

Boundary Guided Context Aggregation for Semantic Segmentation

Haoxiang Ma^{1,2}
mahaoxiang822@buaa.edu.cn

Hongyu Yang^{†3}
hongyuyang@buaa.edu.cn

Di Huang^{1,2}
dhuang@buaa.edu.cn

State Key Laboratory of Software
Development Environment¹
Beihang University, Beijing, China
School of Computer Science and
Engineering²
Beihang University, Beijing, China
Institute of Artificial Intelligence³
Beihang University, Beijing, China

Abstract

The recent studies on semantic segmentation are starting to notice the significance of the boundary information, where most approaches see boundaries as the supplement of semantic details. However, simply combing boundaries and the mainstream features cannot ensure a holistic improvement of semantics modeling. In contrast to the previous studies, we exploit boundary as a significant guidance for context aggregation to promote the overall semantic understanding of an image. To this end, we propose a Boundary guided Context Aggregation Network (BCANet), where a Multi-Scale Boundary extractor (MSB) borrowing the backbone features at multiple scales is specifically designed for accurate boundary detection. Based on which, a Boundary guided Context Aggregation module (BCA) improved from Non-local network is further proposed to capture long-range dependencies between the pixels in the boundary regions and the ones inside the objects. By aggregating the context information along the boundaries, the inner pixels of the same category achieve mutual gains and therefore the intra-class consistency is enhanced. We conduct extensive experiments on the Cityscapes and ADE20K databases, and comparable results are achieved with the state-of-the-art methods, clearly demonstrating the effectiveness of the proposed one. Our code is available at <https://github.com/mahaoxiang822/Boundary-Guided-Context-Aggregation>.

1 Introduction

Semantic segmentation is a fundamental and long-standing task in computer vision, which aims to get a dense prediction of the object category for every pixel in an image. It has been extensively and actively applied to many challenging applications, *e.g.*, autonomous driving [7], image editing [1], and human-machine interaction [15].

Recently, a number of approaches based on fully convolutional networks (FCN) [14] have been proposed for semantic segmentation. Due to the fixed geometric structure and the limited receptive field, however, these methods expose a common drawback in contextual



Figure 1: Comparison of various manners for context aggregation. (a) The original input image; (b) non-adaptive context aggregation; (c) attention-based context aggregation; and (d) the proposed boundary guided context aggregation.

information modeling. To make up for the above deficiency, a number of studies explore contextual dependencies for improved results and the existing methods can be briefly summarized into two categories. One is to use pyramid-based module which combines different scales of atrous convolution or pooling layers to enlarge the receptive field, such as Atrous Spatial Pyramid Pooling (ASPP) in Deeplab [5] and Pyramid Pooling Module (PPM) in PSP-Net [29]. Another trend is to adaptively model long-range dependencies from the aspect of channel or spatial. For example, EncNet [26] learns to weight the feature map on channels and Non-local module [19] exploits the attention-based mechanism to enable one pixel to perceive all the other positions of an image. These alternatives achieve promising results on semantic segmentation. However, either aggregating contextual information in a fixed manner or with a learned response map, the existing context aggregation modules generally lack an explicit prior on the regions to aggregate that undesired inter-class dependency and intra-class ambiguity are also involved, as illustrated in Figure 1(b) and (c). The activated context regions are not properly regularized in the previous methods and this will inevitably impose adverse effect on the contexts.

Considering that object boundary regions contain less mixed dependencies and they are important clues that all the inner pixels share in common, we therefore investigate a boundary constraint for context aggregation in this study, that only the dependencies between the inner pixels and boundaries are built. We notice that boundaries have been actively exploited for improved segmentation performance. In particular, the traditional post-processing modules regard boundary as a useful constraint in domain transform or global energy function, and the recent Boundary-aware Feature Propagation (BFP) module extends this idea into deep models. However, these operations are usually computationally complex and they are not easy to be integrated into an end-to-end model [2, 3, 8]. Moreover, the methods above heavily rely on the closure of the detected boundary, otherwise the context information will be leaked into the neighboring areas and downgrades the inter-class distinction. Such fact indicates that there still leaves much space for improvement.

