

Rotate to Attend: Convolutional Triplet Attention Module

Diganta Misra *
Landskape

mishradiganta91@gmail.com

Ajay Uppili Arasanipalai *
University of Illinois, Urbana Champaign
aua2@illinois.edu

Trikay Nalamada *
Indian Institute of Technology, Guwahati

nalamada.trikay@gmail.com

Qibin Hou
National University of Singapore
andrewhou@gmail.com

Abstract

Benefiting from the capability of building inter-dependencies among channels or spatial locations, attention mechanisms have been extensively studied and broadly used in a variety of computer vision tasks recently. In this paper, we investigate light-weight but effective attention mechanisms and present triplet attention, a novel method for computing attention weights by capturing cross-dimension interaction using a three-branch structure. For an input tensor, triplet attention builds inter-dimensional dependencies by the rotation operation followed by residual transformations and encodes inter-channel and spatial information with negligible computational overhead. Our method is simple as well as efficient and can be easily plugged into classic backbone networks as an add-on module. We demonstrate the effectiveness of our method on various challenging tasks including image classification on ImageNet-1k and object detection on MSCOCO and PASCAL VOC datasets. Furthermore, we provide extensive insight into the performance of triplet attention by visually inspecting the GradCAM and GradCAM++ results. The empirical evaluation of our method supports our intuition on the importance of capturing dependencies across dimensions when computing attention weights. Code for this paper can be publicly accessed at <https://github.com/LandskapeAI/triplet-attention>.

1. Introduction

Over the years of computer vision research, convolutional neural network architectures of increasing depth have demonstrated major success in many computer vision tasks [12, 17, 28, 29, 36]. Numerous recent work [5, 14, 24, 33, 23] have proposed using either channel at-




Figure 1. Abstract representation of triplet attention with three branches capturing cross-dimension interaction. Given the input tensor, triplet attention captures inter-dimensional dependencies by rotating the input tensor followed by residual transformation.

tention, or spatial attention, or both to improve the performance of these neural networks. These attention mechanisms have the capabilities of improving the feature representations generated by standard convolutional layers by explicitly building dependencies among channels or weighted spatial mask for spatial attention. The intuition behind learning attention weights is to allow the network to have the ability to learn where to attend and further focus on the target objects.

One of the most prominent methods is the squeeze-and-excitation networks (SENet) [14]. Squeeze and Excite (SE) module computes channel attentions and provides incremental performance gains at a considerably low cost. SENet was succeeded by Convolutional Block Attention Module (CBAM) [33] and Bottleneck Attention Module (BAM) [24], both of which stressed on providing robust representative attentions by incorporating spatial attention along with channel attention. They provided substantial performance gains over their squeeze-and-excite counter-

*Equal Contribution

part with a small computational overhead.

Different from the aforementioned attention approaches that require a number of extra learnable parameters, the foundation backbone of this paper is to investigate the way of building cheap but effective attentions while maintaining similar or providing better performance. In particular, we aim to stress on the importance of capturing cross-dimension interaction while computing attention weights to provide rich feature representations. We take inspiration from the method of computing attention in CBAM [33] which successfully demonstrated the importance of capturing spatial attention along with channel attention. In CBAM, the channel attention is computed in a similar way as that of SENet [14] except for the usage of global average pooling (GAP) and global max pooling (GMP) while the spatial attention is generated by simply reducing the input to a single channel output to obtain the attention weights. We observe that the channel attention method within CBAM [33] although providing significant performance improvements does not account for cross-dimension interaction which we showcase to have a favorable impact on the performance when captured. Additionally, CBAM incorporates dimensionality reduction while computing channel attention which is redundant to capture non-linear local dependencies between channels.

Based on the above observation, in this paper, we propose *triplet attention* which accounts for cross-dimension interaction in an efficient way. Triplet attention comprises of three branches each responsible for capturing cross-dimension between the spatial dimensions and channel dimension of the input. Given an input tensor with shape $(C \times H \times W)$, each branch is responsible for aggregating cross-dimensional interactive features between either the spatial dimension H or W and the channel dimension C . We achieve this by simply permuting the input tensors in each branch and then passing the tensor through a *Z-pool*, followed by a convolutional layer with kernel size of $k \times k$. The attention weights are then generated by a sigmoid activation layer and then is applied on the permuted input tensor before permuting it back into the original input shape.

Compared to previous channel attention mechanisms [2, 10, 14, 24, 33], our approach offers two advantages. First, our method helps in capturing rich discriminative feature representations at a negligible computational overhead which we further empirically verify by visualizing the Grad-CAM [27] and Grad-CAM++ [3] results. Second, unlike our predecessors, our method stresses the importance of cross-dimension interaction with no dimensionality reduction, thus eliminating indirect correspondence between channels and weights.

