

Channelized Axial Attention – Considering Channel Relation within Spatial Attention for Semantic Segmentation

Ye Huang¹, Di Kang², Wenjing Jia¹, Liu Liu³ and *Xiangjian He¹

¹ University of Technology Sydney

² Tencent AI LAB

³ University of Sydney

edward.ye.huang@qq.com

Abstract

Spatial and channel attentions, modelling the semantic interdependencies in spatial and channel dimensions respectively, have recently been widely used for semantic segmentation. However, computing spatial and channel attentions separately sometimes causes errors, especially for those difficult cases. In this paper, we propose Channelized Axial Attention (CAA) to seamlessly *integrate* channel attention and spatial attention into a single operation with negligible computation overhead. Specifically, we break down the dot-product operation of the spatial attention into two parts and insert channel relation in between, allowing for independently optimized channel attention on each spatial location. We further develop grouped vectorization, which allows our model to run with very little memory consumption without slowing down the running speed. Comparative experiments conducted on multiple benchmark datasets, including Cityscapes, PASCAL Context, and COCO-Stuff, demonstrate that our CAA outperforms many state-of-the-art segmentation models (including dual attention) on all tested datasets.

Introduction

Semantic segmentation is a fundamental task in many computer vision applications, which assigns a class label to each pixel in the image. Most of the existing approaches (Yuan, Chen, and Wang 2020; Yang et al. 2018; Fu et al. 2019; Li et al. 2019) have adopted a pipeline similar to the one that is defined by Fully Convolutional Networks (FCNs) (Long, Shelhamer, and Darrell 2015) using fully convolutional layers to output the pixel-level segmentation results of input images. These approaches have achieved state-of-the-art performance. After the FCNs, there have been many approaches dedicated to extracting enhanced pixel representations from the backbone. Earlier approaches, including PSPNet (Zhao et al. 2017) and DeepLab (Chen et al. 2018), used a Pyramid Pooling Module or an Atrous Spatial Pyramid Pooling module to expand the receptive field to enhance the representation capabilities. Recently, many works focus on using the attention mechanisms to enhance pixel representations. The first attempts in this direction included Squeeze and Excitation Networks (SENeTs) (Hu, Shen, and Sun 2018) that

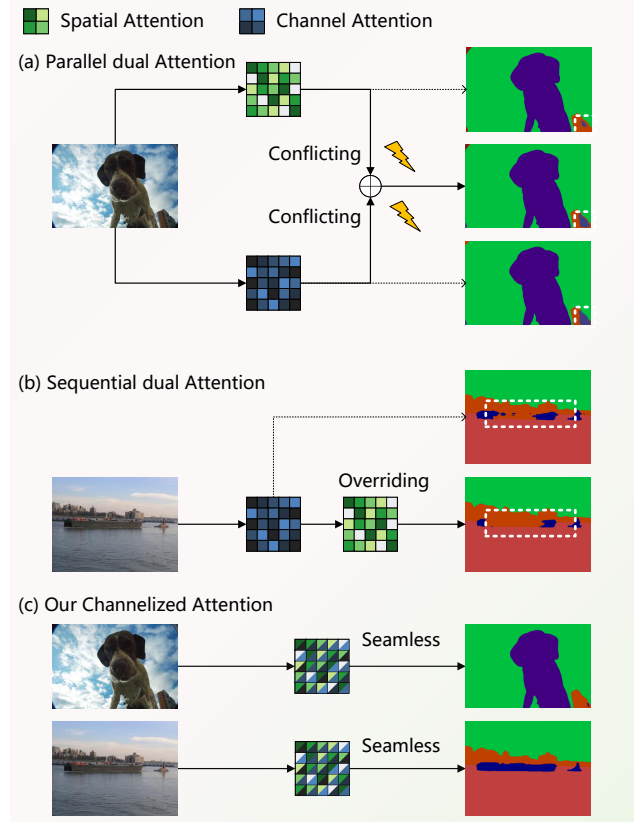


Figure 1: Different dual attention designs: (a) *Parallel dual attention* sums the results from spatial and channel attentions directly, which may cause conflicts because spatial and channel attentions are focusing on different aspects. (b) *Sequential dual attention* performs spatial attention after channel attention, where the spatial attention may override correct features extracted by the channel attention. (c) *Our channelized attention* seamlessly merges the spatial and channel attentions into a single operation, removing the potential conflicting issue caused by a or b.

*Corresponding author.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

introduced a simple yet effective channel attention module to explicitly model the interdependencies between channels. Meanwhile, spatial attention relied on self-attention pro-

posed in (Wang et al. 2018; Vaswani et al. 2017) to model long-range dependencies in spatial domain, so as to produce more correct pixel representations. For each pixel in the feature maps, spatial attention “corrects” its representation with the representations of other pixels depending on their similarity. In contrast, channel attention identifies important channels based on all spatial locations and reweights the extracted features.

Parallel dual attention (e.g., (Fu et al. 2019)) was proposed to gain the advantages of both spatial attention and channel attention. This approach directly fused their results with an element-wise addition (see Fig. 1(a)). Although they have achieved improved performance, the relationship between the contributions of spatial and channel attentions to the final results is unclear. Moreover, calculating the two attentions separately not only increases the computational complexity, but may also result in conflicting importance of feature representations. For example, some channels may appear to be important in spatial attention for a pixel that belongs to a partial region in the feature maps. However, channel attention may have its own perspective, which is calculated by summing up the similarities over the entire feature maps, and weakens the impact of spatial attention.

Sequential dual attention, which combines channel attention and spatial attention in a sequential manner (Fig. 1(a)) has similar issues. For example, channel attention can *ignore* the partial region representation obtained from the overall perspective. However, this partial region representation may be required by spatial attention. Thus, directly fusing the spatial and channel attention results may yield incorrect importance weights for pixel representations.