In contrast to the existing methods, this study directly wires boundary as the context to be aggregated. See Figure 1(d) for an illustration, even though the local surrounded context of the two marked pixels are quite different, their features are more likely to be close to each other once their relationships with the boundary are well established. By aggregating the semantic information along the corresponding boundary region into the inner pixels, any positions in this category harvest similar context, regardless of their conspicuous texture difference. Compared with the previous context solutions that directly model dependencies within a peripheral region or the full-image, the boundary constraint used in this study can not only exclude the undesired inter-class context relationships explicitly, but also alleviate intra-class ambiguity, e.g. an advertisement posted on a bus.

To this end, we propose a novel framework, named Boundary guided Context Aggregation network (BCANet), for image semantic segmentation. It casts boundary detection as an

independent sub-task along with the mainstream feature learning, and develops an attention based mechanism for boundary guided context aggregation. Specifically, the Multi-Scale Boundary extractor (MSB) is introduced to predict the binary boundary of an image, which borrows semantic features from multiple stages of the backbone as its input and feeds back the semantic embedded boundary features to the mainstream for context aggregation. The Boundary guided Context Aggregation module (BCA) modified from the Non-local network calculates the attention map between the boundary features and the mainstream semantics, so that boundaries can be seen as a mutual guidance for aggregating the context information, which enables the pixels of the same class achieve similar gains. Extensive evaluations validate that the proposed method performs favorably against the current state-of-the-art context aggregation approaches.

2 Related Work

2.1 Semantic Segmentation

In the recent years, the Convolutional Neural Networks (CNNs) based methods are dominating the field of semantic segmentation. Fully Convolutional Network (FCN) [14] opens a precedent for the application of CNNs in semantic segmentation by replacing the fully connected layers with convolution layers. More recently, a number of studies focus on capturing richer context information to augment the feature representation, where the multi-scale pyramid modules promote the performance a lot. For example, Zhao *et al.* [29] investigate PPM in PSPNet to aggregate contextual information at different scales, and Chen *et al.* [5] introduce the ASPP module. Besides, to overcome the limitation of fixed size of the kernel, several global pooling based methods [22, 26] and graph convolutional based methods [27, 28] are introduced to perceive the global context. To overcome the limitation of labeled training data, Chen *et al.* [6] develop a semi-supervised method and they achieve the state-of-the-art results. Our study also leverages the context information, but we see boundary as a more effective and explicit constraint.

2.2 Boundary in Semantic Segmentation

As an essential element of image, boundary has been paid much attention in computer vision. In the early research on FCN-based semantic segmentation, [2, 3, 5] use boundaries for post-processing to refine the result at the end of the network. Recently, several methods are starting to explicitly model boundary detection as an independent subtask in parallel with semantic segmentation for sharper results. Takikawa *et al.* [17] and Zhen *et al.* [31] specially design a boundary stream and couple the two tasks of boundary and semantics modeling. Li *et al.* [12] point out that the object boundary and body parts correspond to the high frequency and low frequency information of an image, respectively, based on which they decouple the body and edge with diverse supervisions. Yuan *et al.* [25] propose a model-agnostic method for boundary refinement. Different from the existing studies, our work explicitly exploits the boundary information for context aggregation which will further enhance the semantic representations, rather than just simply combining them. Ding *et al.* [8] investigate a Boundary-aware Feature Propagation module (BFP) to propagate information among the inner pixels of an object, which shares a common concept with our work. However, the full feature map will become over-smooth in BFP once the detected boundary is not closed, thus it greatly downgrades the inter-class distinction.