We showcase this way of computing attention in parallel across branches while accounting for cross-dimension dependencies is extremely effective and cheap in com-

putational terms. For instance, for ResNet-50 [12] with 25.557M parameters and 4.122 GFLOPs, our proposed plug-in triplet attention results in an increase of parameters by 4.8K and GFLOPs by $4.7e-2$ respectively while providing a 2.28% improvement in Top-1 accuracy. We evaluate our method on ImageNet-1k [7] classification and object detection on PASCAL VOC [8] and MS COCO [21] while also providing extensive insight into the effectiveness of our method by visualizing the Grad-CAM [27] and Grad-CAM++ [3] outputs respectively.

2. Related Work

Attention in human perception relates to the process of selectively concentrating on parts of the given information while ignoring the rest. This mechanism helps in refining perceived information while retaining the context of it. Over the last few years, several researched methods have proposed to efficiently incorporate this attention mechanism in deep convolution neural network (CNN) architectures to improve performance on large-scale vision tasks. In the following part of this section, we will review some attention mechanisms that are strongly related to this work.

Residual Attention Network [31] proposes a trunk-and-mask encoder-decoder style module to generate robust three-dimensional attention maps. Due to the direct generation of 3D attention maps, the method is quite computationally complex as compared to the recently proposed methods to compute attention. This was followed by the introduction of Squeeze-and-Excitation Networks (SENet) [14] which as debated by many was the first to successfully implement an efficient way of computing channel attention while providing significant performance improvements. The aim of SENet was to model the cross-channel relationships in feature maps by learning per-channel modulation weights. Succeeding SENet, Convolutional Block Attention Module (CBAM) [33] was proposed, in which they enrich the attention maps by adding max pooled features for the channel attention along with an added spatial attention component. This combination of spatial attention and channel attention demonstrated substantial improvement in performance as compared to SENet. More recently, Double Attention Networks (A^2 -Nets) [6] introduced a novel relation function for Non-Local (NL) blocks. NL blocks [32] were introduced to capture long range dependencies via non-local operations and were designed to be lightweight and easy to use in any architecture. Global Second order Pooling Networks (GSoP-Net) [10] uses second-order pooling for richer feature aggregation. The key idea is to gather important features from the entire input space using second order pooling and subsequently distributing them to make it easier for further layers to recognize and propagate. Global-Context Networks (GC-Net) [2] propose a novel NL-block integrated with a SE block in which they aimed to combine contextual




Figure 2. **Comparisons with different attention modules:** (a) Squeeze Excitation (SE) Module; (b) Convolutional Block Attention Module (CBAM); (c) Global Context (GC) Module; (d) triplet attention (ours). The feature maps are denoted as feature dimensions, e.g. $C \times H \times W$ denotes a feature map with channel number C , height H and width W . \otimes represents matrix multiplication, \odot denotes broadcast element wise multiplication and \oplus denotes broadcast element-wise addition.

representations with channel weighting more efficiently. Instead of simple downsampling by GAP as in the case of SENet [14], GC-Net uses a set of complex permutation-based operations to reduce the feature maps before passing it to the SE block.

Attention mechanisms have also been successfully used for image segmentation and fine grained image classification. Criss-Cross Networks (CCNet) [15] and SPNet [13] present novel attention blocks to capture rich contextual information using intersecting strips. Xiao *et al.* [35] propose a pipeline integrated with one bottom-up and two top-down attention for fine grained image classification. Cao *et al.* [1] introduce the 'Look and Think Twice' mechanism which is based on a computational feedback process inspired from the human visual cortex which helps in capturing visual attention on target objects even in distorted background conditions.

Most of the above methods have significant shortcomings which we address in our method. Our triplet attention module aims to capture cross-dimension interaction and thus is able to provide significant performance gains at a justified negligible computational overhead as compared to the above described methods where none of them account for cross-dimension interaction while allowing some form of dimensionality reduction which is unnecessary to capture cross-channel interaction.

3. Proposed Method

In this section, we first revisit CBAM [33] and analytically diagnose the efficiency of the shared MLP structure within the channel attention module of CBAM. Subsequently, we propose our triplet attention module where we demonstrate the importance of cross-dimension dependencies and further compare the complexity of our method with other standard attention mechanisms. Finally, we conclude

by showcasing how to adapt triplet attention into standard deep CNN architectures for different challenging tasks in the domain of computer vision.

