation data.

$\delta \in \{0.3, 0.5, 0.7\}$ to show the fineness of localization. 2) *GT-known Loc (IOU 50)*: use the GT label during inference to measure performance at fixed IOU threshold $\delta = 0.5$. 3) *Top1-Loc*: This metric provides a positive value when classification is correct and IOU is greater than 50%. 4) *Top1-CLS*: Top1 classification performance.

Key Metric: *MaxBoxAccV2* metric measures the ability of the model to localize the objects of interest, *which we consider as the key criterion to compare different WSOL approaches*. Choe *et al.* [8] argue that the *goal of WSOL is to localize objects and not classify images and therefore they advocate the use of MaxBoxAccV2* as the desired performance metric. We adopt this recently fixed version of the localization metric. Moreover, only an increase in the classification accuracy without any observed improvement in the localization performance will increase Top1-localization accuracy value, however, *this does not translate to the performance improvements on the localization task*.

4.3. Implementation Details

We use DeiT-B and DeiT-S architectures for our task. We use $s = 197$, $K = 12$, patch size of 16x16 and image resolution of 224x224 in all our experiments. DeiT-S and DeiT-B uses embedding dimension of 384 and 768, and attention heads 6 and 12 respectively in each ViTOL encoder block. In p-ADL, we use a drop threshold (λ) of 0.9 and embedding drop-rate (α) as 0.75 across all experiments.

Training and Testing details: On ImageNet-1K, for baseline models, we use the DeiT-B and DeiT-S ImageNet pre-trained weights and evaluate on all the methods, namely, AR, GAR and LRP as stated in Table 1, 2 and 3. In experiments with ViTOL, on ImageNet-1K, we train the model with the p-ADL layer for 10 epochs using a learning rate of 0.0001, a weight decay (wd) of 0, and a learning rate (lr) decay of 0.1 after every 3 epochs. For experiments on CUB dataset, to create a new baseline model, we take the ImageNet pre-trained DeiT-B model but initialize a new

Architecture	Method	ImageNet				CUB			
		MaxboxAccV2	IOU30	IOU50	IOU70	MaxboxAccV2	IOU30	IOU50	IOU70
DeiT-B	AR	45.94	70.41	44.39	23.03	49.75	83.27	49.37	16.6
DeiT-B+p-ADL	AR	46.51	70.71	45.33	23.47	57.42	88.45	60.87	22.93
DeiT-B	GAR	62.82	83.58	64.89	40.00	70.79	96.15	78.2	38.03
DeiT-B+p-ADL	GAR	69.17	86.95	71.32	49.25	72.42	96.68	80.89	39.69

Table 3. **Effect of p-ADL and Comparison of Attention Rollout v/s Grad Attention Rollout**

classifier layer with 200 class output. We trained the model for 50 epochs with a lr of 0.0001, wd of 0, and lr decay of 0.1 every 10 epochs. In experiments with ViTOL, we add the p-ADL layer and continue training this model. Here, we use lr of 0.00001, wd of 0, and lr decay of 0.1 after every 5 epochs. Following [8], we use last saved checkpoint model for all the evaluations in Table 1 and 2. During inference, p-ADL layer is not used.

4.4. Quantitative Results

4.4.1 Comparison to the State-of-the-Art

Table 1 and Table 2 compare our proposed approach ViTOL against the state-of-the-art (SOTA) approaches for WSOL. In Table 1, we use the MaxBoxAccV2 metric [8] to quantify the localization performance on ImageNet and CUB datasets. Here, we compare against all SOTA approaches which report the aforementioned metrics. The *Architecture* column mentions the network architecture across all approaches based on the best performing backbones on each dataset. We observe that our approaches with a DeiT-B backbone with p-ADL + (a) GAR and (b) LRP significantly outperform the other WSOL approaches. Using the MaxBoxAccV2 metric we observe gains by a margin of 6.57% and 6.67%, on ImageNet and CUB datasets respectively.

