

Adaptive Context Encoding Module for Semantic Segmentation

Congcong Wang¹, Faouzi Alaya Cheikh¹, Azeddine Beghdadi², and Ole Jakob Elle^{3,4}

¹ Norwegian Colour and Visual Computing Lab, Norwegian University of Science & Technology, Norway. ² L2TI-Institut Galilée, Université Paris 13, Sorbonne Paris Cité, Villetaneuse, France. ³ The Intervention Centre, Oslo University Hospital, Oslo, Norway.

⁴ The Department of Informatics, University of Oslo, Oslo, Norway.

{congcong.wang, faouzi.cheikh}@ntnu.no, azeddine.beghdadi@univ-paris13.fr, oelle@ous-hf.no

Abstract

The object sizes in images are diverse, therefore, capturing multiple scale context information is essential for semantic segmentation. Existing context aggregation methods such as pyramid pooling module (PPM) and atrous spatial pyramid pooling (ASPP) employ different pooling size or atrous rate, such that multiple scale information is captured. However, the pooling sizes and atrous rates are chosen empirically. Rethinking of ASPP leads to our observation that learnable sampling locations of the convolution operation can endow the network learnable field-of-view, thus the ability of capturing object context information adaptively. Following this observation, in this paper, we propose an adaptive context encoding (ACE) module based on deformable convolution operation where sampling locations of the convolution operation are learnable. Our ACE module can be embedded into other Convolutional Neural Networks (CNNs) easily for context aggregation. The effectiveness of the proposed module is demonstrated on Pascal-Context and ADE20K datasets. Although our proposed ACE only consists of three deformable convolution blocks, it outperforms PPM and ASPP in terms of mean Intersection of Union (mIoU) on both datasets. All the experimental studies confirm that our proposed module is effective compared to the state-of-the-art methods.

Introduction

Semantic segmentation is a pixel wise classification problem, where class prediction is assigned for each pixel of an image. The development of deep learning brings semantic segmentation into a new era. Starting from the Fully Convolutional Network (FCN) [1], we have seen a rapid increase in the research field of semantic segmentation based on Convolutional Neural Networks (CNNs) [2–4]. Those methods boost the field and push the state-of-the-art performance on several semantic segmentation benchmarks.

However, multiple scale objects context understanding is still a challenging problem. Some approaches have been proposed to handle this problem. Following similar criteria presented in [5], we group those methods into three categories. First, image pyramid based methods: the input image is decomposed to image pyramid, then DCNN (Deep Convolutional Neural Network) is applied separately to every resolution level of the image pyramid input [6–8]. In this way, different scale objects are captured from different level feature maps. Second, encoder-decoder based methods: for encoder, convolution and pooling operations are applied hierarchically to extract features, then the spatial resolution is recovered in the decoder path by hierarchically applying up-sampling and convolution operations. The most representa-

tive architecture is U-Net [2] which has been achieving promising results in the medical image processing field. And many other encoder-decoder based architectures [4, 9–12]. Third, spatial pyramid pooling strategy based approaches where the feature maps are aggregated by pooling operations or by atrous convolutions with multiple rates. The atrous spatial pyramid pooling (ASPP) proposed in DeepLabs [3, 5] and the pyramid pooling module (PPM) presented in PSPNet [13] are two representative work of this group.

The performance of Deeplabs [5, 14] and PSPNet [13] on some benchmarks shows the effectiveness of their pyramid pooling module. However, the rates of ASPP and PPM are empirically selected, which is not an image content adaptive encoding. In this paper, our goal is to investigate if there is a way to aggregate the feature maps adaptively.

Rethinking of ASPP and atrous convolution, we observed that setting different values of the atrous rate of atrous convolution operation endows the network with multiple effective fields-of-view, thus the ability of capturing multi-scale context information. Therefore, if we can adaptively adjust the field-of-view of the convolution operations in the aggregation part, it would be possible for the network to aggregate contextual information adaptively based on input.