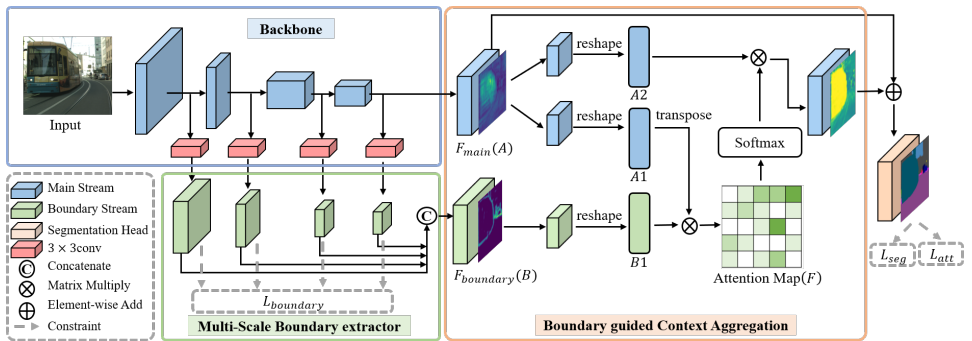


Figure 2: An overview of the proposed Boundary guided Context Aggregation Network (BCANet). With the semantic embedded boundary features generated by the Multi-Scale Boundary module (MSB), the Boundary guided Context Aggregation module (BCA) employs a cross-stream attention for more accurate context aggregation.

2.3 Attention Mechanism in Semantic Segmentation

Attention mechanism has been actively exploited in deep neural networks. For semantic segmentation, Chen *et al.* [4] present a pioneering work that utilizes attention to reweight the multi-scale features. Inspired by the use of self-attention in machine translation [18], Wang *et al.* [19] propose the Non-local module to capture global dependencies by calculating the correlation matrix of pixels at all positions and apply it to generate new feature representations. Based on Non-local module, a number of methods [9, 11, 20, 23, 24] are proposed for more accurate semantic segmentation. Our BCA module is also inspired by attention mechanism, we present a different but more effective network, where it mainly focuses on the positions in boundary regions. This solution is more powerful in reinforcing the intra-class correlation than the previous methods when intra-class ambiguity occurs.

3 Methodology

The proposed method explicitly models boundaries as the context, ensuring that inner pixels within the same object achieve similar gains after context aggregation. See Figure 2 for an overview. Specifically, the method employs the residual networks [10] as its backbone, where the vanilla convolutions are replaced by the dilated ones to enlarge the receptive field. Furthermore, we introduce the Multi-Scale Boundary extractor (MSB) to model boundary prediction as an independent sub-task along with the mainstream feature learning. The features from each stage of the backbone with different scales are concatenated together to capture boundary information with diverse semantics. With the semantic embedded boundary features generated by MSB, the Boundary guided Context Aggregation module (BCA) is carefully designed to perform semantic context aggregation. We build relationships between each pixel in the semantic feature maps and the boundary features to generate an attention map where the object boundary regions of the same category are highly activated. By aggregating the context information along the boundaries, pixels of the same category can achieve similar gains, thus enhancing the semantic consistency. We detail the method subsequently.

3.1 Multi-Scale Boundary (MSB) extractor

Considering that the boundary features are used to guide the process of semantic context aggregation and they interact with the mainstream features, MSB thus directly borrows the

intermediate representations from the backbone as its input. We add multiple connections between the mainstream and the boundary streams to enable different types of information to flow across the network, so that both low-level details and high-level semantics can be utilized. Specifically, the feature maps from the last residual block of each stage of the backbone are fed into MSB and their channels are unified to 256 by a 3×3 convolutional layer. Since the boundary features interact with the mainstream, thus they convey semantic-related information. To make the boundary regions salient, the features of multi-scales are directly supervised by the binary boundary labels generated from the segmentation ground-truth. Here, a 1×1 convolutional layer and a Sigmoid function is used to map the boundary features to edge maps, and all scales of boundary features are resized to $1/8$ size of the input image and they are concatenated together for the following context aggregation.