3.1. Revisiting Channel Attention in CBAM

We first revisit the channel attention module used in CBAM [33] in this subsection. Let $\chi \in \mathbb{R}^{C \times H \times W}$ be the output of a convolutional layer and the subsequent input to the channel attention module of CBAM where C , H and W represent the channels of the tensor or the number of filters, height, and width of the spatial feature maps, respectively. The channel attention in CBAM can be represented by the following equation:

$$\omega = \sigma(f_{(\mathbf{w}_0, \mathbf{w}_1)}(g(\chi)) + f_{(\mathbf{w}_0, \mathbf{w}_1)}(\delta(\chi))) \quad (1)$$

where $\omega \in \mathbb{R}^{C \times 1 \times 1}$ represent the learnt channel attention weights which are then applied to the input χ , $g(\chi)$ is the global average pooling (GAP) function as formulated as follows:

$$g(\chi) = \frac{1}{W \times H} \sum_{i=1}^H \sum_{j=1}^W \chi_{i,j} \quad (2)$$

and $\delta(\chi)$ represents the global max pooling (GMP) function written as:

$$\delta(\chi) = \max_{H,W}(\chi) \quad (3)$$

The above two pooling functions make up the two methods of spatial feature aggregation in CBAM. Symbol σ represents the sigmoid activation function. Functions $f_{(\mathbf{w}_0, \mathbf{w}_1)}(g(\chi))$ and $f_{(\mathbf{w}_0, \mathbf{w}_1)}(\delta(\chi))$ are two transformations. Thus, after expanding $f_{(\mathbf{w}_0, \mathbf{w}_1)}(g(\chi))$ and $f_{(\mathbf{w}_0, \mathbf{w}_1)}(\delta(\chi))$, we have the following form of ω :

$$\omega = \sigma(\mathbf{W}_1 \text{ReLU}(\mathbf{W}_0 g(\chi)) + \mathbf{W}_1 \text{ReLU}(\mathbf{W}_0 \delta(\chi))) \quad (4)$$

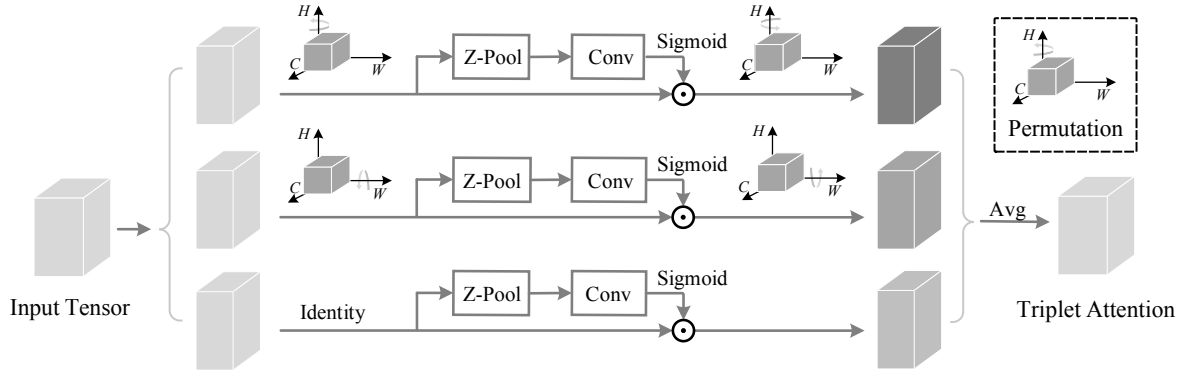


Figure 3. Illustration of the proposed triplet attention which has three branches. The top branch is responsible for computing attention weights across the channel dimension C and the spatial dimension W . Similarly, the middle branch is responsible for channel dimension C and spatial dimension H . The final branch at the bottom is used to capture spatial dependencies (H and W). In the first two branches, we adopt rotation operation to build connections between the channel dimension and either one of the spatial dimension. Finally, the weights are aggregated by simple averaging. More details can be found in Sec. 3.2

where ReLU represents the Rectified Linear Unit and \mathbf{W}_0 and \mathbf{W}_1 are weight matrices, the size of which are defined to be $C \times \frac{C}{r}$ and $\frac{C}{r} \times C$, respectively. Here, r represents the reduction ratio in the bottleneck of the MLP network which is responsible for dimensionality reduction. Larger r results in lower computational complexity and vice versa. To note, the weights of the MLP: \mathbf{W}_0 and \mathbf{W}_1 are shared in CBAM for both the inputs: $g(\chi)$ and $\delta(\chi)$. In Eq. (4), the channel descriptors are projected into a lower dimensional space and then maps them back which causes loss in inter-channel relation because of the indirect weight-channel correspondence.

3.2. Triplet Attention

As demonstrated in Sec. 1, the goal of this paper is to investigate how to model cheap but effective channel attention while not involving any dimensionality reduction. In this subsection, unlike CBAM [33] and SENet [14], which require a certain number of learnable parameters to build inter-dependencies among channels, we present an almost parameter-free attention mechanism to model channel attention and spatial attention, namely triplet attention.