Interestingly, Deformable Convolution Networks (DCN) were proposed recently in [15, 16]. For deformable convolution, the sampling locations are learnable, which can be highly integrated into our purpose. Therefore, in this paper, we propose an adaptive context encoding (ACE) module based on deformable convolution. More precisely, we replace the ASPP module or PPM by three deformable convolution blocks as shown in Figure 2 (b).

This idea is evaluated on Pascal-Context [17] and ADE20K [18] datasets for semantic segmentation. We experimentally demonstrate that our proposed method improves the segmentation results consistently over the baseline methods: ASPP and PPM. Notably, more robust performance is shown under different batch size settings during the training process. Moreover, even though our goal is to find a better multi-scale aggregation module compared to ASPP and PPM, our method achieves the state-of-the-art on Pascal-Context dataset with 53.6% mIoU and promising results on ADE20K dataset with a final score of 0.5535. Furthermore, our ACE module can be easily embedded into other networks for further improvement.

The remainder of this paper is organized as follows. First, we review the related work on semantic segmentation as well as the highly relevant ASPP and PPM. Next, we describe our proposed method. Then, the experimental results and analysis are

presented. Finally, discussion and conclusions are drawn.

Related Work

Deep learning based semantic segmentation is rapidly developing and significant progress has already been achieved. DCNN with pooling and convolution operations is invariant to local image transformations, thus can extract abstractions of data hierarchically [3, 19]. On one hand, this ability is beneficial for high-level vision tasks such as classification. On the other hand, it can bound the performance of pixel wise dense prediction tasks where spatial information is important [20]. Semantic segmentation thus is challenging as it needs to perform classification and localization simultaneously.

There are many works proposed to improve semantic segmentation which can be briefly divided into two directions: Resolution Enlarging and Context Extraction.

Resolution Enlarging

Atrous convolution which is inspired by the atrous algorithm [21] is claimed to be useful for extracting denser feature maps and can further alleviate the detail information loss. Thus it is widely used in semantic segmentation to enlarge the receptive-field-of-view and extract denser feature maps [3, 14, 22]. Besides, encoder-decoder architectures employ decoder to up-sample the feature maps and increase its resolution hierarchically and partially recover some of the information loss of encoder [2, 9, 23]. They will be introduced in the following paragraphs.

Atrous convolution based methods: Typically, for DCNN, such as Resnet [24], the spatial resolution of the output feature maps of the final layer is 32 times smaller compared to the resolution of input images, which is harmful to pixel wise tasks. Atrous convolution is used to enlarge the receptive field while preserving the resolution of the feature map. DeepLabs especially DeepLabv2 [3] and DeepLabv3 [5] are a series of methods that investigate atrous convolution for semantic segmentation and are considered as one of the state-of-the-art techniques. Similar feature extraction DCNN backbone is also used in PSPNet [13], where the resolution of the final layer feature maps is 8 times smaller. Atrous convolution is an effective solution for spatial information loss. However, the larger feature maps and larger convolution kernels make the network computational demanding. Recently, in [22], Wu *et al.* proposed a Joint Pyramid Up-sampling (JPU) module to reduce the memory and time consuming atrous convolutions, while keeping the ability of extracting high resolution feature maps.

Encoder-Decoder based methods: In an encoder-decoder network, the spatial resolution is gradually up-sampled at the decoder part. DeconvNet [25] uses deconvolutional layers [25] to recover the resolution which can get full resolution final prediction by a complex decoder part. U-Net [2] introduces skip connections from encoder to decoder, thus the information from the skip connection is used to compensate for the information loss. RefineNet [9] elaborately design an up-sampling path to fuse low level and high level features. DeepLabv3+ [14] employs both skip connection and atrous convolution, thus reaching the state-of-the-art performance on some benchmarks to date.

Context Extraction

Scene context is important for extracting semantics. There are many approaches proposed to extract useful context information. Spatial pyramid pooling is proven to be effective for extracting context information [3, 13, 26]. Moreover, attention mechanism is proposed to learn the object context map in [23]. In [27], Peng *et al.* use convolution operations with large kernel size to extract classification information. Among those approaches, the spatial pyramid pooling based methods are popular. Spatial pyramid pooling aims at extracting multiple scale context information from feature maps. For semantic segmentation, pyramid pooling module (PPM) [13] investigates pooling operation as a tool for multiple scale context aggregation and atrous spatial pyramid pooling (ASPP) [5] exploits atrous convolution for pyramid pooling. These two modules will be described in detail as they are highly related to the proposed approach.