3.2 Boundary guided Context Aggregation (BCA) module

We exploit the intrinsic partitioning capability of boundaries and use them as the harvested contexts to reinforce intra-class consistency. Inspired by the success of Non-local module in modeling long-range dependency in semantic segmentation, an attention-based module is developed, *i.e.*, Boundary guided Context Aggregation module (BCA), to aggregate the context semantics. Unlike the original Non-local module that calculates self-attention, the boundary features serve as the *Key* in the proposed BCA modules, where multi-scale features with rich semantics are used and only the boundaries regions are salient. As such, pixels from the same object activate similar attention areas, while that from different objects share less similarities. Figure 2 shows the structure of the proposed BCA module. Specifically, given a semantic feature map $A \in \mathbb{R}^{C_1 \times H \times W}$ generated from the backbone and a boundary feature map $B \in \mathbb{R}^{C_2 \times H \times W}$ from the MSB, they are processed by two convolutional layers to generate two new feature maps $\{A_1, B_1\} \in \mathbb{R}^{C \times H \times W}$, where $C = 256$. The features are then reshaped to $\mathbb{R}^{C \times N}$, where $N = H \times W$ is the number of pixels. We conduct matrix multiplication between the transpose of reshaped A_1 and B_1 and then apply a Softmax function. The whole operation can be described as:

$$F(i, j) = \frac{\exp(B_{1i} \cdot A_{1j}^T)}{\sum_{i=1}^N \exp(B_{1i} \cdot A_{1j}^T)} \quad (1)$$

where F is the boundary-semantic similarity map and $F(i, j)$ indicates the effect of i^{th} position in the boundary feature map B on the j^{th} position in the semantic feature map A .

The semantic context along the boundaries in A will be aggregated into inner pixels. We therefore process A with two convolutional layers of kernel size 1×1 and reshape the output A_2 to $\mathbb{R}^{C \times N}$, then perform matrix multiplication between A_2 and the boundary-semantic similarity map F , followed by an element-wise sum operation with feature A . For the j^{th} pixel, its new feature can be calculated by:

$$D_j = A_j + \sum_{i=1}^N F(i, j) \cdot A_{2i} \quad (2)$$

where D is the enhanced feature map after boundary context aggregation. Since the positions in the boundary regions of the same category will be activated with much higher weights than the other irrelevant ones, thus Equation (2) can be approximately written as:

$$D_j = A_j + \sum_{i \in \text{Boundary}} F(i, j) \cdot A_{2i} \quad (3)$$

which indicates that intra-class pixels achieve similar updates. Compared with the previous Non-local module based methods that directly model the dependency between pixels, boundary information can better guarantee the intra-class consistency when ambiguity occurs.

3.3 Training Loss

The proposed method consists of two main components which aim to generate the segmentation masks and boundaries, respectively. For semantic segmentation, the standard cross-entropy L_{seg} is exploited to assign every pixel a category label. For boundary prediction, the binary cross-entropy loss $L_{boundary}$ is applied on MSB module, written as:

$$L_{boundary} = - \sum_i (b_i \log \hat{b}_i + (1 - b_i) \log (1 - \hat{b}_i)) \quad (4)$$

where b_i and \hat{b}_i are the ground-truth and predicted boundaries, respectively. To enhance the consistency of the two collaborative tasks, we also utilize another boundary-aware loss, *i.e.* L_{att} [12, 17], to describe the segmentation accuracy of the pixels along the boundary regions:

$$L_{att} = - \sum_i \sum_c \mathbb{1}[\hat{b}_i > t] \cdot (s_{i,c} \log \hat{s}_{i,c}) \quad (5)$$

where t is the threshold of positive boundary, $s_{i,c}$ and $\hat{s}_{i,c}$ are the ground-truth and segmentation result of the i -th pixel of class c , respectively. Inspired by [29], we further apply an auxiliary cross-entropy loss L_{aux} to the intermediate feature representations of the backbone to accelerate model convergence. Finally, the system training loss can be written as:

$$L = \lambda_1 L_{seg} + \lambda_2 L_{boundary} + \lambda_3 L_{att} + \lambda_4 L_{aux} \quad (6)$$

where $\lambda_1, \lambda_2, \lambda_3$ and λ_4 are the hyper-parameters, we empirically set $\lambda_1 = 1, \lambda_2 = 20, \lambda_3 = 1$ and $\lambda_4 = 0.4$.