In order to overcome the aforementioned issues, we propose Channelized Axial Attention (CAA), which breaks down the axial attention into more basic parts and inserts channel attention into them, combining spatial attention and channel attention together seamlessly and efficiently. Specifically, when applying the axial attention maps to the input signal (Wang et al. 2018), we capture the intermediate results of the dot product before they are summed up along the corresponding axes. Capturing these intermediate results allows channel attention to be integrated for each column and each row, instead of computing on the mean or sum of the features in the entire feature maps. We also develop a novel grouped vectorization approach to maximize the computation speed in limited GPU memory.

In summary, our contributions in this paper include:

- We are the first to explicitly *identify* the potential conflicts between spatial and channel attention in existing dual attention designs by *visualizing* the effects of each attention on the final result.
- We propose a novel Channelized Axial Attention, which breaks down the axial attention into more basic parts and inserts channel attention in between, integrating spatial attention and channel attention together seamlessly and efficiently, with only a minor computation overhead compared to the original axial attention.
- To balance the computation speed and GPU memory usage, a grouped vectorization approach for computing the

channelized attentions is proposed. This is particularly advantageous when processing large images.

- Experiments on three challenging benchmark datasets, including PASCAL Context (Everingham et al. 2009), COCO-Stuff (Caesar, Uijlings, and Ferrari 2018) and Cityscapes (Marius et al. 2016), demonstrate the superiority of our approach over the SOTA approaches.

Related Work

Spatial attention. Non-local networks (Wang et al. 2018) and Transformer (Vaswani et al. 2017) introduced the self-attention mechanism to examine the pixel relationship in the spatial domain. It usually calculates dot-product similarity or cosine similarity to obtain the similarity measurement between every two pixels in feature maps, and recalculates the feature representation of each pixel according to its similarity with others. Self-attention has successfully addressed the feature map coverage issue of multiple fixed-range approaches (Chen et al. 2017; Zhao et al. 2017), but it has also introduced huge computation costs for computing the complete feature map. This means that, for each pixel in the feature maps, its attention similarity affects all other pixels. Recently, many approaches (Huang et al. 2020; Zhu et al. 2019) have been developed to optimize the GPU memory costs of spatial self-attention.

Channel Attention. Channel attention (Hu, Shen, and Sun 2018) examined the relationships between channels, and enhanced the important channels so as to improve performance. SENets (Hu, Shen, and Sun 2018) conducted a global average pooling to get mean feature representations, and then went through two fully connected layers, where the first one reduced channels and the second one recovered the original channels, resulting in channel-wise weights according to the importance of channels. In DANet (Fu et al. 2019), channel-wise relationships were modelled by self-attention, similar to the it is used in the spatial domain, except that it computed the attention with a dimension of $C \times C$ rather than $(H \times W) \times (H \times W)$ (here, C represents the number of channels, and H and W represent the height and width of the feature maps, respectively).

Spatial Attention + Channel Attention. Combining spatial attention and channel attention can provide fully optimized pixel representations in a feature map. However, it is not easy to enjoy both advantages seamlessly. In DANet (Fu et al. 2019), the results of the channel attention and spatial attention are directly added together. Supposing that there is a pixel belonging to a semantic class that has a tiny region in the feature maps, spatial-attention can find its similar pixels. However, channel representation of the semantic class with a partial region of the feature maps may not be important in the perspective of entire feature maps, so it may be ignored when conducting channel attention computations. Computing self-attention and channel attention separately (as illustrated in Fig. 1(a)) can cause conflicting results, and thus weaken their performance when both results are summarized together. Similarly, in the cascaded model (see Fig. 1(b)), the spatial attention module after the channel attention module may pick up an incorrect pixel representation enhanced by

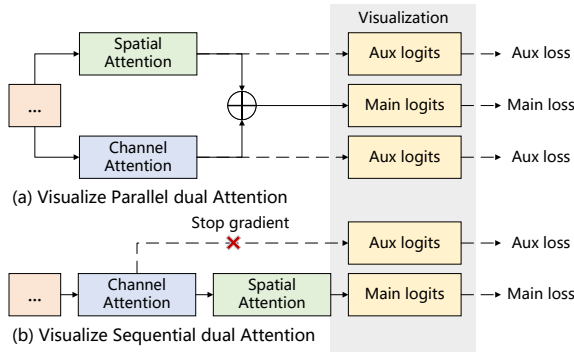


Figure 2: Our designs for visualizing the effects of dual attentions in parallel and sequential.

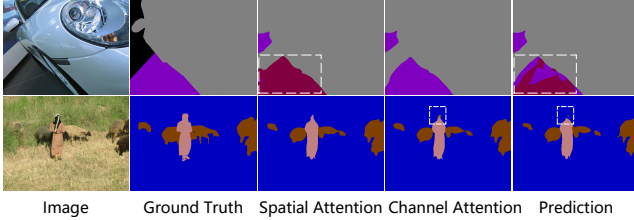


Figure 3: Conflicting features in parallel dual attention designs. *Top*: The bad spatial attention representation negatively influences the good channel attention representation. *Bottom*: The bad channel attention representation negatively influences the good spatial attention representation. See the boxed areas.

channel attention, because channel attention computes channel importance according to the entire feature maps.

Exploring Conflicting Features

As we have analyzed earlier, computing spatial and channel attentions separately can cause conflicting features. In our experiments, to illustrate this feature conflicting issue faced by existing dual attention approaches, we designed a simple way to visualize the effects of spatial attention and channel attentions on pixel representation.

Visualizing Conflicts For a parallel dual attention design such as DANet (Fu et al. 2019), since it has two auxiliary losses for each of spatial attention and channel attention, we directly use their logits during inference to generate corresponding segmentation results and compare them with the result generated by the main logits. For a sequential dual attention design, we add an extra branch that directly uses the pixel representation obtained from channel attention to perform the segmentation logits. Note that, since the original sequential design does not have independent logits after the channel attention module, we stop the gradient from back-propagating to the main branch, to ensure that our newly added branch has no effect on the main branch.