Pyramid Pooling Module: Global Average Pooling (GAP) is used to obtain global contextual prior in ParseNet [26] for semantic segmentation. However, as pointed out in [13], fusing one feature map into one single value may cause information loss. Thus, in [13], Zhao *et al.* propose to hierarchically apply pooling operations with four scales as illustrated in Figure 1 (a), resulting in feature maps with four levels of resolution. The coarsest level is obtained by applying GAP on the feature maps and gives a single vector output. For the other levels, the feature maps are first divided into sub-regions, then a global pooling is applied to every sub-region. The numbers of sub-region are set to 2×2 , 3×3 , 6×6 for each level respectively as illustrated in Figure 1 (a) [13]. This PPM can thus extract the information at different scales for context aggregation.

Atrous Spatial Pyramid Pooling: ASPP module is first proposed in [3] and further revised in [5]. In ASPP module, as shown in Figure 1 (b), different atrous rates are used to extract multiple scale information. Besides, to capture global context prior, similar to ParseNet [26] and PPM [13], GAP is applied. In summary, one 1×1 convolution block and three 3×3 atrous convolution blocks with different atrous rates (6, 12, 18 respectively), and one GAP block are employed in parallel.

While DeepLabs and PSPNet reach the state-of-the-art performance on different benchmarks when they are proposed, and they still have influence on semantic segmentation, it is important and interesting to investigate two following aspects: (1) It is observed that the numbers of sub-region of PPM in PSPNet and the atrous rates of ASPP module from DeepLabs are selected empirically. The choice of these parameters need to be adjusted according to the application, such as in [28], 2×2 , 3×3 , 5×5 , 6×6 sub-regions are chosen for their RMP (Residual Multi-kernel Pooling) module which is similar to PPM. It is essential to avoid choosing those parameters empirically. (2) PPM and ASPP both extract the context information by sampling from rigid rectangular regions which contain pixels from different object categories. However, for a certain pixel, the surrounding pixels which belong to the same category should contribute more. As pointed out in [29], Yuan and Wang hold similar opinion and they define object context as the set of pixels which belongs to the same category. While Yuan and Wang utilize self-attention [30] mechanism to exploit context from pixels that from the same object class, in this work, inspired by the original ASPP, we will investigate the possibility to aggregate multi-scale information by adjusting the field-of-

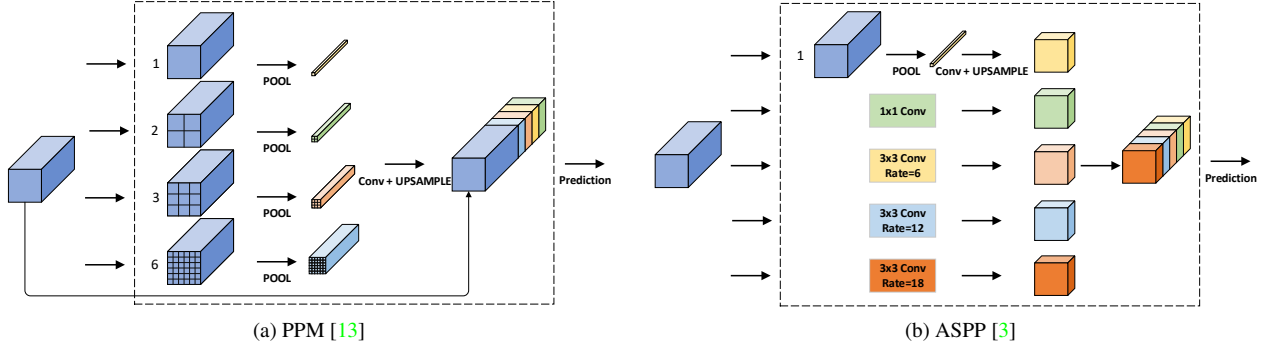


Figure 1: (a) Pyramid pooling module (PPM) proposed in PSPNet [13]. (b) Atrous spatial pyramid pooling (ASPP) module proposed in DeepLabs [3,5].

view of the convolution operation adaptively.