In Table 2, we use the Top-1 localization accuracy metric to compare ViTOL against the SOTA approaches for WSOL on the ImageNet dataset. We outperform the SOTA approaches by achieving a Top-1 localization accuracy of 58.64%, a GT-Known value of 72.51% and Top-1 classification (CLS) accuracy of 77.08%.

Architecture	MMp	MpM	ImageNet	CUB
DeiT-B+ p-ADL		✓	70.41	71.99
DeiT-B+ p-ADL	✓		70.47	73.17

Table 4. **Changing p-ADL position inside encoder blocks:** MMp [MSA \ll MLP \ll p-ADL] denotes that p-ADL is placed after MLP inside each encoder block and MpM [MSA \ll p-ADL \ll MLP] denotes that p-ADL is inserted in between MLP and MSA layers. MaxBoxAccv2 is compared on both ImageNet and CUB.

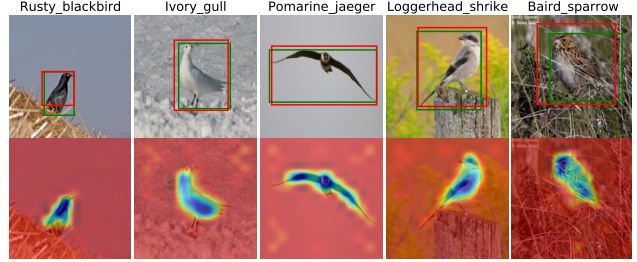


Figure 5. **Qualitative results on CUB.** GT bounding box is shown in green and predicted bounding box is shown in red.

4.4.2 Ablation Study

Impact of p-ADL: In Table 3, we compare the performance of DeiT-B with and without p-ADL using AR and GAR map generation mechanisms. We observe that, for the AR map-generation method, the DeiT-B+p-ADL achieves a performance gain of 0.57% and 7.67% on ImageNet and CUB datasets respectively against DeiT-B. Similarly, for the GAR method, DeiT-B+p-ADL achieves a gain of 6.35% and 1.63% respectively against DeiT-B. These performance gains can be attributed to the ability of the p-ADL layer to highlight discriminative and non-discriminative regions of the object of interest.

Position of p-ADL: We experiment with the position of the p-ADL layer inside each encoder block to understand its impact on the localization performance. We observe from Table 4 that we get minor performance improvements if we place the p-ADL after the MLP layer as shown in Fig 2. This minor difference can be intuitively attributed to the MLP layer prioritizing classification performance over the localization in the case of p-ADL placed before MLP.

Improvements with GAR: The usage of gradients (equation 7) provides additional information about the semantics of the class of the object. This change introduces class dependent behavior for the attention maps. We observe a huge performance gain attributed to this behavior in Table 3. A gain of 22.66% and 15% is observed for ImageNet and CUB datasets respectively for the DeiT+p-ADL model.

Other ablations: Some other ablation studies include i) comparison of GAR and LRP across different transformer backbones, ii) detailed study of patch drop masks for the p-ADL layer and iii) effect of changing embedding drop rate

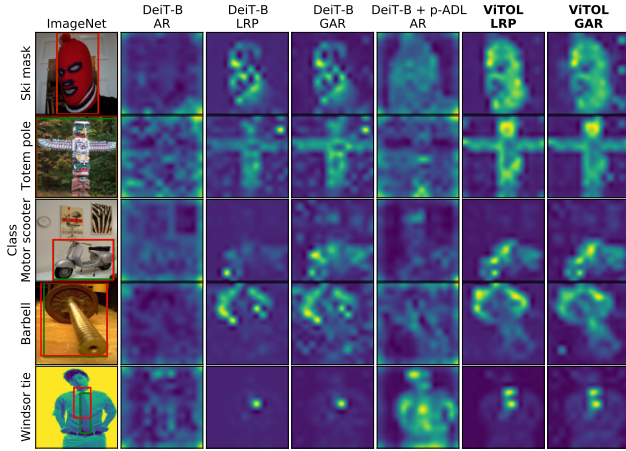


Figure 6. **Qualitative ablation for map generation methods.** Comparison of final localization map for baseline DeiT-B (Column 2,3,4) and DeiT-B + p-ADL layer (Column 5, 6, 7) is shown here for map generation methods AR, GAR and LRP. GT bounding box is shown in green. Predicted bounding box (red) is calculated using the ViTOL-GAR.

and drop threshold hyper-parameters for p-ADL layer. We refer the reader to supplementary for these ablations.

