

Global Aggregation then Local Distribution in Fully Convolutional Networks

Xiangtai Li¹

lxtpku@pku.edu.cn

Li Zhang²

lz@robots.ox.ac.uk

Ansheng You¹

youansheng@pku.edu.cn

Maoke Yang³

maokeyang@deepmotion.ai

Kuiyuan Yang³

kuiyuanyang@deepmotion.ai

Yunhai Tong¹

yhtong@pku.edu.cn

¹ Key Laboratory of Machine Perception,
School of EECS,
Peking University

² Department of Engineering Science,
Torr Vision Group,
University of Oxford

³ DeepMotion AI Research

Abstract

It has been widely proven that modelling long-range dependencies in fully convolutional networks (FCNs) via global aggregation modules is critical for complex scene understanding tasks such as semantic segmentation and object detection. However, global aggregation is often dominated by features of large patterns and tends to oversmooth regions that contain small patterns (*e.g.*, boundaries and small objects). To resolve this problem, we propose to first use *Global Aggregation* and then *Local Distribution*, which is called GALD, where long-range dependencies are more confidently used inside large pattern regions and vice versa. The size of each pattern at each position is estimated in the network as a per-channel mask map. GALD is end-to-end trainable and can be easily plugged into existing FCNs with various global aggregation modules for a wide range of vision tasks, and consistently improves the performance of state-of-the-art object detection and instance segmentation approaches. In particular, GALD used in semantic segmentation achieves new state-of-the-art performance on Cityscapes test set with mIoU 83.3%. Code is available at: <https://github.com/lxtGH/GALD-Net>

1 Introduction

Detection and segmentation tasks have made steady progress with more powerful representations learned from Fully Convolutional Networks (FCNs). Since stacking more convolutional layers is not an effective way to achieve large receptive fields for long-range dependency modeling [20, 25], several Global Aggregation (GA) modules have been proposed to resolve this problem.



Figure 1: Our proposed GALD framework for semantic segmentation task. The imbalanced spread of information from small and large patterns in GA module is appropriately handled through LD module.

In contrast to a standard convolutional layer which aggregates features in a small local window, GA modules use long-range operators such as averaging pooling [8, 33] and spatial-wise feature propagation over the whole image [11, 22, 28]. FCNs coupled with GA modules have consistently improved basic FCNs especially for large objects.

Unfortunately, the advantage of GA modules for large objects is a disadvantage for small patterns such as object boundaries and small objects, where features from GA modules tends to oversmooth the predictions for these small patterns. Thus, a straightforward idea is using GA features conditionally on the pattern size of each position. Accordingly, we propose a Local Distribution (LD) module after a GA module (together as GALD for short) to adaptively distribute GA features at each position as illustrated in Fig. 1. The adaptive process is controlled by a set of mask maps, where each mask map is estimated from a feature map that records activations of some latent pattern over the whole image.

LD is a simple and universal module, and can be combined with existing GA modules to form different GALD modules for various detection and segmentation tasks. In our experiment, LD is verified on GA modules such as PSP [33], ASPP [6], Non-Local [22] and CGNL [29], and achieves consistent performance improvement. We also extensively verify GALD on three vision benchmarks, including Cityscapes for semantic segmentation, Pascal VOC 2007 for object detection, and MS COCO for both object detection and instance segmentation, and all achieve notable improvement. In particular, for semantic segmentation evaluated on Cityscapes test set, GALD achieves mIoU of 83.3% with single model and ResNet101 as our backbone network, which surpasses all previously best published single-model results using ResNet101 as backbone network.

2 Related Work

To keep spatial information required by detection and segmentation tasks, convolutional networks designed for image classification are modified to FCNs by removing global information aggregation layers such as global average pooling layer and fully-connected layers [19]. To quickly increase receptive field size while keep the spatial resolution, filters in top convolutional layers are enlarged by dilation [2, 6].

To further enlarge the receptive field to the whole image, several methods are proposed recently. Global average pooled features are concatenated into existing feature maps in [18]. In PSPnet [33], average pooled features of multiple window sizes including global average pooling are upsampled to the same size and concatenated together to enrich global information. The DeepLab series of papers [2, 3, 5] propose atrous or dilated convolutions and atrous spatial pyramid pooling (ASPP) to increase the effective receptive field. DenseASPP [25] improves on [5] by densely connecting convolutional layers with different dilation rates to further increase the receptive field of network. In addition to concatenating global information into feature maps, multiplying global information into feature maps also shows better

performance [23, 27, 29, 30]. In particular, EncNet [30] and DFN [27] use attention along the channel dimension of the convolutional feature map to account for global context such as the co-occurrences of different classes in the scene. CBAM[23] explores channel and spatial attention in cascade way to learn task specific representation.