Method

In this section, we will discuss our proposed ACE module in detail. The most relevant atrous and deformable convolution operations are first introduced, then the ACE module and the network architecture are demonstrated.

Convolution Operation

Atrous convolution is chosen as the tool for the context aggregation module in ASPP [5]. For two dimensional signals such as images, atrous convolution can be written as:

$$y[i] = \sum_{k \in K} x[i + r \cdot k] \cdot w[k], \quad (1)$$

where y indicates the output after atrous convolution operation, i is the location, x is the input signal, r is the atrous rate. K defines the convolutional kernel, for example, for a 3×3 standard convolutional kernel, $K = \{(-1, -1), (-1, 0), \dots, (0, 1), (1, 1)\}$. $w[k]$ denotes the weight of the k -th location, and k enumerates the locations in K . When $r = 1$, the equation stands for standard convolution. The value of r controls the sampling locations of atrous convolution. In ASPP, different sizes of field-of-view are obtained by setting different r values. This observation leads to our claim that a learnable field-of-view can thus be obtained by learnable sampling locations of the convolution operation.

It is interesting to notice that the recent proposed deformable convolution [15, 16] meets our requirement where convolution with learnable sampling locations is proposed. Dai *et al.* propose Deformable Convolutional Networks (DCNv1) in [15] and Zhu *et al.* propose a revised version (DCNv2) in [16]. Eq. 2 presents the deformable convolution operation from DCNv1:

$$y[i] = \sum_{k \in K} x[i + k + \Delta k] \cdot w[k], \quad (2)$$

where Δk denotes offsets. The regular sampling locations k is then augmented with the irregular offset Δk .

It is clearly observed that the main difference between Eq. 1 and Eq. 2 is that the sampling locations of atrous convolution are always regular grid. For example, the sampling grid is square for a 3×3 kernel, no matter what the value of atrous rate r is. But the offset Δk is input dependent, without regular shape constraint. Besides, compared to the manually set atrous rate r , the offset

Δk is learned by the network. In [15], Zhu *et al.* further investigate deformable convolution and find out that the spatial support of the deformable convolution operation from DCNv1 can extend beyond the pertinent region. Therefore, they propose DCNv2 to let the network better focus on relevant image content by introducing a modulation mechanism to manipulate the spatial support region. This modulated deformable convolution can be expressed as follows:

$$y[i] = \sum_{k \in K} x[i + k + \Delta k] \cdot w[k] \cdot \Delta m_k, \quad (3)$$

where Δm_k is the learnable modulation value with a range of [0,1]. This modulation value can further adjust the sampled pixel's contribution, thus the spatial support regions are adjusted better.

Therefore, in this work, we propose to employ the deformable convolution operation from DCNv2 as the tool for context aggregation.

Network Architecture

A brief illustration of context extraction based semantic segmentation pipeline is depicted in Figure 2 (a). For a given input image, convolution networks such as ResNet [24], Xception [31], are applied to extract feature maps, then feature aggregation is employed to extract context information. Based on the aggregated feature information, the final prediction is made and up-sampled to the original input spatial resolution.

In this paper, we only focus on feature aggregation part. Our proposed ACE module is shown in Figure 2 (b). After feature extraction, the input image is represented by feature maps with the size of $H \times W \times C$ where H and W are the height and width of the feature maps respectively, and C indicates the number of feature channels. In ACE module, three deformable convolution blocks (DCB) are applied to aggregate the feature maps. Each block is consists of ‘‘Deformable Convolution (DConv) \rightarrow BN (Batch-Norm) \rightarrow ReLU (Rectified Linear Unit)’’ operations. The outputs size after each deformable convolution block are $H \times W \times \frac{C}{4}$, $H \times W \times \frac{C}{8}$, and $H \times W \times \frac{C}{8}$ separately. After the ACE module, one 1×1 convolution operation is applied for the final segmentation map prediction. Then the predicted result is up-sampled by ‘bilinear’ up-sampling operation to the original image spatial resolution directly.