4 Experiments

We conduct extensive experiments on the benchmark datasets of Cityscapes [7] and ADE20K [32], and make fair comparison with the state-of-the-art counterparts.

4.1 Implementation Details

The model is implemented in Pytorch. ResNet-101 is used as the backbone and the dilation strategy is applied as in [5]. After the BCA module, the network is tailed with two 3×3 convolutional layers and one 1×1 convolutional layer to generate the segmentation mask. The resolution of the features to generate the final segmentation result is 1/8 of the initial input and bilinear interpolation is used to upsample the prediction. Data augmentation is performed during training as in [29]. The initial learning rate is set to 0.01, and momentum is set to 0.9 with weight decay of 0.0005. We use the polynomial learning rate adjustment policy $\gamma = \gamma_0 \times (1 - \frac{iter}{max_{iter}})^{power}$, where γ and γ_0 are the current and initial learning rate, respectively, and power is set to 0.9. The batch size is set to 8 for Cityscapes and 12 for ADE20K, and the models are trained for 200 and 120 epochs on these two datasets, respectively. During inference, we follow the multi-scale inference strategy as in [8, 26, 29]. The open source code provided by [17] is utilized to generate the ground-truth boundaries from the segmentation masks and the boundary radius is set to 2.

4.2 Model Analysis

Ablation study. As shown in Table 1, simply incorporating the MSB module into the network has little influence on the performance of the dilated FCN baseline model. When introducing the boundary guided context aggregation mechanism but with a single-scale boundary (SSB) extractor, the performance has significantly improved by 3.54%, demonstrating that the proposed information aggregation manner can indeed aggregate more reasonable context. By utilizing the proposed MSB module to perform boundary detection, the segmentation performance continues to improve and 80.03% mIoU has been achieved, which clearly validates the significance of the multi-scale design of the proposed MSB module. Besides, the boundary-aware loss leads to an additional mIoU improvement of 0.89%, clearly validating its effectiveness.

MSB	SSB	BCA	L_{att}	mIoU%
				75.52
✓				74.95
	✓	✓		79.06
✓		✓		80.03
✓		✓	✓	80.92

Table 1: Ablation study on MSB and BCA modules and the boundary-aware training loss.

Comparison of different context aggregation modules. The results are shown in Table 2. Specifically, ASPP [5] applies parallel convolutions to capture multi-scale context information, its potential in guiding the mainstream context aggregation is not fully exploited. Our model achieves 0.74% accuracy improvement than ASPP. Compared to the previous self-attention based methods, such as Non Local [19], RCCA [11] and DNL [20], BCA can effectively alleviate the adverse effect caused by intra-class ambiguity, leading to significant mIoU improvements of 1.43%, 1.59% and 0.65%, respectively. Moreover, different from BFP [8] that exploits boundaries as a constraint for feature propagation, the proposed BCA is not sensitive to the closure of boundaries and it achieves 4.15% mIoU improvement.

Metric	Dilated FCN	+ASPP[5]	+Non Local[19]	+RCCA[11]	+DNL[20]	+MSB-BFP[8]	+MSB-BCA
mIoU%	75.52	79.29	78.60	78.44	79.38	75.88	80.03

Table 2: Comparison of context aggregation modules.

Performance along the boundaries. In Table 3, we evaluate the segmentation accuracies in the boundary and interior regions, and compare the proposed BCANet with one state-of-the-art method, *i.e.*, SegFix [25], which is specially designed to refine the segmentation accuracy along the boundaries. Following [25], we employ the boundary F-score with threshold of 0.0003 and the interior F-score to measure the performance. Compared to dilated FCN, both SegFix and BCANet improve the accuracies, while our method performs more favorably in the interior regions and SegFix works better along the boundaries, clearly validating the context aggregation ability of our method. Besides, when combing SegFix with our method, higher accuracy is reported which confirms the complementary of our method and SegFix.