Overview: The diagram of the proposed triplet attention can be found in Fig. 3. As the name implies, triplet attention is made up of three parallel branches, two of which are responsible for capturing cross-dimension interaction between the channel dimension C and either the spatial dimension H or W . The remaining final branch is similar to CBAM [33], used to build spatial attention. The outputs from all three branches are aggregated using simple averaging. In the following, before specifically describing the proposed triplet attention, we first introduce the intuition of building cross-dimension interaction.

Cross-Dimension Interaction: Traditional ways of com-

puting channel attention involve computing a singular weight, often a scalar for each channel in the input tensor and then scaling these feature maps uniformly using the singular weight. Though this process of computing channel attention has been proven to be extremely lightweight and quite successful, there is a significant missing piece in considering this method. Usually, to compute these singular weights for channels, the input tensor is spatially decomposed to one pixel per channel by performing global average pooling. This results in a major loss of spatial information and thus the inter-dependence between the channel dimension and the spatial dimension is absent when computing attention on these single pixel channels. CBAM [33] introduced spatial attention as a complementary module to the channel attention. In simple terms, the spatial attention tells 'where in the channel to focus' and the channel attention tells 'what channel to focus on'. However, the shortcoming in this process is that the channel attention and spatial attention are segregated and computed independent of each other. Thus, any relationship between the two is not considered. Motivated by the way of building spatial attention, we present the concept of *cross dimension interaction*, which addresses this shortcoming by capturing the interaction between the spatial dimensions and the channel dimension of the input tensor. We introduce cross-dimension interaction in triplet attention by dedicating three branches to capture dependencies between the (C, H) , (C, W) and (H, W) dimensions of the input tensor respectively.

Z-pool: The Z-pool layer here is responsible for reducing the zeroth dimension of the tensor to two by concatenating the average pooled and max pooled features across that dimension. This enables the layer to preserve a rich representation of the actual tensor while simultaneously shrinking its depth to make further computation lightweight. Mathe-

matically, it can be represented by the following equation:

$$Z\text{-pool}(\chi) = [\text{MaxPool}_{0d}(\chi), \text{AvgPool}_{0d}(\chi)], \quad (5)$$

where $0d$ is the 0th-dimension across which the max and average pooling operations take place. For instance, the Z-Pool of a tensor of shape $(C \times H \times W)$ results in a tensor of shape $(2 \times H \times W)$.

Triplet Attention: Given the above defined operations, we define *triplet attention* as a three branched module which takes in an input tensor and outputs a refined tensor of the same shape. Given an input tensor $\chi \in \mathbb{R}^{C \times H \times W}$, we first pass it to each of the three branches in the proposed triplet attention module. In the first branch, we build interactions between the height dimension and the channel dimension. To achieve so, the input χ is rotated 90° anti-clockwise along the H axis. This rotated tensor denoted as $\hat{\chi}_1$ is of the shape $(W \times H \times C)$. $\hat{\chi}_1$ is then passed through Z-pool and is subsequently reduced to $\hat{\chi}_1^*$ which is of shape $(2 \times H \times C)$. $\hat{\chi}_1^*$ is then passed through a standard convolutional layer of kernel size $k \times k$ followed by a batch normalization layer, which provides the intermediate output of dimensions $(1 \times H \times C)$. The resultant attention weights are then generated by passing the tensor through a sigmoid activation layer (σ). The attention weights generated are subsequently applied to $\hat{\chi}_1$ and then rotated 90° clockwise along the H axis to retain the original input shape of χ .

Similarly, in the second branch, we rotate χ 90° anti-clockwise along the W axis. The rotated tensor $\hat{\chi}_2$ can be represented with dimension of $(H \times C \times W)$ and is passed through a Z-pool layer. Thus, the tensor is reduced to $\hat{\chi}_2^*$ of the shape $(2 \times C \times W)$. $\hat{\chi}_2^*$ is passed through a standard convolutional layer defined by kernel size $k \times k$ followed by a batch normalization layer which outputs a tensor of the shape $(1 \times C \times W)$. The attention weights are then obtained by passing this tensor through a sigmoid activation layer (σ) which are then simply applied on $\hat{\chi}_2$ and the output is subsequently rotated 90° clockwise along the W axis to retain the same shape as input χ .

For the final branch, the channels of the input tensor χ are reduced to two by Z-pool. This reduced tensor $\hat{\chi}_3$ of shape $(2 \times H \times W)$ is then passed through a standard convolution layer defined by kernel size k followed by a batch normalization layer. The output is passed through sigmoid activation layer (σ) to generate the attention weights of shape $(1 \times H \times W)$ which are then applied to the input χ . The refined tensors of shape $(C \times H \times W)$ generated by each of the three branches are then aggregated by simple averaging.