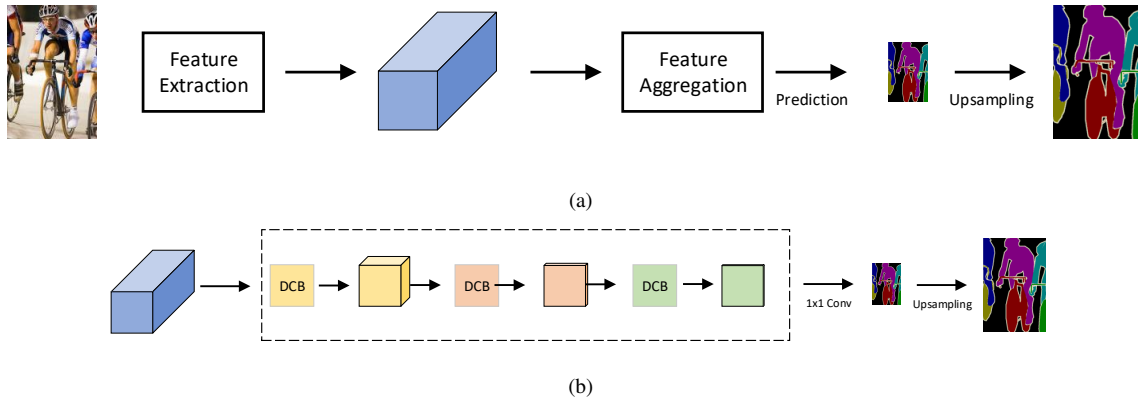


Figure 2: (a) A brief illustration of context extraction based semantic segmentation pipeline. (b) Proposed adaptive context aggregation (ACE) module.

Experiment

In this section, we validate our proposed module on two public datasets Pascal-Context [17] and ADE20K [18]. We first introduce the implementation details. Then the experiment results are presented and analyzed on these two datasets. The performance of the proposed method is evaluated in terms of two common measures, namely pixel accuracy (pixAcc) and mean Intersection of Union (mIoU).

In order to illustrate the effectiveness of the proposed method, we will compare it with ASPP and PPM. It is worth noticing that DeepLabv3 and PSPNet utilize atrous convolution for feature extraction which is memory and time consuming. In [22], Wu *et al.* propose a Joint Pyramid Up-sampling (JPU) module to replace the heavy feature extraction module. Their method (FastFCN) reduces more than three times the computation complexity and reaches slightly better performance. As a result of the limited computational resources we have, we adopt FastFCN’s feature extraction part as the backbone for comparison. In other words, only the feature aggregation (*head*) part is replaced with ASPP, PPM and our proposed ACE module.

Implementation Details

The implementation is based on the PyTorch [32] implementation of FastFCN [22]¹ which is similar to [33]² and the implementation of Deformable ConvNets^{3,4}. For a fair comparison, we adopt the original training strategy of FastFCN [22]. Specifically, “poly” learning rate policy is used: $lr = baselr * (1 - \frac{iter}{total_iter})^{power}$, where $power = 0.9$. The initial base learning rate is set to 0.001 for batch size 16 for PASCAL-Context [17] and 0.01 for ADE20K [18]. Besides, it is adjusted relatively to batch size value if other batch size is chosen: $baselr_adjusted = \frac{baselr}{16} * batch_size$. The networks are trained for 80 epochs with SGD for PASCAL-Context [17] and 120 epochs for ADE20K [18]. The momentum is 0.9 and the weight decay is set to 0.0001. For data augmentation, the image is randomly flipped and scaled between 0.5 to 2. Then the image is cropped to a fixed size (480

$\times 480$). Pixel-wise cross-entropy loss and auxiliary loss as presented in [13, 22] are used, the weight for auxiliary loss is set to 0.2.

Due to limited access to multi-GPUs computational resource, our experiment includes training on a single GeForce RTX 2080 GPU for small batch sizes and training on $4 \times$ GeForce GTX 1080 GPUs for batch size 16.