Metrics	Dilated FCN	+SegFix	+MSB+BCA	+MSB+BCA+SegFix
F-score (boundary)	57.93	61.70	60.38	63.92
F-score (interior)	75.07	75.26	77.02	77.13
mIoU	75.52	76.39	80.03	80.86

Table 3: Comparison of segmentation accuracy in the boundary and interior regions.

4.3 Visualization

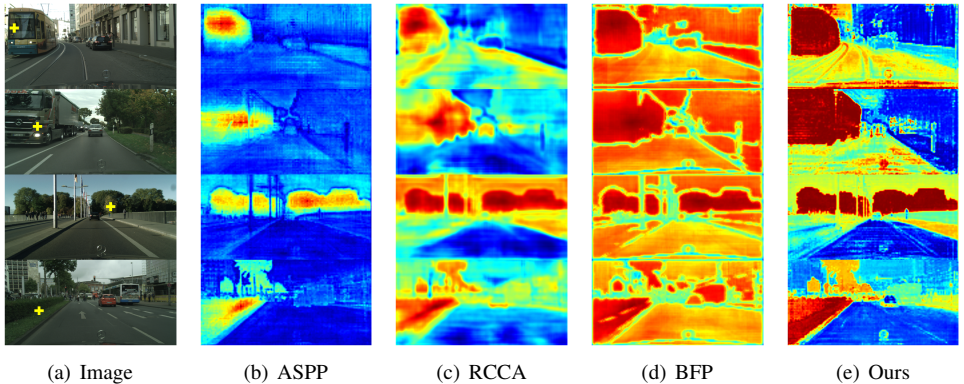


Figure 3: Visualization of the cosine similarity between the marked pixel and the whole feature map after context aggregation. (a) The input image; (b) ASPP [5]; (c) RCCA [11], which is an improved version of Non-local module; (d) BFP [8]; and (e) the proposed BCA module.

Feature similarity visualization. BCA is designed to make the pixels from the identical category achieve mutual gains. Accordingly, we visualize the feature similarity after context aggregation to see if the method is able to enhance the intra-class consistency and inter-class discrimination. We use cosine distance to measure the similarity between two features, and the results are visualized in Figure 3, where the reference pixel is marked by the yellow cross. In particular, ASPP [5] aggregates pixels of the same category efficiently, but the intra-class consistency tends to be weak for large objects, mainly because it neglects the long-range context. The RCCA module [11] in CCNet is an improved version of Non-local module which aggregates the full-image contexts. In this counterpart, the pixels from the same category indeed have larger similarities; however, it also introduces many unreliable contexts, indicating that the Non-local based module cannot exclude inter-class dependencies. In terms of BFP [8], it can be obviously observed that inter-class discrimination is weak that the context information will flow across the whole image when the boundary closure cannot be guaranteed. Comparatively, our method performs more favorably, demonstrating that boundary context is effective for semantic segmentation.

Attention map visualization. In the proposed BCA module, when aggregating the contexts into a marked reference point, positions along the corresponding boundary regions are expected to be activated with much higher weights. To validate this assumption, we visualize the attention map of different positions in Figure 4. In particular, the red and green points belong to the same class and they have similar attention maps, but the attention map of yellow point is significantly different. The visualization results show that the proposed method




Figure 4: Visualization of the attention maps. (a) The input image; attention maps of (b) the red point, (c) the green point, and (d) the yellow point.

indeed assists to enhance the intra-class consistency and inter-class discrimination.

Segmentation result visualization. Figure 5 shows the final segmentation results achieved by BCANet and the other methods modeling long-range contexts. Take the images in the first row for example, BCANet successfully segments the complete bus while the compared counterparts are struggling with the complex interior variations. In second row, the proposed model is able to separate the ground and the tree more clearly and it segments the truck more accurately in the third row. All the results have validated the effectiveness of our method.




Figure 5: Comparison of the final segmentation results. (a) The initial image; (b) the ground-truth; (c) CCNet [11]; (d) BFP [8]; (e) the proposed BCANet.