Summarizing, the process to obtain the refined attention-applied tensor y from triplet attention for an input tensor $\chi \in \mathbb{R}^{C \times H \times W}$ can be represented by the equation:

$$y = \frac{1}{3}(\overline{\hat{\chi}_1 \sigma(\psi_1(\hat{\chi}_1^*))} + \overline{\hat{\chi}_2 \sigma(\psi_2(\hat{\chi}_2^*))} + \chi \sigma(\psi_3(\hat{\chi}_3))), \quad (6)$$

Attention Mechanism	Parameters	Overhead (ResNet-50)
SE [14]	$2C^2/r$	2.514M
CBAM [33]	$2C^2/r + 2k^2$	2.532M
BAM [24]	$C/r(3C + 2k^2C/r + 1)$	0.358M
GC [2]	$2C^2/r + C$	2.548M
Triplet Attention	$6k^2$	0.0048M

Table 1. Comparisons of various attention modules based on their parameter complexity and overhead using ResNet-50 backbone.

where σ represents the sigmoid activation function; ψ_1 , ψ_2 and ψ_3 represent the standard two-dimensional convolutional layers defined by kernel size k in the three branches of triplet attention. Simplifying Eq.(6), y becomes:

$$y = \frac{1}{3}(\overline{\hat{\chi}_1 \omega_1} + \overline{\hat{\chi}_2 \omega_2} + \chi \omega_3) = \frac{1}{3}(\overline{y_1} + \overline{y_2} + y_3), \quad (7)$$

where ω_1 , ω_2 and ω_3 are the three cross-dimensional attention weights computed in triplet attention. The $\overline{y_1}$ and $\overline{y_2}$ in Eq. (7) represents the 90° clockwise rotation to retain the original input shape of $(C \times H \times W)$.

Complexity Analysis: In Tab. 1, we empirically verify the parameter efficiency of triplet attention as compared to other standard attention mechanisms. C represents the number of input channels to the layer, r represents the reduction ratio used in the bottleneck of the MLP while computing the channel attention and the kernel size used for 2D convolution is represented by k ; $k \lll C$. We show that the parameter overhead brought along by different attention layers is much higher as compared to our method. We calculate the overhead on a ResNet-50 [12] by adding the attention layers in each block while fixing r to be 16. k was fixed at 7 for CBAM [33] and triplet attention while for BAM [24] k was set to be 3. The reason for the lower overhead cost for BAM as compared to CBAM, GC [2] and SE [14] is because unlike the latter mentioned attention layers being used in every block, BAM was used only three times across the architecture in total according to the default setting for BAM.

4. Experiments

In this section, we provide the details for experiments and results that demonstrate the performance and efficiency of triplet attention, and compare it with previously proposed attention mechanisms on several challenging computer vision tasks like ImageNet-1k [7] classification and object detection on PASCAL VOC [8] and MS COCO [21] datasets using standard network architectures like ResNet-50 [12] and MobileNetV2 [26]. To further validate our results, we provide the Grad-CAM [27] and Grad-CAM++ [3] results for sample images to showcase the ability of triplet attention to capture more deterministic feature-rich representations.

Method	Backbone	Parameters	FLOPs	Top-1 (%)	Top-5 (%)
ResNet [12]	ResNet-18	11.69M	1.82G	30.20	10.90
	ResNet-50	25.56M	4.12G	24.56	7.50
	ResNet-101	44.46M	7.85G	22.63	6.44
SENet [14]	ResNet-18	11.78M	1.82G	29.41	10.22
BAM [24]		11.71M	1.83G	28.88	10.01
CBAM [33]		11.78M	1.82G	29.27	10.09
Triplet Attention (Ours)		11.69M	1.83G	28.91	10.01
SENet [14]	ResNet-50	28.07M	4.13G	23.14	6.70
BAM [24]		25.92M	4.21G	24.02	7.18
CBAM [33]		28.09M	4.13G	22.66	6.31
GSoP-Net1 [10]		28.29M	6.41G	22.02	5.88
A^2 -Nets [6]		33.00M	6.50G	23.00	6.50
GCNet [2]		28.10M	4.13G	22.30	6.34
GALA [22]		29.40M	-	22.73	6.35
ABN [9]		43.59M	7.66G	23.10	-
SRM [18]		25.62M	4.12G	22.87	6.49
Triplet Attention (Ours)		25.56M	4.17G	22.52	6.32
SENet [14]	ResNet-101	49.29M	7.86G	22.38	6.07
BAM [24]		44.91M	7.93G	22.44	6.29
CBAM [33]		49.33M	7.86G	21.51	5.69
SRM [18]		44.68M	7.85G	21.53	5.80
Triplet Attention (Ours)		44.56M	7.95G	21.97	6.15
MobileNetV2 [26]	MobileNetV2	3.51M	0.32G	28.36	9.80
SENet [14]		3.53M	0.32G	27.58	9.33
CBAM [33]		3.54M	0.32G	30.07	10.67
Triplet Attention (Ours)		3.51M	0.32G	27.38	9.23