Pascal-Context

Dataset: Pascal-Context dataset [17] is based on PASCAL VOC 2010 with additional annotations that provides annotations for the whole scene. Training images are 4,998 (*pascal-train*) and testing images contain 5,105 images (*pascal-val*). Following the prior work [3, 22, 33], the semantic labels we used in this paper are the 59 categories with one background class.

Experimental Results: Table 1 illustrates the performance on *pascal-val* of ASPP, PPM and the proposed ACE based FastFCN [22]. We trained the models on different batch sizes: 4, 6 and 16. 6 is the maximum batch size for our single GPU. The methods employ ResNet-50 [24] as backbone. The reported results are obtained for 59 classes without multi-scale evaluation. All the methods are trained by our machine for fair comparison. Obviously, our proposed ACE outperforms ASPP and PSP on all the different batch sizes training settings. It’s worth mentioning that, ASPP based FastFCN’s accuracy influenced severely by batch size, as also pointed out in the paper of Deeplabv3 [5]: training DeepLabv3 model with small batch size is inefficient. Note that our proposed method not only reaches the best result, but also shows its stable performance for different batch sizes. The absolute improvements of mIoU of ACE compared to ASPP are 4.45%, 2.83% and 1.31% for batch sizes 4, 6, 16 respectively. And 2.39%, 1.04%, and 0.81% compared to PPM.

Table 2 illustrates the performance compared to the state-of-the-art methods. For a fair comparison, the reported result of our method is calculated with background class and multi-scale evaluation where the network prediction is averaged through multiple scales as in [13, 22, 26, 33]. The results of the other methods are obtained from the corresponding papers. Our proposed method achieves 53.6% mIoU, which outperforms the previous methods.

¹<https://github.com/wuhuikai/FastFCN>

²<https://github.com/zhanghang1989/PyTorch-Encoding>

³<https://github.com/chengdazhi/Deformable-Convolution-V2-PyTorch>

⁴<https://github.com/open-mmlab/mmdetection>

Batch Size	Head	pixAcc%	mIoU%
4	ASPP	75.42	43.62
	PPM	75.58	45.68
	Proposed	77.68	48.07
6	ASPP	77.19	46.53
	PPM	77.45	48.32
	Proposed	78.35	49.36
16	ASPP	78.68	49.04
	PPM	78.41	49.54
	Proposed	78.85	50.35

Table 1: Segmentation results on PASCAL-Context dataset (*pascal-val*).

Method	mIoU%
FCN-8s [1]	37.8
ParseNet [26]	40.4
Piecewise [8]	43.3
DeepLabv2 (Res101-COCO) [3]	45.7
RefineNet (Res152) [9]	47.3
PSPNet (Res101) [13]	47.8
EncNet (Res101) [33]	51.7
DANet (Res101) [34]	52.6
FastFCN (Res101,EncNet)* [22]	53.1
Proposed (Res101)	53.6

* FastFCN backbone with EncNet head.

Table 2: Segmentation results on PASCAL-Context dataset (*pascal-val*) of the state-of-the-art methods.

ADE20K

Dataset: ADE20K is used in ImageNet Scene parsing challenge 2016 and it contains 150 object categories. It is divided into 20k/2K/3K images for training (*a-train*), validation (*a-val*) and testing (*a-test*) respectively.

Experimental Results: Table 3 demonstrates the results for *a-val* set without multi-scale evaluation. Note, one GeForce RTX 2080 GPU maximum can fit 4 batches, thus the results reported in Table 3 are obtained on batch size 4 and ResNet-50 as backbone. Our proposed method achieves the best result compared to ASPP and PPM based FastFCN with an absolute improvement of 1.4% and 0.71% for ASPP and PPM in terms of mIoU.

Batch Size	Head	pixAcc%	mIoU%
4	ASPP	78.11	37.11
	PPM	77.39	37.80
	Proposed	78.62	38.51

Table 3: Segmentation results on ADE20K dataset (*a-val*).