4.4 Comparison with the State-of-the-art

On Cityscapes Test Set. We conduct experiments on Cityscapes test set and quantitatively evaluate the proposed method on the official evaluation server. To ensure the comparison fairness, multi-scale inference and flipping strategies are exploited during testing. The compared methods include [8, 9, 11, 13, 16, 17, 21, 22, 23, 24, 29, 30, 31, 33] and Table 4 shows the results. As can be seen, the proposed BCANet achieves 81.7% mIoU on Cityscapes test set, which outperforms most of the existing studies. GSCNN [17] achieves the highest segmentation accuracy mainly because a stronger backbone and external dataset are utilized. RPCNet [31] has cascaded more ResNet blocks to obtain the features of a larger resolution (1/4 initial input) to generate final results, which tends to benefit small objects, while our method and the other studies generally use the features of 1/8 scale of the given image.

Model	Backbone	Cityscapes	ADE20K	
		mIoU%	mIoU%	pixAcc%
RefineNet[13]	ResNet-101	73.6	40.2	-
GCN[6]	ResNet-101	76.9	-	-
PSPNet[29]	ResNet-101	78.4	43.29	81.39
BiSeNet[21]	ResNet-101	78.9	-	-
DFN[22]	ResNet-101	79.3	-	-
PSANet[30]	ResNet-101	80.1	43.77	81.51
EncNet[26]	ResNet-101	-	44.65	81.69
ANL[53]	ResNet-101	81.3	45.24	-
DANet[9]	ResNet-101	81.5	-	-
CCNet[11]	ResNet-101	81.4	45.22	-
GSCNN[17]	WideResNet38	82.8	-	-
BFP[8]	ResNet-101	81.4	-	-
OCR[24]	ResNet-101	81.8	45.28	-
RPCNet[61]	ResNet-101	81.8	-	-
CPNet[23]	ResNet-101	81.3	46.27	81.85
BCANet	ResNet-101	81.7	45.62	82.35

Table 4: Segmentation results on Cityscapes test set and ADE20K validation set.

Among the models with fair settings, our results are comparable with the current state-of-the-art.

On ADE20K Validation Set. We compare BCANet with [11, 13, 23, 24, 26, 29, 30, 53]. As shown in Table 4, the proposed method achieves the best performance of 82.35% under the metric of pixel accuracy (pixAcc). In terms of mIoU, it achieves comparable result with most state-of-the-art methods while is a little inferior to CPNet [23]. It is mainly because both intra-class context prior and the reversed prior are considered in CPNet, which favors the performance of all the classes. Comparatively, our model is more conducive to the situation that ambiguity occurs inside the object with relatively large spatial spans, thus it performs more favorably in terms of pixel accuracy.

5 Conclusion

In this paper, we propose a novel approach to semantic segmentation, which explicitly exploits boundary as the aggregated context to improve intra-class consistency. In particular, MSB and BCA modules are specially designed to perform such boundary guided context aggregation. On the two widely applied benchmarks of Cityscapes and ADE20K, we have experimentally demonstrated that the proposed model performs more favorably than the existing attention-based methods by leveraging the partitioning capability of boundary. It should be noted that it may theoretically incur errors when a dominating foreground object shares large boundary areas with the background, since its features tend to leak into the background. We observe the datasets and this case rarely happens, but we keep it in mind and will work for such a corner case in the future.

Acknowledgements

This work is partly supported by the National Natural Science Foundation of China (No. 62022011, U20B2069), the Research Program of State Key Laboratory of Software Development Environment (SKLSDE-2021ZX-04), and the Fundamental Research Funds for the Central Universities.