Table 2. Single-crop error rate (%) on the ImageNet validation set and complexity comparisons in terms of network parameters (in millions) and floating point operations per second (FLOPs). Other than reporting results on heavy-weight ResNets, we also show results based on light-weight mobile networks. With a negligible increase of learnable parameters, our approach works much better than the baselines and is also comparable to the state-of-the-art methods that need large additional parameters and computations, like GSoP-Net1 [10].

All ImageNet models were trained using 8 Nvidia Tesla V100 GPUs, and all object detection models were trained with 4 Nvidia Tesla P100 GPUs. We did not observe any substantial difference in total wall time between the baseline models and those augmented with triplet attention.

4.1. ImageNet

To train our ResNet [12] based models, we add triplet attention layers at the end of each bottleneck block. We follow the exact training configuration as [12, 14] for consistent and fair comparison with other methods. Similarly, we follow the approach of [26] to train our MobileNetV2-based architecture.

Our results for the validated architectures are shown in Tab. 2. Triplet attention is able to match or outperform other similar techniques, while simultaneously introducing the fewest number of additional model parameters.

A ResNet50-based model augmented with triplet attention achieves a 2.04% improvement in top-1 error rate on ImageNet while only increasing the number of parameters by approximately 0.02% and increasing the FLOPs by $\approx 1\%$. The only comparable model that outperforms triplet attention is GSoP-Net, which provides a 0.5% gain over triplet attention at the cost of 10.7% more parameters and 53.6% more FLOPs.

We observe similar trend in performance in the smaller ResNet-18 model where triplet attention provides a 0.5% improvement in top-1 error rate while only increasing the parametric complexity by 0.02%.

For ResNet-101 based models, triplet attention outperforms both vanilla and SENet variants by 0.66% and 0.41%, respectively. While SRM [18] and CBAM were able to obtain marginally better results than triplet attention, our approach is still the lightest in terms of parameters.

Backbone	Detectors	Parameters	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
ResNet-50 [12]	Faster R-CNN [25]	41.7M	36.4	58.4	39.1	21.5	40.0	46.6
ResNet-101 [12]		60.6M	38.5	60.3	41.6	22.3	43.0	49.8
SENet-50 [14]		44.2M	37.7	60.1	40.9	22.9	41.9	48.2
ResNet-50 + CBAM [33]		44.2M	39.3	60.8	42.8	24.1	43.0	49.8
ResNet-50 + Triplet Attention (Ours)		41.7M	39.3	60.8	42.7	23.4	42.8	50.3
ResNet-50 [12]	RetinaNet [25]	38.0M	35.6	55.5	38.3	20.0	39.6	46.8
SENet-50 [14]		40.5M	37.1	57.2	39.9	21.2	40.7	49.3
ResNet-50 + CBAM [33]		40.5M	38.0	57.7	40.6	22.1	41.9	50.2
ResNet-50 + Triplet Attention (Ours)		38.0M	37.6	57.3	40.0	21.7	41.1	49.7
ResNet-50 [12]	Mask RCNN [11]	44.3M	37.3	59.0	40.2	21.9	40.9	48.1
SENet-50 [14]		46.8M	38.7	60.9	42.1	23.4	42.7	50.0
ResNet-50 + 1 NL block [32]		46.5M	38.0	59.8	41.0	-	-	-
GCNet [10]		46.9M	39.4	61.6	42.4	-	-	-
ResNet-50 + Triplet Attention (Ours)		44.3M	39.8	61.6	42.8	24.3	42.9	51.3

Table 3. **Object detection mAP(%) on the MS COCO validation set.** Triplet Attention results in higher performance gain with minimal computational overhead.

With MobileNetV2, triplet attention provides a 0.98% improvement in top-1 error rate on ImageNet while only increasing parameters by approximately 0.03%. We also observed that CBAM hurts model performance in case of a MobileNetV2 where it drops accuracy by 1.71%. The experimental results demonstrate that the proposed triplet attention works well for both heavy and light-weight models with a negligible increase in parameters and computations. In the following subsection and supplementary materials, we will show the effectiveness of our triplet attention module when applied to other vision tasks, like object detection, instance segmentation, and human key-point detection.