Examples of Conflicting Features To visualize the impact

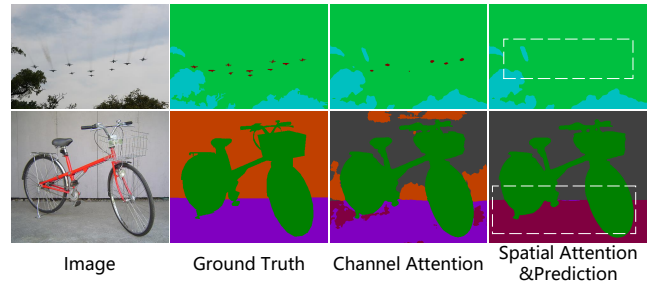


Figure 4: In sequential dual attention designs, the spatial attention representation (the 4th column) ignores the correct channel attention representation (the 3rd column).

of the feature conflicting issue in the existing dual attention designs, we present examples of the segmentation results obtained with the conflicting features in the parallel dual attention design (Fig. 3) and the sequential dual attention design (Fig. 4). As observed from Fig. 3, the parallel design of dual attention directly sums up the pixel representations obtained from spatial attention and channel attention. With this approach, the advantages of the pixel representations obtained from one can be weakened by the other. The sequential way of combining the dual attentions avoids taking their average but still has its own problem. As shown in Fig. 4, the pixel representation obtained from the spatial attention ignores the correct representation obtained from the channel attention, and worsens the prediction.

Methods

Preliminaries

Formulation of the Spatial Self-attention Following Non Local (Wang et al. 2018) and Stand Alone Self Attention (Ramachandran et al. 2019), a 2D self-attention operation in spatial domain can be defined by:

$$\mathbf{y}_{i,j} = \sum_{\forall m,n} f(\mathbf{x}_{i,j}, \mathbf{x}_{m,n})g(\mathbf{x}_{m,n}). \quad (1)$$

Here, a pairwise function f computes the similarity between the pixel representation $\mathbf{x}_{i,j}$ (*query*) at position (i, j) and the pixel representation $\mathbf{x}_{m,n}$ (*key*) at all other possible positions (m, n) . The unary function g maps the original representation at position (m, n) to a new domain (*value*). In our work, we use the similarity function (Wang et al. 2018) as f , i.e.,

$$f(\mathbf{x}_{i,j}, \mathbf{x}_{m,n}) = \text{softmax}_{m,n}(\theta(\mathbf{x}_{i,j})^T \theta(\mathbf{x}_{m,n})), \quad (2)$$

where θ is a 1×1 convolution layer transforming the feature maps \mathbf{x} to a new domain to calculate dot-product similarity (Wang et al. 2018) between every two pixels. Note that, following a common practice (Li et al. 2020), we use the same 1×1 convolution weights for both query and key. Then, these similarities are used as the weights (Eq. (1)) to aggregate features of all pixels, producing an enhanced pixel representation $\mathbf{y}_{i,j}$ at position (i, j) .

Formulation of the Axial Attention From the above equations, we can see the computational complexity of the self-attention module is $O(H^2W^2)$, which requires large

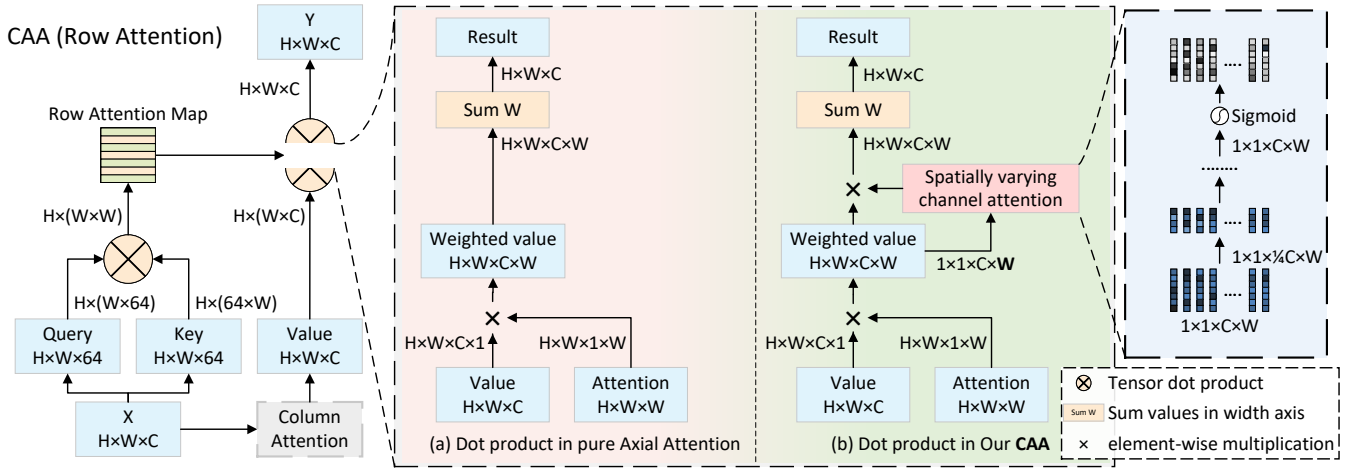


Figure 5: The detailed architecture of the proposed CAA (Row Attention). We present the way to apply channel attention seamlessly in (b). We mark the independent spatial dimension in **bold style**. This allows channel attention to also consider spatial unique information. *Note that*, in our design, the “value” for row attention is obtained from the result of column attention. See Eq. 11 for details.

computation resources and prevents real-time applications, such as autopilot. Several subsequent works (Huang et al. 2020; Ho et al. 2019) focused on reducing the computational complexity while maintaining high accuracy. In this work, we adopt axial-attention to perform spatial attention. In axial attention, the attention map is calculated for the column and row that cover the current pixel, reducing the computational complexity to $O(HW^2 + H^2W)$.