References

- [1] Yağiz Aksoy, Tae-Hyun Oh, Sylvain Paris, Marc Pollefeys, and Wojciech Matusik. Semantic soft segmentation. *ACM Transactions on Graphics*, 37(4):1–13, 2018.
- [2] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. Semantic segmentation with boundary neural fields. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3602–3610, 2016.
- [3] Liang-Chieh Chen, Jonathan T Barron, George Papandreou, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with task-specific edge detection using cnns and a discriminatively trained domain transform. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 4545–4554, 2016.
- [4] Liang-Chieh Chen, Yi Yang, Jiang Wang, Wei Xu, and Alan L Yuille. Attention to scale: Scale-aware semantic image segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3640–3649, 2016.
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2017.
- [6] Liang-Chieh Chen, Raphael Gontijo Lopes, Bowen Cheng, Maxwell D Collins, Ekin D Cubuk, Barret Zoph, Hartwig Adam, and Jonathon Shlens. Naive-student: Leveraging semi-supervised learning in video sequences for urban scene segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 695–714. Springer, 2020.
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3213–3223, 2016.
- [8] Henghui Ding, Xudong Jiang, Ai Qun Liu, Nadia Magnenat Thalmann, and Gang Wang. Boundary-aware feature propagation for scene segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6819–6829, 2019.
- [9] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3146–3154, 2019.
- [10] 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.
- [11] Zilong Huang, Xinggang 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.

- [12] Xiangtai Li, Xia Li, Li Zhang, Guangliang Cheng, Jianping Shi, Zhouchen Lin, Shao-hua Tan, and Yunhai Tong. Improving semantic segmentation via decoupled body and edge supervision. In *Proceedings of the European Conference on Computer Vision*, pages 435–452. Springer, 2020.
- [13] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1925–1934, 2017.
- [14] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.
- [15] Markus Oberweger, Paul Wohlhart, and Vincent Lepetit. Hands deep in deep learning for hand pose estimation. *arXiv preprint arXiv:1502.06807*, 2015.
- [16] Chao Peng, Xiangyu Zhang, Gang Yu, Guiming Luo, and Jian Sun. Large kernel matters—improve semantic segmentation by global convolutional network. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 4353–4361, 2017.
- [17] Towaki Takikawa, David Acuna, Varun Jampani, and Sanja Fidler. Gated-scnn: Gated shape cnns for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5229–5238, 2019.
- [18] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- [19] 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.
- [20] Minghao Yin, Zhuliang Yao, Yue Cao, Xiu Li, Zheng Zhang, Stephen Lin, and Han Hu. Disentangled non-local neural networks. In *Proceedings of the European Conference on Computer Vision*, pages 191–207. Springer, 2020.
- [21] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 325–341, 2018.
- [22] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Learning a discriminative feature network for semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1857–1866, 2018.
- [23] Changqian Yu, Jingbo Wang, Changxin Gao, Gang Yu, Chunhua Shen, and Nong Sang. Context prior for scene segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 12416–12425, 2020.

- [24] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Object-contextual representations for semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 173–190. Springer, 2020.
- [25] Yuhui Yuan, Jingyi Xie, Xilin Chen, and Jingdong Wang. Segfix: Model-agnostic boundary refinement for segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 489–506. Springer, 2020.
- [26] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaoang Wang, Amrith Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7151–7160, 2018.
- [27] Li Zhang, Xiangtai Li, Anurag Arnab, Kuiyuan Yang, Yunhai Tong, and Philip HS Torr. Dual graph convolutional network for semantic segmentation. In *Proceedings of the British Machine Vision Conference*, 2019.
- [28] Li Zhang, Dan Xu, Anurag Arnab, and Philip HS Torr. Dynamic graph message passing networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3726–3735, 2020.
- [29] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaoang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2881–2890, 2017.
- [30] Hengshuang Zhao, Yi Zhang, Shu Liu, Jianping Shi, Chen Change Loy, Dahua Lin, and Jiaya Jia. Pscanet: Point-wise spatial attention network for scene parsing. In *Proceedings of the European Conference on Computer Vision*, pages 267–283, 2018.
- [31] Mingmin Zhen, Jinglu Wang, Lei Zhou, Shiwei Li, Tianwei Shen, Jiaxiang Shang, Tian Fang, and Long Quan. Joint semantic segmentation and boundary detection using iterative pyramid contexts. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 13666–13675, 2020.
- [32] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 633–641, 2017.
- [33] Zhen Zhu, Mengde Xu, Song Bai, Tengting Huang, and Xiang Bai. Asymmetric non-local neural networks for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 593–602, 2019.