4.2. PASCAL VOC

Method	Detector	AP	AP ₅₀	AP ₇₅
ResNet-50 [12]	Faster R-CNN [25]	46.956	77.521	48.903
ResNet-50 + CBAM [33]		51.398	80.409	54.919
ResNet-50 + TA (Ours)		53.919	80.932	58.810

Table 4. **Object detection mAP(%) on the PASCAL VOC 2012 test set.** Triplet attention results in providing significant improvement in performance with negligible overhead as compared to its counterparts. TA represents Triplet Attention.

For object detection, we utilize our pre-trained ResNet-50 model described in Sec. 4.1 in conjunction with Faster R-CNN [25] with FPN [19] on the Pascal VOC dataset [8]. We adopt default training configuration for the detectron2 toolkit [34] to train a baseline ResNet-50 [12] and ResNet-50 with CBAM [33]. For all models, we train on the 2007 and 2012 versions of the training set and validate on the 2007 validation set as described in [8].

The results can be found in Tab. 4. When compared to the baseline model and its corresponding CBAM variant, our triplet attention module is able to produce a distinct improvement in AP score, beating the baseline ResNet50 by 6.9%, and CBAM by 2.6% while having a backbone that consumes fewer FLOPs and parameters.

4.3. MS COCO

As in Sec. 4.2, using the ImageNet models augmented with triplet attention as backbones, we train Faster-RCNN [25], Mask-RCNN [11], and RetinaNet [20] models to apply our attention module to object detection tasks on the COCO dataset [21]. We use the training procedure described in [20, 25], implemented in the mmdetection framework[4], to ensure a fair test. Our results for the COCO dataset results are summarized in Tab. 3. We observe that triplet attention outperforms most of the similar layers, achieving a higher AP score in multiple categories. Across all architectures, adding a triplet attention module improves the AP score by over 2 points in AP over the baseline model while using the same ImageNet backbone described in Sec. 4.1 that adds a negligible computational overhead. The improvement in performance observed in the experiments showcase the benefit of our cross-dimension interaction strategy in triplet attention.

4.4. Ablation Study on Branches

We further validate the importance of cross-dimension interaction by conducting ablation experiments to observe the impact of the branches in the triplet attention module. In Tab. 5, *spatial off* indicates that the third branch, where the input tensor is not permuted, is switched off, and *channel*

Model	Parameters	FLOPs	Top-1 Accuracy (%)
ResNet-32 [12]	0.464M	3.404G	93.12
ResNet-32 + TA (channel off)	0.466M	3.437G	93.27
ResNet-32 + TA (spatial off)	0.467M	3.415G	93.29
ResNet-32 + TA (full)	0.469M	3.448G	93.56
VGG-16 + BN [28]	15.254M	0.315G	93.25
VGG-16 + BN + TA (channel off)	15.255M	0.315G	93.59
VGG-16 + BN + TA (spatial off)	15.256M	0.32G	93.15
VGG-16 + BN + TA (full)	15.257M	0.32G	93.78
MobileNet-v2 [26]	2.297M	0.095G	93.11
MobileNet-v2 + TA (channel off)	2.302M	0.096G	92.94
MobileNet-v2 + TA (spatial off)	2.308M	0.12G	93.22
MobileNet-v2 + TA (full)	2.313M	0.122G	93.51

Table 5. Effect of different branches in triplet attention on performance in CIFAR-10 classification.

off indicates that the two branches, which involve permutations of the input tensor, are switched off. As shown, the results support our intuition with triplet attention having all three branches switched on, denoted as *full*, to be performing consistently better than the vanilla version and its two counterparts.

4.5. Grad-CAM Visualization




Figure 4. **Visualization of Grad-CAM and Grad-CAM++ results.** The results were obtained for two random samples from the ImageNet validation set and were compared for a baseline ResNet-50, ResNet-50 + CBAM and a ResNet-50 + triplet attention. Ground truth (G.T) labels for the images are provided below the original samples and the networks prediction and confidence scores are provided in the corresponding boxes.

We hypothesize that the cross-dimensional interaction provided by triplet attention helps the network learn more meaningful internal representations of the image. To validate this claim, we provide sample visualizations from the Grad-CAM [27] and Grad-CAM++ [3] techniques, which visualize the gradients of the top-class prediction with respect to the input image as a colored overlay. As shown in Fig. 4, triplet attention is able to capture tighter and more relevant bounds on images from the ImageNet dataset [7]. In certain cases, when using triplet attention, a ResNet50 is able to identify classes that the baseline model fails at predicting correctly. More Grad-CAM based results are presented in the supplementary section.