4.5. Qualitative Results

In this section, we qualitatively compare the performance of ViTOL in Fig. 4. The attention maps of our model are visualized for CUB in Fig. 5 and ImageNet in Fig. 6.

Comparison with SOTA: In Figure 4, we highlight the localization maps for ViTOL and compare against state-of-the-art WSOL methods, VGG-CAM [35], ResNet50-ADL [9], TS-CAM [11], for a few randomly sampled images from ImageNet. Our method, ViTOL, shows superior localization performance in highlighting the entire object region across all the sample images. Moreover, our localization maps reconstruct the structure of the sorrel and centipede classes more accurately as compared with other methods.

Qualitative Ablation on ImageNet: In Figure 6, we generate attention maps for five random example images from the ImageNet validation split. For each of the example images, we visualize the baseline DeiT-B with AR, GAR and LRP in the second, third and fourth columns, and, DeiT-B + p-ADL with AR, ViTOL with LRP and ViTOL with GAR attention maps in the final three columns.

We observe that attention maps generated for ViTOL with GAR/LRP show dependency with the class, are noise-free and cover the complete object of interest. Our method is superior in localizing the objects in the image as compared to the DeiT-B baseline models. In Figure 6, row five, we observe that the attention maps using DeiT-B highlight the whole person. However, the object of interest is the tie worn by the person. The class conditional behavior of ViTOL enables the models to approximately localize the *tie* in

the image. Similarly in row three, our method highlights the desired class in an image containing multiple objects. The first four rows showcase that the entire object of interest is highlighted in the ViTOL attention maps. In the second row we observe that the negative contributions from the noisy background observed in the attention maps of DeiT-B + p-ADL + AR are alleviated in ViTOL-GAR/LRP.

Visualization on CUB: In Figure 5, first row, we showcase five random example images from the CUB dataset. In the second row, we overlay the attention maps obtained from ViTOL-LRP for these images. In the first column, half the bird is occluded and our model only captures the region of bird that is visible. In the second column, our method successfully localizes a bird camouflaged into its background. In the third column our method captures the bird’s narrow wings in its attention map. In the fourth column, the tail of the bird looks similar to the stem of the plants in the background. Our method succeeds in highlighting the entire bird along with its thin tail. In last column, our method localizes the visible portion of the bird which is resting in a noisy background.

5. Conclusion

In this work, we explore the problem of WSOL. We discuss in detail the common challenges that image classification models encounter when localizing objects of interest. To alleviate these challenges, we introduced our proposed method ViTOL. We discuss simple yet effective changes to a baseline transformer architecture and introduced a new map generation method for WSOL. We employed patch-based attention dropout layer (p-ADL) which aided in increasing the coverage of the localization map. We also proposed a post-hoc map generation method which enforced class-dependent attention map generation for WSOL. Finally, we showcased state-of-the-art performance across different metrics, datasets and methods for object localization. Our experimental results establish the significance of using ViTOL in object localization. The clear resemblance of the actual shape of the object in the attention map visualization shows great potential to leverage it for other tasks such as weakly-supervised semantic segmentation, temporal action localization, etc. In order to apply the proposed p-ADL and GAR map generation methods for CNN based architectures, the p-ADL layer and rollout mechanism can potentially be applied to those layers that have dense features and high receptive fields. In future work, we plan to explore variants of global attention in transformers to capture fine-grained details as opposed to the patch-based attention which we currently employ.

Acknowledgements: This work was funded by Mercedes-Benz Research and Development, India. Special thanks to Brijesh Pillai for the research support and Sujay Babruwad for the assistance in our work.