For convenience, we call the attention values calculated along the Y axis “column attention”, and the attention values calculated along the X axis “row attention”. Similar to Eq. 2, we define axial similarity functions by:

$$A_{\text{col}}(\mathbf{x}_{i,j}, \mathbf{x}_{m,j}) = \text{softmax}_m \left(\theta(\mathbf{x}_{i,j})^T \theta(\mathbf{x}_{m,j}) \right), \quad m \in [H]^1, \quad (3)$$

$$A_{\text{row}}(\mathbf{x}_{i,j}, \mathbf{x}_{i,n}) = \text{softmax}_n \left(\phi(\mathbf{x}_{i,j})^T \phi(\mathbf{x}_{i,n}) \right), \quad n \in [W]. \quad (4)$$

Note that we use different feature transformations (θ, ϕ) for the row and column attention calculations.

With the column and row attention maps A_{col} and A_{row} , the final value weighted by the column and row attention maps can be represented by:

$$\mathbf{y}_{i,j} = \sum_{\forall n} \left(A_{\text{row}}(\mathbf{x}_{i,j}, \mathbf{x}_{i,n}) \left(\sum_{\forall m} A_{\text{col}}(\mathbf{x}_{i,j}, \mathbf{x}_{m,j}) g(\mathbf{x}_{m,n}) \right) \right). \quad (5)$$

Channelized Axial Attention

In order to address the feature conflicting issue of the existing dual attention designs, we propose a novel *Channelized Axial Attention* (CAA), which seamlessly combines the advantages of spatial attention and channel attention.

As mentioned in the above sections, feature conflicts may be caused by the different interests of spatial and channel attentions. Ideally, channel attention should not ignore the regional features that are interesting to spatial attention. Conversely, spatial attention should consider channel relation as

¹We use $i \in [n]$ to denote that i is generated from $[n] = \{1, 2, \dots, n\}$.

well. Thus, we propose to compute channel attention within spatial attention. Specifically, we firstly break down spatial attention into more basic parts. Then, spatially varying channel attention is generated with $\alpha_{i,j,m,n}$ and $\beta_{i,j,n}$. In this way, channel attention is incorporated into spatial attention and spatial attention will not be ignored when small objects exist, seamlessly and effectively combining spatial and channel attention together.

Breaking Down Axial Attention. For convenience, we firstly define two variables $\alpha_{i,j,m,n}$ and $\beta_{i,j,n}$ to represent the *intermediate weighted features* as follows:

$$\alpha_{i,j,m,n} = A_{\text{col}}(\mathbf{x}_{i,j}, \mathbf{x}_{m,j}) g(\mathbf{x}_{m,n}) \quad (6)$$

$$\beta_{i,j,n} = A_{\text{row}}(\mathbf{x}_{i,j}, \mathbf{x}_{i,n}) \sum_{\forall m} \alpha_{i,j,m,n}. \quad (7)$$

Thus, Eq. (5) can be rewritten as:

$$\mathbf{y}_{i,j} = \sum_{\forall n} \beta_{i,j,n} = \sum_{\forall n} A_{\text{row}}(\mathbf{x}_{i,j}, \mathbf{x}_{i,n}) \left(\sum_{\forall m} \alpha_{i,j,m,n} \right). \quad (8)$$

Eqs. (6), (7) and (8) show that the computation of the dot product is composed of two steps: 1) *Reweighting*: reweighting features on selected locations by column attention as in Eq. (6) and row attention as in Eq. (7), and 2) *Summation*: summing the elements along row and column axes in Eq. (8). Note that, this breakdown is also applicable to regular self-attention (see Table 3 and **Appendix**).

Spatially Varying Channel Attention. With the intermediate results $\alpha_{i,j,m,n}$ and $\beta_{i,j,n}$ in Eqs. (6) and (7), channel relation can be applied inside spatial attention, seamlessly combining them into one operation. In addition, channel attention is now independently conducted on each column or row (on each pixel in regular self-attention) and provides spatial perspective for the channel relation modeling, resulting in our *spatially varying channel attention*. Enhanced

with spatially varying channel attentions, now C_{col} and C_{row} are written as:

$$C_{\text{col}}(\alpha_{i,j,m,n}) = \text{Sigmod} \left(\text{ReLU} \left(\frac{\sum_{\forall m,j} (\alpha_{i,j,m,n})}{H \times W} \omega_{c1} \right) \omega_{c2} \right) \alpha_{i,j,m,n}, \quad (9)$$

$$C_{\text{row}}(\beta_{i,j,n}) = \text{Sigmod} \left(\text{ReLU} \left(\frac{\sum_{\forall i,n} (\beta_{i,j,n})}{H \times W} \omega_{r1} \right) \omega_{r2} \right) \beta_{i,j,n}, \quad (10)$$

where $\text{Sigmod}(\cdot)$ is the learned channel attention, and ω_{c1} , ω_{c2} , ω_{r1} and ω_{r2} are the learned relationships between different channels according to $\alpha_{i,j,m,n}$ and $\beta_{i,j,n}$.

Thus, instead of directly using $\alpha_{i,j,m,n}$ and $\beta_{i,j,n}$ as in Eq. (8), for each column and row, we obtain the channelized axial attention features, where the intermediate results $\alpha_{i,j,m,n}$ and $\beta_{i,j,n}$ are weighted by the spatially varying channel attention defined in Eqs. (9) and (10) as:

$$y_{i,j} = \sum_{\forall n} C_{\text{row}} \left(A_{\text{row}}(\mathbf{x}_{i,j}, \mathbf{x}_{i,n}) \left(\sum_{\forall m} C_{\text{col}}(\alpha_{i,j,m,n}) \right) \right). \quad (11)$$

Note that the spatially varying channel attention keeps a W dimension after averaging $H \times W$ during the channel attention (Fig. 5). Now each row has its own channel attention thanks to the breaking down of spatial axial attention.

Going Deeper in Channel Attention. Similar to the work in (Hu, Shen, and Sun 2018), we use two fully connected layers, followed by ReLU and sigmoid activations respectively, to first reduce the channel number and then increase it to the original channel number. To further boost performance, we explore the design of more powerful channel attention modules for our channelization since our attention module keeps the spatial dimension, and thus contains more information than a regular SE module ($1 \times 1 \times C \times W$ or $H \times 1 \times 1 \times C$, see Fig. 5).