The visualizations support our understanding of the intrinsic capability of triplet attention to capture richer and more discriminative contextual information for a particular target class. This property of triplet attention is extremely favorable and helpful in improving the performance of deep neural network architectures as compared to their baseline counterparts.

5. Conclusion

In this work, we propose a novel attention layer, triplet attention, which captures the importance of features across dimensions in a tensor. Triplet attention uses an efficient attention computation method that does not have any information bottlenecks. Our experiments demonstrate that triplet attention improves the baseline performance of architectures like ResNet and MobileNet on tasks like image classification on ImageNet and object detection on MS COCO, while only introducing a minimal computational overhead.

We expect that other novel and robust techniques of capturing cross-dimension dependencies when computing attention may improve upon our results while reducing cost. In the future, we plan to investigate the effects of adding triplet attention to more sophisticated architectures like EfficientNets [30] and extend our intuition in the domain of 3D vision.

6. Acknowledgements

The authors would like to offer sincere gratitude to everyone who supported during the timeline of this project including Himanshu Arora from Montreal Institute for Learning Algorithms (MILA), Jaegul Choo and Sanghun Jung from Korea Advanced Institute of Science and Technology (KAIST). This work utilizes resources supported by the National Science Foundation’s Major Research Instrumentation program [16], grant #1725729, as well as the University of Illinois at Urbana-Champaign.

References

- [1] Chunshui Cao, Xianming Liu, Yi Yang, Yinan Yu, Jiang Wang, Zilei Wang, Yongzhen Huang, Liang Wang, Chang Huang, Wei Xu, et al. Look and think twice: Capturing top-down visual attention with feedback convolutional neural networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2956–2964, 2015.
- [2] Yue Cao, Jiarui Xu, Stephen Lin, Fangyun Wei, and Han Hu. Gcnet: Non-local networks meet squeeze-excitation networks and beyond. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [3] 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.
- [4] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. Mmdetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019.
- [5] Long Chen, Hanwang Zhang, Jun Xiao, Liqiang Nie, Jian Shao, Wei Liu, and Tat-Seng Chua. Sca-cnn: Spatial and channel-wise attention in convolutional networks for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5659–5667, 2017.
- [6] Yunpeng Chen, Yannis Kalantidis, Jianshu Li, Shuicheng Yan, and Jiashi Feng. A²-nets: Double attention networks. In *Advances in Neural Information Processing Systems*, pages 352–361, 2018.
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [8] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- [9] Hiroshi Fukui, Tsubasa Hirakawa, Takayoshi Yamashita, and Hironobu Fujiyoshi. Attention branch network: Learning of attention mechanism for visual explanation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [10] Zilin Gao, Jiangtao Xie, Qilong Wang, and Peihua Li. Global second-order pooling convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3024–3033, 2019.
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [12] 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.
- [13] Qibin Hou, Li Zhang, Ming-Ming Cheng, and Jiashi Feng. Strip pooling: Rethinking spatial pooling for scene parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4003–4012, 2020.
- [14] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Jun 2018.
- [15] Zilong Huang, Xinggong Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Ccnet: Criss-cross attention for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 603–612, 2019.
- [16] Volodymyr Kindratenko, Dawei Mu, Yan Zhan, John Maloney, Sayed Hadi Hashemi, Benjamin Rabe, Ke Xu, Roy Campbell, Jian Peng, and William Gropp. Hal: Computer system for scalable deep learning. In *Practice and Experience in Advanced Research Computing*, pages 41–48. 2020.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [18] HyunJae Lee, Hyo-Eun Kim, and Hyeonseob Nam. Srm: A style-based recalibration module for convolutional neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1854–1862, 2019.
- [19] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [20] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [22] Drew Linsley, Dan Shiebler, Sven Eberhardt, and Thomas Serre. Learning what and where to attend. In *International Conference on Learning Representations*, 2018.
- [23] Jiang-Jiang Liu, Qibin Hou, Ming-Ming Cheng, Changhu Wang, and Jiashi Feng. Improving convolutional networks with self-calibrated convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10096–10105, 2020.
- [24] Jongchan Park, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Bam: Bottleneck attention module. *arXiv preprint arXiv:1807.06514*, 2018.
- [25] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [26] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the*

- IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [27] 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.
 - [28] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
 - [29] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
 - [30] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019.
 - [31] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3156–3164, 2017.
 - [32] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018.
 - [33] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *The European Conference on Computer Vision (ECCV)*, September 2018.
 - [34] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
 - [35] Tianjun Xiao, Yichong Xu, Kuiyuan Yang, Jiaying Zhang, Yuxin Peng, and Zheng Zhang. The application of two-level attention models in deep convolutional neural network for fine-grained image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 842–850, 2015.
 - [36] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016.