In order to compare with the state-of-the-art methods, we further train our model with ResNet-101 backbone on 4×GeForce 1080 GPUs with batch size 16. Table 4 shows the obtained result and results reported in the corresponding papers of the other approaches. The proposed method provides better results compared to PSPNet with an absolute improvement of 0.52% of mIoU. EncNet achieves the best result. Except for the methodology itself, some part of the performance gap could be from the training strategy, such as EncNet is trained with an image size of 576×576 and our method is trained with 480×480.

Method	pixAcc%	mIoU%
FCN [1]	71.32	29.39
SegNet [4]	71.00	21.64
DilatedNet [35]	73.55	32.31
CascadeNet [18]	74.52	34.90
RefineNet (Res152) [9]	-	40.7
PSPNet (Res101) [13]	81.39	43.29
EncNet (Res101) [33]	81.69	44.65
FastFCN (Res101,EncNet)* [22]	80.99	44.34
Proposed	81.07	43.81

* FastFCN backbone with EncNet head.

Table 4: Segmentation results on ADE20K dataset (*a-val*) of the state-of-the-art methods.

Moreover, we fine-tune our trained model for another 20 epochs on *a-train a-val* set with a smaller learning rate 0.001, then submit the *a-test* set result to the evaluation website⁵. Our method obtains 72.99% (pixAcc) and 37.71% (mIoU) with a final score of 0.5535 which is not the best but is an encouraging result.

Discussion and Conclusion

In summary, in this work, we revisited the atrous convolution operation and pyramid pooling modules and propose an effective feature aggregation method based on deformable convolution to extract multiple scale context adaptively for the final segmentation map prediction. Based on the experimental validation, our method outperforms the ASPP module and PPM on Pascal-Context and ADE20K datasets. Noticeably, although the goal for this work is to propose a better multiple scale context aggregation module, rather than to obtain the best results on the benchmarks, our proposed approach achieves state-of-the-art result 53.6% mIoU on Pascal-Context and encouraging result 0.5535 on ADE20K.

All the experiments confirm that an adaptive context encoding (ACE) module is benefit for semantic segmentation which deserves further research. In this work, we directly use deformable convolution as the tool for ACE and simply cascaded three deformable convolution blocks, sophisticated design of the architecture is essential. We believe that further exploration of the usage and improvement of our feature aggregation idea is promising and necessary in the design of an efficient semantic segmentation approach.

Acknowledgment

The authors would like to acknowledge Junfeng Ge for the support of computational resources. This work is funded by the Research Council of Norway through project no. 247689 IQ-MED: Image Quality enhancement in MEDical diagnosis, monitoring and treatment.

References

- [1] J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2015, pp. 3431–3440.
- [2] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional net-