We experimented with increased depth and/or width of hidden layers to enhance the capacity of spatial varying channel attention. We find that deeper hidden layers allow channel attention to find a better relationship between channels for our spatially varying channel attention. Moreover, increasing layer width is not as effective as adding layer depth (see Table 1).

Furthermore, in spatial domain, each channel of a pixel contains unique information that can lead to a unique semantic representation. We find that using Leaky ReLU (Mass, Hannun, and Ng 2013) is more effective than ReLU in preventing the loss of information along deeper activations (Sandler et al. 2018). Apparently, this replacement only works in spatially varying channel attention.

Grouped Vectorization. Computing spatial attention row by row and column by column can save computation but it is still too slow (about 2.5 times slower on a single V100 with feature map size = 33×33) even with parallelization. Full vectorization can achieve a very high speed but it has a high requirement on GPU memory (about 2 times larger GPU memory usage than no vectorization on a single V100 with feature map size = 33×33) for storing the intermediate partial axial attention results α (which has a dimension of $H \times H \times W \times C$) and β (which has a dimension of $W \times H \times W \times C$) in Eqs. (6) and (7). To enjoy the high speed

benefit of vectorization with limited GPU memory usage, in our implementation we propose *grouped* vectorization to dynamically batch rows and columns into multiple groups, and then perform vectorization for each group individually.

Experiments

To demonstrate the effectiveness for accuracy of the proposed CAA, comprehensive experimental results are compared with the state-of-the-art methods on three benchmark datasets, *i.e.*, PASCAL Context (Everingham et al. 2009), COCO-Stuff (Caesar, Uijlings, and Ferrari 2018) and Cityscapes (Marius et al. 2016). Experimental results show that our CAA outperforms the state-of-the-art methods on all tested datasets.

Implementation Details

Architecture Our network is built on ResNet-101 (He et al. 2016) pre-trained on ImageNet, and then followed by our CAA module (or axial + SE for ablation study), as well as the image-level block and auxiliary head in (Zhang et al. 2019; Fu et al. 2019). The original ResNet results in a feature map of $1/32$ of the input size. Following other works (Chen et al. 2018; Li et al. 2019), we apply dilated convolution at the output stride = 16 for ablation experiments if not specified. We conduct experiments with the output stride = 8 to compare with the state-of-the-art methods.

Naive Upsampling Unless otherwise specified, we directly bi-linearly upsampled the logits to the input size without refining using any low-level and high resolution features.

Training Settings We employ stochastic gradient descent (SGD) for optimization, where the polynomial decay learning rate policy $(1 - \frac{iter}{maxiter})^{0.9}$ is applied with an initial learning rate = 0.01. We use synchronized batch normalization with batch size = 16 (8 for Cityscapes) during training. For data augmentation, we only apply the most basic data augmentation strategies as in (Chen et al. 2018), including random flip, random scale and random crop.

Results on PASCAL Context

The PASCAL Context (Mottaghi et al. 2014) dataset has 59 classes with 4,998 images for training and 5,105 images for testing. We train our CAA on the training set for 40k iterations. In the following, we first present a series of ablative experiments to show the effectiveness of our method. Then, quantitative and qualitative comparisons with other state-of-the-art methods are presented.

Effectiveness of the Proposed Channelization We first report the impact of adding channelized axial attention and with different depth and width in Table 1, where ‘-’ for the baseline result indicates no channelization is performed.

As can be seen from Table 2, our proposed channelization improves the mIOU over the baseline regardless of the layer counts and the number of channels used. In particular, the best performance is achieved when the Layer Counts = 5 and the number of Channels = 128.

We also compare our model with the sequential design of ‘‘Axial Attention + SE’’, as shown in Table 2. We find the

Layer Counts	# of Channels	mIOU (%)	FLOPs
-	-	50.27(± 0.2)	68.7G
1	128	50.75(± 0.2)	+0.00024G
3	128	50.85(± 0.2)	+0.00027G
5	128	51.06(± 0.2)	+0.00030G
7	128	50.40(± 0.3)	+0.00043G
5	64	50.12(± 0.2)	+0.00015G
5	256	50.35(± 0.4)	+0.00098G

Table 1: Results w/o using channelization (Row 1) and w/ channelization with different layer counts and channel numbers. Numbers in parentheses indicate standard deviations.

Axial Attention	+ SE	+ Our Channelization
50.27(± 0.2)	50.37(± 0.2)	51.06(± 0.2)

Table 2: Result comparison between axial attention, axial attention + SE and channelized axial attention.

sequential design brings only marginal contributions to performance, showing that our proposed channelization method can combine the advantages of both spatial attention and channel attention more effectively. In Table 5, results obtained with *other backbones* are provided to demonstrate the effectiveness and robustness of CAA.

Channelized Self-Attention In this section, we conduct experiments on the PASCAL Context by applying channelization to the original self-attention. We report its single-scale performance in Table 3 with ResNet-101. Our channelized method can also further improve the performance of self-attention by 0.67% (51.09% vs 50.42%).

Impact of the Testing Strategies We compare the performance and computation cost of our proposed model against the baseline and DANet (Fu et al. 2019) with different testing strategies in Table 4. Using the same settings as in other works (Fu et al. 2019), we add multi-scale, left-right flip and auxiliary loss during inference. The accuracies of CAA are further boosted with output stride = 8 since the channel attention can learn and optimize three times more pixels.

Comparison with the State-of-the-art Finally, in Table 6, we compare our approach with the state-of-the-art approaches. Like other similar works, we apply multi-scale and left-right flip during inference. For a fair comparison, we only compare with the methods that use ResNet-101 and naive upsampling in the main paper. More results using alternative backbones are included in Table 5.