Recently, advanced global information modeling approaches initiated from non-local network [22] are showing promising results on scene understanding tasks. In contrast to convolutional operator where information are aggregated locally defined by filters, non-local operator aggregates information from the whole image based on an affinity matrix calculated among all positions around the image. Using non-local operator, impressive results are achieved in OCNet [28], CoCurNet [51], DANet [11], A2Net [7], CCnet [12] and Compact Generalized Non-Local Net [29]. OCNet [28] uses non-local blocks to learn pixel-wise relationship while CoCurNet [51] adds extra global average pooling path to learn whole scene statistic. DANet [11] explores orthogonal relationships in both channel and spatial dimension using non-local operator. CCnet [12] models the long range dependencies by considering its surrounding pixels on the criss-cross path through a recurrent way to save both computation and memory cost. Compact Generalized non-local Net [29] considers channel information into affinity matrix. Another similar work to model the pixel-wised relationship is PSANet [54]. It captures pixel-to-pixel relations using an attention module that takes the relative location of each pixel into account.

Another way to get global representation is using graph convolutional networks, and do reasoning in a non-euclidean space [8, 14] where messages are passing between each node before projection back to each position. Glore [8] projects the feature map into interaction space using learned projection matrix and does graph convolution on projected fully connected graph. BeyondGrids [14] learns to cluster different graph nodes and does graph convolution in parallel.

All previous work focus on global context modeling, our work also utilizes global information modeling but takes a further step to better distribute the global information to each position, and further improves GA modules on both detection and segmentation tasks.

3 Method

3.1 Model Overview

Our method, Global Aggregation (GA) then Local Distribution (LD), dubbed GALD, exploits the long-range contextual information of the feature $\mathbf{F} \in \mathbb{R}^{H \times W \times C}$ from a fully-convolution network (FCN), and then adaptively distributed the global context to each spatial and channel position of the output feature, $\mathbf{F}_{GALD} \in \mathbb{R}^{H \times W \times C}$. To be noted, one can choose any one of the methods discussed in Section 2 as GA.

3.2 GALD

Global Aggregation. To calculate a feature vector for each position, GA module takes feature vectors of \mathbf{F} in a large window even the whole feature map depending on different GA designs. Take the Compact Generalized Non-Local (CGNL) [29] as an example, similar to non-local [22], it aggregates contextual information from all spatial and channel positions in the same group. Specifically, a global statistics is calculated for each group and multiplied back to the features in the same group, which forms \mathbf{F}_{GA} . In our implementation, we

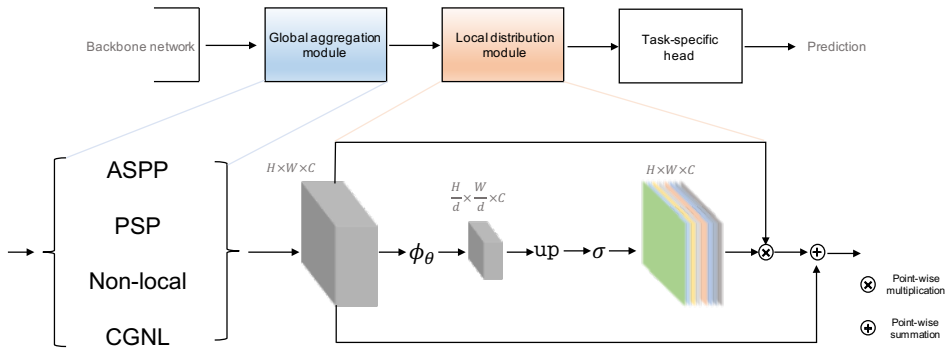


Figure 2: Schematic illustration of GALD, which contains two main components: Global Aggregation (GA) and Local Distribution (LD). GALD receives a feature map from the backbone network and outputs a feature map with same size with global information appropriately assigned to each local position.

downsample \mathbf{F} by a factor of 2 for saving memory and computation cost without observing performance degradation, which also demonstrates the coarse property of global aggregation.