⁵<http://sceneparsing.csail.mit.edu/>

- works for biomedical image segmentation,” in *International Conference on Medical Image Computing and Computer-assisted Intervention*. Springer, 2015, pp. 234–241.
- [3] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, “Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 4, pp. 834–848, 2018.
- [4] V. Badrinarayanan, A. Kendall, and R. Cipolla, “Segnet: A deep convolutional encoder-decoder architecture for image segmentation,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 12, pp. 2481–2495, 2017.
- [5] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, “Rethinking atrous convolution for semantic image segmentation,” *arXiv preprint arXiv:1706.05587*, 2017.
- [6] C. Farabet, C. Couprie, L. Najman, and Y. LeCun, “Learning hierarchical features for scene labeling,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 8, pp. 1915–1929, 2013.
- [7] P. H. Pinheiro and R. Collobert, “Recurrent convolutional neural networks for scene labeling,” in *International Conference on Machine Learning*, 2014.
- [8] G. Lin, C. Shen, A. Van Den Hengel, and I. Reid, “Efficient piecewise training of deep structured models for semantic segmentation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2016, pp. 3194–3203.
- [9] G. Lin, A. Milan, C. Shen, and I. Reid, “Refinenet: Multi-path refinement networks for high-resolution semantic segmentation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2017, pp. 1925–1934.
- [10] M. Amirul Islam, M. Rochan, N. D. Bruce, and Y. Wang, “Gated feedback refinement network for dense image labeling,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2017, pp. 3751–3759.
- [11] A. Mohammed, S. Yildirim, I. Farup, M. Pedersen, and Ø. Hovde, “Y-net: A deep convolutional neural network for polyp detection,” *arXiv preprint arXiv:1806.01907*, 2018.
- [12] A. Mohammed, S. Yildirim, I. Farup, M. Pedersen, and Ø. Hovde, “Streoscennet: surgical stereo robotic scene segmentation,” in *Medical Imaging 2019: Image-Guided Procedures, Robotic Interventions, and Modeling*, vol. 10951. International Society for Optics and Photonics, 2019, p. 109510P.
- [13] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, “Pyramid scene parsing network,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2017, pp. 2881–2890.
- [14] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, “Encoder-decoder with atrous separable convolution for semantic image segmentation,” in *Proceedings of the European Conference on Computer Vision*. Springer, 2018, pp. 801–818.
- [15] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei, “Deformable convolutional networks,” in *Proceedings of the IEEE International Conference on Computer Vision*. IEEE, 2017, pp. 764–773.
- [16] X. Zhu, H. Hu, S. Lin, and J. Dai, “Deformable convnets v2: More deformable, better results,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9308–9316.
- [17] A. Oliva and A. Torralba, “The role of context in object recognition,” *Trends in cognitive sciences*, vol. 11, no. 12, pp. 520–527, 2007.
- [18] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, “Scene parsing through ade20k dataset,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2017, pp. 633–641.
- [19] M. D. Zeiler and R. Fergus, “Visualizing and understanding convolutional networks,” in *Proceedings of the European Conference on Computer Vision*. Springer, 2014, pp. 818–833.
- [20] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, “Semantic image segmentation with deep convolutional nets and fully connected crfs,” in *International Conference on Learning Representations*, 2015.
- [21] S. Mallat, *A wavelet tour of signal processing*. Elsevier, 1999.
- [22] H. Wu, J. Zhang, K. Huang, K. Liang, and Y. Yu, “Fastfcn: Rethinking dilated convolution in the backbone for semantic segmentation,” *arXiv preprint arXiv:1903.11816*, 2019.
- [23] C. Yu, J. Wang, C. Peng, C. Gao, G. Yu, and N. Sang, “Learning a discriminative feature network for semantic segmentation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2018, pp. 1857–1866.
- [24] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2016, pp. 770–778.
- [25] H. Noh, S. Hong, and B. Han, “Learning deconvolution network for semantic segmentation,” in *Proceedings of the IEEE International Conference on Computer Vision*. IEEE, 2015, pp. 1520–1528.
- [26] W. Liu, A. Rabinovich, and A. C. Berg, “ParseNet: Looking wider to see better,” *arXiv preprint arXiv:1506.04579*, 2015.
- [27] C. Peng, X. Zhang, G. Yu, G. Luo, and J. Sun, “Large kernel matters—improve semantic segmentation by global convolutional network,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2017, pp. 1743–1751.
- [28] Z. Gu, J. Cheng, H. Fu, K. Zhou, H. Hao, Y. Zhao, T. Zhang, S. Gao, and J. Liu, “Ce-net: Context encoder network for 2d medical image segmentation,” *IEEE transactions on medical imaging*, 2019.
- [29] Y. Yuan and J. Wang, “Ocnets: Object context network for scene parsing,” *arXiv preprint arXiv:1809.00916*, 2018.
- [30] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [31] F. Chollet, “Xception: Deep learning with depthwise separable convolutions,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2017, pp. 1251–1258.
- [32] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, “Automatic differentiation in pytorch,” 2017.
- [33] H. Zhang, K. Dana, J. Shi, Z. Zhang, X. Wang, A. Tyagi, and A. Agrawal, “Context encoding for semantic segmentation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2018, pp. 7151–7160.
- [34] J. Fu, J. Liu, H. Tian, Y. Li, Y. Bao, Z. Fang, and H. Lu, “Dual attention network for scene segmentation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2019, pp. 3146–3154.
- [35] F. Yu and V. Koltun, “Multi-scale context aggregation by dilated convolutions,” in *International Conference on Learning Representations*, 2016.

JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!



- **SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •**
- **INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •**

www.electronicimaging.org