As shown in this table, our proposed CAA outperforms all listed state-of-the-art models that are trained with an output stride = 8. Our CAA also performs better than EMANet and SPYGR that are trained with output stride = 16. Note that, in this and the following tables, we report the best results of our approach obtained in experiments.

In Fig. 6, we show some results obtained by our CAA model, FCN, and Dual attention. Our model is able to han-

Attention Base	Eval OS	Channelized	mIOU (%)
Axial Attention	16 16	✓	50.27 51.06
Self Attention	16 16	✓	50.42 51.09

Table 3: Ablation study of applying our Channelized Attention on self-attention with ResNet-101. *Eval OS*: Output strides (Chen et al. 2018) during evaluation.

Methods	OS	MF	Aux	mIOU (%)	FLOPs
ResNet-101	16 8		- -	- -	59.85G 190.70G
DANet	8 8	✓	✓	52.60	+101.25G -
Axial Attention	16 16 8 8	✓		50.27(± 0.2) 52.01(± 0.2) 51.24(± 0.2) 52.51(± 0.2)	+8.85G - +34.33G -
Our CAA	16 16 8 8	✓		51.06(± 0.2) 53.09(± 0.3) 52.73(± 0.1) 54.05(± 0.1)	+8.85G - +34.33G -
Our CAA + Aux loss	16 16 8 8	✓	✓ ✓ ✓	51.80(± 0.2) 53.52(± 0.2) 53.48(± 0.3) 54.65(± 0.4)	+8.85G - +34.33G -

Table 4: Comparison results with different testing strategies. *OS*: Output stride in training and inference. *MF*: Apply multi-scale and left-right flipping during inference. *Aux*: Add auxiliary loss during training. “+” refers to the extra FLOPs over the baseline FLOPs of ResNet-101.

dle previous failure cases better, especially when some class A covering only a smaller region is surrounded by another class B covering a much larger region.

Results on the COCO-Stuff 10K

Following the other works (Fu et al. 2019), we evaluate our CAA on COCO-Stuff 10K dataset (Caesar, Uijlings, and Ferrari 2018), which contains 9,000 training images and 1,000 testing images with 171 classes. Our model is trained for 40k iterations. As shown in Table 7, our proposed CAA outperforms all other state-of-the-art approaches by a large margin of 1.3%, demonstrating that our model can better handle complex images with a large number of classes.

COCO-Stuff 10K results in Fig. 7 compare FCN, DANet, and our CAA. As it can be seen, our CAA can segment common objects such as building, human, or sea very well.

Results on the Cityscapes

The Cityscapes dataset (Marius et al. 2016) has 19 classes. Following previous works (Fu et al. 2019), we train 80k iteration on the *fine* set with a crop size of 769×769 . We report

Backbone	OS	AA	C	mIOU (%)
ResNet-50 (He et al. 2016)	16	✓	✓	49.73
	16	✓	✓	50.23
Xception65 (Chollet 2017)	16	✓	✓	52.42
	16	✓	✓	52.65
EfficientNetB7 (Tan and Le 2019)	16	✓	✓	57.24
	16	✓	✓	57.93
	8	✓	✓	58.40

Table 5: Ablation study of other backbones. All results are obtained in single scale without flipping. *OS*: Output strides during evaluation. *AA*: Axial Attention. *C*: Channelized.

Methods	Backbone	mIOU (%)
ENCNet (Zhang et al. 2018)	ResNet-101	51.7
ANNet (Zhu et al. 2019)	ResNet-101	52.8
EMANet (Li et al. 2019)	ResNet-101	53.1
SPYGR (Li et al. 2020)	ResNet-101	52.8
CPN (Yu et al. 2020)	ResNet-101	53.9
CFNet (Zhang et al. 2019)	ResNet-101	54.0
DANet (Fu et al. 2019)	ResNet-101	52.6
Our CAA (OS = 16)	ResNet-101	53.7
Our CAA (OS = 8)	ResNet-101	55.0

Table 6: Comparisons with other state-of-the-art approaches on the PASCAL Context test set. For a fair comparison, all compared methods used ResNet-101 and naive upsampling.

our results on the *test* set in Table 8. Results show our CAA is also working well on high-resolution images. In Fig. 8, we compare the results on Cityscapes validation set predicted by DANet and our CAA. Key areas of difference are highlighted with white boxes. Results show that many errors produced by DANet no longer exist in our CAA.

Effectiveness of Our Grouped Vectorization

We introduced the grouped vectorization to split tensors into multiple groups so as to reduce the GPU memory usage when performing channel attention inside spatial attention. As we use more groups in group vectorization, the proportionally less GPU memory is needed for the computation. However, longer running time is required. In this section, we conduct experiments to show the variation of the inference time (seconds/image) when different numbers of groups are used. Fig. 9 shows the results of three different input resolutions. As shown in this graph, when splitting the vectorization into smaller numbers of groups, *e.g.*, 2 (or 4), our grouped vectorization achieves similar inference speed but requires one half (or one quarter) of the original spatial complexity. For example, separating into 4 groups has similar inference speed with no separation (*i.e.*, 1 group).

Stronger Backbone on the PASCAL Context

In this section, we report the performance of our CAA with EfficientNet-B7 (Tan and Le 2019) in Table 9. The “sim-

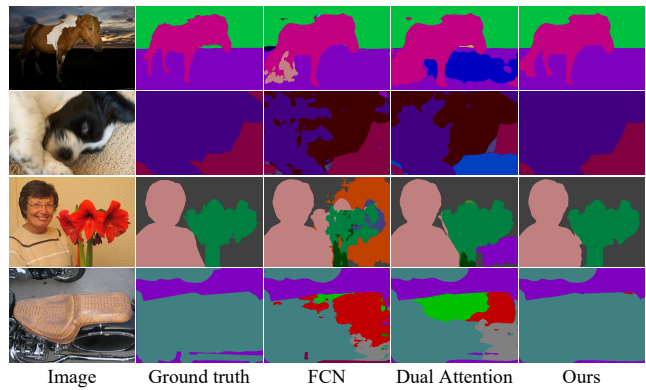


Figure 6: Examples of the segmentation results obtained on the PASCAL Context dataset using FCN, DANet and CAA.