Since GA modules calculate global statistics of features in large windows, which are easily biased towards features from large patterns as they contain more samples. Then the global information distributed to each position is also biased towards large patterns, which causes over smoothing results for small patterns. One can refer to Section 4.3.1 for more detailed visualization results.

Local Distribution. LD is proposed to adaptively use \mathbf{F}_{GA} considering patterns on each position. Without explicit supervision, the required patterns are latently described by C channels in \mathbf{F}_{GA} . For each pattern $c \in \{1, \dots, C\}$, a spatial operator is learned to recalculate the spatial extent of the pattern in an image based on the activation map $\mathbf{F}_{GA}[:, :, c]$ sliced from \mathbf{F}_{GA} . Intuitively, spatial operators for large patterns would shrink the spatial extent more while shrink less even expand for small patterns.

The spatial operators for each pattern/channel is modeled as a set of depth-wise convolutional layers with \mathbf{F}_{GA} as input, i.e.,

$$\mathbf{M} = \sigma(\text{upsample}(\mathbf{W}_d \mathbf{F}_{GA})), \quad (1)$$

where $\mathbf{M} \in [0, 1]^{H \times W \times C}$ contains the mask maps for each pattern and describes the recalculated spatial extents of each pattern, $\sigma(\cdot)$ is the sigmoid function, \mathbf{W}_d is the weights of those depth-wised convolutional filters with d as the downsampling rate by stride convolution. The output mask \mathbf{M} is sensitive to both spatial and channel and it is upsampled using bilinear interpolation. With the mask maps \mathbf{M} , \mathbf{F}_{GA} is refined into \mathbf{F}_{GALD} by

$$\mathbf{F}_{GALD} = \mathbf{M} \odot \mathbf{F}_{GA} + \mathbf{F}_{GA}, \quad (2)$$

where \odot the element-wise multiplication, and elements in \mathbf{F}_{GA} are weighted according to the estimated spatial extent of each pattern at each position. In summary, LD predicts local weights \mathbf{M} for each position of GA features and avoids issues of coarse feature representation.

As a common practice [63], original feature \mathbf{F} and global aggregated feature \mathbf{F}_{GA} are concatenated together for final task-specific head, i.e.,

$$\begin{aligned}\mathbf{F}_o &= \text{concat}(\mathbf{F}_{GALD}, \mathbf{F}) \\ &= \text{concat}(\mathbf{M} \odot \mathbf{F}_{GA} + \mathbf{F}_{GA}, \mathbf{F}),\end{aligned}\quad (3)$$

where \mathbf{M} adds point-wise trade-off between global information \mathbf{F}_{GA} and local detailed information \mathbf{F} . Note that since the lack of details in GA, LD module only changes the proportion and distribution of coarse features in GA and leads to a fine-grained feature representation output \mathbf{F}_o .

3.3 Overall Architecture

Fig. 2 illustrates the overall architecture with GALD. For semantic segmentation, GALD is added right after a FCN, features from Eq. 3 are used for final prediction. To further boost the performance, Online Hard Example Mining (OHEM) loss [24] is used for training, where only top-K ranked pixels according their losses are used during back-propagation.

For object detection and instance segmentation task, GALD is added at the end of stage4 of a ResNet backbone, FPN [16] is used to build a strong baseline with a feature pyramid for multi-scale object detection. \mathbf{F}_{GALD} sits on top of FPN and passes information from the top-down pathway.

4 Experiment

In this section, we verify GALD on three scene understanding tasks including semantic segmentation, object detection and instance segmentation.

4.1 Benchmarks

Cityscapes: Cityscapes [9] is a benchmark that densely annotated for 19 categories in urban scenes, which contains 5000 fine annotated images in total and is divided into 2975, 500, and 1525 images for training, validation and testing, respectively. In addition, 20,000 coarse labeled images are also provided to enrich the training data. Images of this dataset are all with the same high resolution, i.e., 1024×2048 . Following the standard protocol [9], mean Intersection over Union (mIoU) of all categories on validation set and test set is used for performance comparison.

MS COCO: MS COCO [15] is built for detecting and segmenting objects found in everyday life in their natural environment. The dataset for detection consists of three sets for 80 common object categories, i.e., the training set has 118,287 images, validation set has 5,000 images and test-dev set has more than 20,000 images.