References

- [1] Samira Abnar and Willem H Zuidema. Quantifying attention flow in transformers. In *ACL*, 2020. 1, 2, 3, 4
- [2] Sadbhavana Babar and Sukhendu Das. Where to look?: Mining complementary image regions for weakly supervised object localization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1010–1019, 2021. 6
- [3] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one*, 10(7):e0130140, 2015. 3
- [4] 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. 2
- [5] Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, Klaus-Robert Müller, and Wojciech Samek. Layer-wise relevance propagation for neural networks with local renormalization layers. In *International Conference on Artificial Neural Networks*, pages 63–71. Springer, 2016. 3
- [6] Aditya Chattopadhyay, 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. 3
- [7] Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 782–791, 2021. 1, 2, 3, 4, 5
- [8] Junsuk Choe, Seong Joon Oh, Sanghyuk Chun, Zeynep Akata, and Hyunjung Shim. Evaluation for weakly supervised object localization: Protocol, metrics, and datasets. *arXiv preprint arXiv:2007.04178*, 2020. 5, 6, 7
- [9] 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. 2, 4, 5, 8
- [10] 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. In *International Conference on Learning Representations*, 2020. 1, 2
- [11] 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. 2, 6, 8
- [12] Jindong Gu, Yinchong Yang, and Volker Tresp. Understanding individual decisions of cnns via contrastive backpropagation. In *Asian Conference on Computer Vision*, pages 119–134. Springer, 2018. 3
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [14] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 4
- [15] Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. A benchmark for interpretability methods in deep neural networks. *Advances in neural information processing systems*, 32, 2019. 3
- [16] Vadim Kantorov, Maxime Oquab, Minsu Cho, and Ivan Laptev. Contextlocnet: Context-aware deep network models for weakly supervised localization. In *European Conference on Computer Vision*, pages 350–365. Springer, 2016. 1
- [17] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *ACM Computing Surveys (CSUR)*, 2021. 2
- [18] Krishna Kumar Singh and Yong Jae Lee. Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3524–3533, 2017. 2, 5
- [19] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013. 2
- [20] Weizeng Lu, Xi Jia, Weicheng Xie, Linlin Shen, Yicong Zhou, and Jinming Duan. Geometry constrained weakly supervised object localization. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16*, pages 481–496. Springer, 2020. 2, 6
- [21] 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. 2, 6
- [22] Grégoire Montavon, Sebastian Lapuschkin, Alexander Binder, Wojciech Samek, and Klaus-Robert Müller. Explaining nonlinear classification decisions with deep taylor decomposition. *Pattern Recognition*, 65:211–222, 2017. 3
- [23] Muhammad Muzammal Naseer, Kanchana Ranasinghe, Salman H Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *Advances in Neural Information Processing Systems*, 34, 2021. 2
- [24] Xingjia Pan, Yingguo Gao, Zhiwen Lin, Fan Tang, Weiming Dong, Haolei Yuan, Feiyue Huang, and Changsheng Xu. Unveiling the potential of structure preserving for weakly supervised object localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11642–11651, 2021. 6
- [25] 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. 2, 5
- [26] 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. 1, 3
- [27] Miaojing Shi and Vittorio Ferrari. Weakly supervised object localization using size estimates. In *European Conference on Computer Vision*, pages 105–121. Springer, 2016. 1
- [28] 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. 2
- [29] 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. 2, 3
- [30] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-ucsd birds 200. 2010. 2, 5
- [31] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6023–6032, 2019. 2, 5
- [32] 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. 2, 5
- [33] Xiaolin Zhang, Yunchao Wei, Guoliang Kang, Yi Yang, and Thomas Huang. Self-produced guidance for weakly-supervised object localization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 597–613, 2018. 2, 5
- [34] Xiaolin Zhang, Yunchao Wei, and Yi Yang. Inter-image communication for weakly supervised localization. 2020. 2, 6
- [35] 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. 1, 2, 5, 8