Methods	Backbone	mIOU (%)
SVCNet (Ding et al. 2019)	ResNet-101	39.6
EMANet (Li et al. 2019)	ResNet-101	39.9
SPYGR (Li et al. 2020)	ResNet-101	39.9
OCR (Yuan, Chen, and Wang 2020)	ResNet-101	39.5
DANet (Fu et al. 2019)	ResNet-101	39.7
Our CAA	ResNet-101	41.2

Table 7: Comparisons with other state-of-the-art approaches on the COCO-Stuff 10K test set. For a fair comparison, all compared methods use ResNet-101 and naive upsampling.

ple decoder” used in our method merges the low level features from output stride = 4, during the final upsampling (see (Chen et al. 2018) for details). Note that, the listed methods were not trained under the same settings, or using the same backbone. Our method outperforms the SOTA transformer-based hybrid models, including SETR (Sixiao et al. 2021) and DPT (Ranftl, Bochkovski, and Koltun 2021).

Results on the COCOStuff-164k

Segformer (Xie et al. 2021) used COCOStuff-164k (164,000 images), *i.e.*, the full set of COCOStuff-10k to validate its performance for the first time. In this section, we use “EfficientNet-B5 + CAA” to verify the robustness of our CAA on COCOStuff-164k. Table 10 shows that our method outperforms Segformer and SETR (Sixiao et al. 2021) by a large margin, indicating that our CAA maintains its superior performance well on large-scale training data.

Conclusion

In this paper, we have proposed a novel and effective Channelized Axial Attention, effectively combining the advantages of the popular spatial-attention and channel attention. Experiments on the three popular benchmark datasets have demonstrated the effectiveness of our proposed CAA.

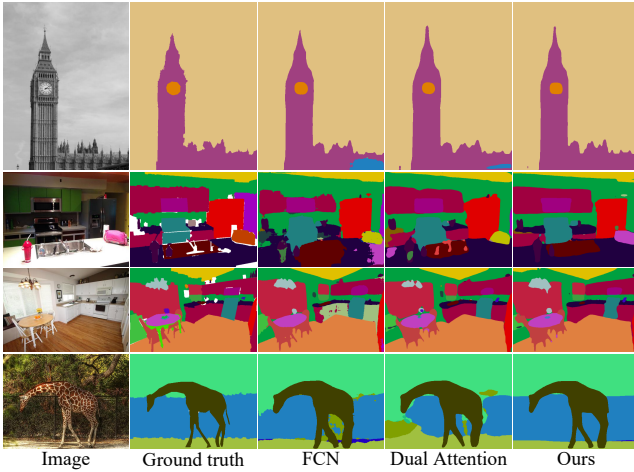


Figure 7: Examples of the results obtained on the COCO-Stuff 10K dataset with our proposed CAA in comparison to the results obtained with FCN, DANet and the ground truth.

Methods	Backbone	mIOU (%)
CFNet (Zhang et al. 2019)	ResNet-101	79.6
ANNet (Zhu et al. 2019)	ResNet-101	81.3
CCNet (Huang et al. 2020)	ResNet-101	81.4
CPN (Yu et al. 2020)	ResNet-101	81.3
SPYGR (Li et al. 2020)	ResNet-101	81.6
OCR (Yuan, Chen, and Wang 2020)	ResNet-101	81.8
DANet (Fu et al. 2019)	ResNet-101	81.5
Our CAA	ResNet-101	82.6

Table 8: Comparisons with other state-of-the-art approaches on the Cityscapes Test set. For a fair comparison, all compared methods use ResNet-101 and naive upsampling.

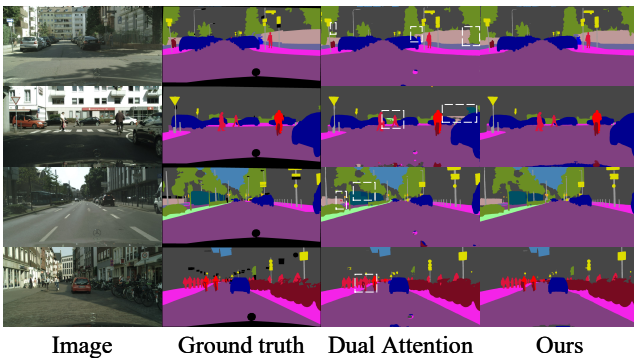


Figure 8: Examples of the segmentation results obtained on the Cityscapes validation set with our proposed CAA in comparison to the results obtained with DANet and the ground truth.

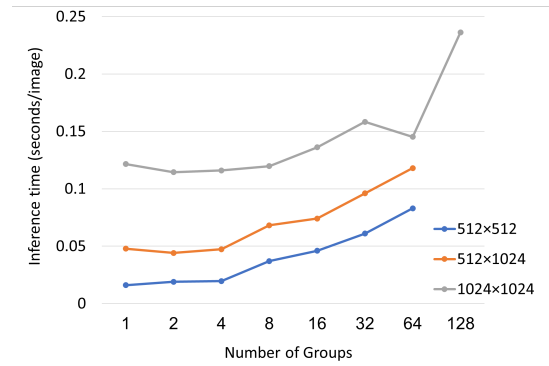


Figure 9: Inference time (seconds/image) tested on three different input resolution when applying different numbers of groups in our proposed grouped vectorization.

Methods	mIOU (%)
CTNet (Li, Sun, and Tang 2021) + JPU	55.5
SETR-MLA (Sixiao et al. 2021)	55.83
ResNeSt-269 (Zhang et al. 2020) + DeepLab V3+	58.9
HRNetV2 + OCR + RMI	59.6
DPT (Ranftl, Bochkovskiy, and Koltun 2021)	60.46
Our CAA (EfficientNet-B7, w/o decoder)	60.12
Our CAA (EfficientNet-B7 + simple decoder)	60.50

Table 9: Result comparison with the state-of-the-art approaches on the PASCAL Context test set for multi-scale prediction. Note that, the listed methods were not trained under the same settings, or using same backbone.

Methods	mIOU (%)
ResNet-50 + DeepLabV3+ (Chen et al. 2018)	38.4
HRNetV2 + OCR	42.3
SETR (Sixiao et al. 2021)	45.8
Segformer-B5 (Xie et al. 2021)	46.7
Our CAA (EfficientNet-B5)	47.30

Table 10: Result comparison with the state-of-the-art approaches on the COCO-Stuff-164K test set.

Acknowledgments

We thank TPU Research Cloud (TRC) for providing TPUs.

References

- Caesar, H.; Uijlings, J.; and Ferrari, V. 2018. COCO-Stuff: Thing and Stuff Classes in Context. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2017. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Chen, L.-C.; Zhu, Y.; Papandreou, G.; Schroff, F.; and Adam, H. 2018. Encoder-Decoder with atrous separable convolution for semantic image segmentation. In *European Conference on Computer Vision*.
- Chollet, F. 2017. Xception: Deep Learning With Depthwise Separable Convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Ding, H.; Jiang, X.; Shuai, B.; Liu, A. Q.; and Wang, G. 2019. Semantic Correlation Promoted Shape-Variant Context for Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Everingham, M.; Gool, L. V.; K.I.Williams, C.; Winn, J.; and Zisserman, A. 2009. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision*.
- Fu, J.; Liu, J.; Tian, H.; Li, Y.; Bao, Y.; Fang, Z.; and Lu, H. 2019. Dual Attention Network for Scene Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Ho, J.; Kalchbrenner, N.; Weissenborn, D.; and Salimans, T. 2019. Axial Attention in Multidimensional Transformers. *arXiv:1912.12180*.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Huang, Z.; Wang, X.; Wei, Y.; Huang, L.; Shi, H.; Liu, W.; and Huang, T. S. 2020. CCNet: Criss-Cross Attention for Semantic Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Li, X.; Yang, Y.; Zhao, Q.; Shen, T.; Lin, Z.; and Liu, H. 2020. Spatial Pyramid Based Graph Reasoning for Semantic Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Li, X.; Zhong, Z.; Wu, J.; Yang, Y.; Lin, Z.; and Liu, H. 2019. Expectation-Maximization Attention Networks for Semantic Segmentation. In *International Conference on Computer Vision*.
- Li, Z.; Sun, Y.; and Tang, J. 2021. CTNet: Context-based Tandem Network for Semantic Segmentation. *arXiv preprint arXiv:2104.09805*.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Marius, C.; Mohamed, O.; Sebastian, R.; Timo, R.; Markus, E.; Rodrigo, B.; Uwe, F.; Roth, S.; and Bernt, S. 2016. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Mass, A. L.; Hannun, A. Y.; and Ng, A. Y. 2013. Rectifier Nonlinearities Improve Neural Network Acoustic Models. In *International Conference on Machine Learning*.
- Mottaghi, R.; Chen, X.; Liu, X.; Cho, N.-G.; Lee, S.-W.; Fidler, S.; Urtasun, R.; and Yuille, A. 2014. The Role of Context for Object Detection and Semantic Segmentation in the Wild. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Ramachandran, P.; Parmar, N.; Vaswani, A.; Bello, I.; Levskaya, A.; and Shlens, J. 2019. Stand-Alone Self-Attention in Vision Models. In *Conference on Neural Information Processing Systems*.
- Ranftl, R.; Bochkovskiy, A.; and Koltun, V. 2021. Vision Transformers for Dense Prediction. In *International Conference on Computer Vision*.
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. MobileNetV2: Inverted Residuals and Linear Bottlenecks. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Sixiao, Z.; Jiachen, L.; Hengshuang, Z.; Xiatian, Z.; Zekun, L.; Yabiao, W.; Yanwei, F.; Jianfeng, F.; Tao, X.; H.S., T. P.; and Li, Z. 2021. Rethinking Semantic Segmentation from a Sequence-to-Sequence Perspective with Transformers. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Tan, M.; and Le, Q. 2019. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In *International Conference on Machine Learning*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Łukasz Kaiser; and Polosukhin, I. 2017. Attention Is All You Need. In *Conference on Neural Information Processing Systems*.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers. *arXiv preprint arXiv:2105.15203*.
- Yang, M.; Yu, K.; Zhang, C.; Li, Z.; and Yang, K. 2018. DenseASPP for semantic segmentation in street scenes. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Yu, C.; Wang, J.; Gao, C.; Yu, G.; Shen, C.; and Sang, N. 2020. Context Prior for Scene Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Yuan, Y.; Chen, X.; and Wang, J. 2020. Object-Contextual Representations for Semantic Segmentation. In *European Conference on Computer Vision*.

Zhang, H.; Dana, K.; Shi, J.; Zhang, Z.; Wang, X.; Tyagi, A.; and Agrawal, A. 2018. Context Encoding for Semantic Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*.

Zhang, H.; Wu, C.; Zhang, Z.; Zhu, Y.; Zhang, Z.; Lin, H.; Sun, Y.; He, T.; Muller, J.; Manmatha, R.; Li, M.; and Smola, A. 2020. ResNeSt: Split-Attention Networks. *arXiv preprint arXiv:2004.08955*.

Zhang, H.; Zhan, H.; Wang, C.; and Xie, J. 2019. Semantic Correlation Promoted Shape-Variant Context for Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*.

Zhao, H.; Shi, J.; Qi, X.; Wang, X.; and Jia, J. 2017. Pyramid Scene Parsing Network. In *IEEE Conference on Computer Vision and Pattern Recognition*.

Zhu, Z.; Xu, M.; Bai, S.; Huang, T.; and Bai, X. 2019. Asymmetric Non-local Neural Networks for Semantic Segmentation. In *International Conference on Computer Vision*.