Pascal VOC: Pascal VOC [10] is a widely used public benchmark for semantic segmentation and object detection covering 20 object categories including the background. We use VOC 2007 and VOC 2012 trainval set as training set and report results on VOC 2007 test set.

4.2 Implementation Details

Semantic Segmentation We employ Fully Convolutional Networks (FCNs) as baseline, where ResNet pretrained on ImageNet is chosen as the backbone following the same setting




Figure 3: Ablation studies on combinations of GA and LD. (a)-(d) shows the different GA modules with LD. (e)-(g) shows the different arrangements of GA and LD. (h)-(f) represents using GA and LD respectively.

as PSPNet [63], the proposed GALD is appended to the backbone with random initialization. For optimization, we also keep the same setting as PSPNet, where mini-batch SGD with momentum 0.9 and initial learning rate 0.01 is used to train all models with 50K iterations, using mini-batch size of 8 and crop size of 769. During training, “poly” learning rate scheduling policy where power = 0.9 is used to adjust the learning rate. Synchronized batch normalization [60] is used for better mean / variance estimation across GPUs.

Object Detection and Instance Segmentation For object detection and instance segmentation, mmdetection [10] is used as our baseline implementation for fair comparison. GALD is evaluated for object detection on Pascal VOC based on Faster R-CNN, and for both object detection and instance segmentation on MS COCO based on Mask R-CNN. FPN [16] is used as default setting in all these experiments. For fair comparison, we report all the results that we re-implemented in our framework.

4.3 Results on Cityscapes

Two groups of experiments are conducted on Cityscapes, the first group of experiments verifies the effectiveness of our GALD framework by ablation studies. The second group of experiments compares GALD to the state-of-the-art methods.

4.3.1 Ablation Studies

Comparison with baseline We explore our LD module with four different GA modules as illustrated in Fig. 3 (a)-(d). Table 1(a) first reports the performances of adding four GA modules to the baseline FCN, where all methods are using the same backbone ResNet50 for fair comparison. Obviously, all GA modules significantly improves the baseline FCN on semantic segmentation task, where CGNL performs better than other three GA modules. Table 1(b) reports the results by adding our proposed LD module. Directly using LD alone improves the baseline FCN by 3.8%, which demonstrates that features from FCN have the similar problem as features from GA modules. LD together with four different GA modules consistently improves the corresponding GA module. Comparing with baseline, the combination of CGNL+LD achieves the best performance, and we mainly choose CGNL as our GA module in following experiments.

Arrangements of LD and GA Considering LD module can also improve the baseline, we further study different arrangements of LD and GA as illustrated in Fig. 3 (e)-(g). (f) and (g) represent LDGA and Parallel in Table. 1(c) respectively. LDGA means first doing LD

Method	mIoU(%)	Δa
FCN (Baseline)	73.7	-
+ASPP [9]	77.2	3.5 \uparrow
+NL [10]	78.0	4.3 \uparrow
+PSP [11]	76.2	2.5 \uparrow
+CGNL [12]	78.2	4.5 \uparrow

(a) Ablation study on different GA modules using ResNet50 as backbone.

Method	mIoU(%)	Δa	Δb
FCN (Baseline)	73.7	-	-
+LD	77.5	3.8 \uparrow	-
+PSP + LD	78.9	5.2 \uparrow	2.7 \uparrow
+ASPP + LD	79.5	5.4 \uparrow	2.3 \uparrow
+NL + LD	79.2	5.3 \uparrow	1.2 \uparrow
+CGNL + LD	79.6	5.9 \uparrow	1.4 \uparrow

(b) Ablation study on LD applied on different GA modules using ResNet50 as backbone.

Method	mIoU(%)	Δa
FCN (Baseline)	73.7	-
+Parallel	77.5	3.8 \uparrow
+LDGA	78.1	4.4 \uparrow
+GALD	79.6	5.9 \uparrow

(c) Ablation study on different arrangements of GA and LD using ResNet50 as backbone.

Method	mIoU(%)	Backbone	Δa
FCN (Baseline)	73.7	ResNet50	-
FCN (Baseline)	75.3	ResNet101	-
+CGNL	79.7	ResNet101	4.4 \uparrow
+CGNL+LD	79.6	ResNet50	5.9 \uparrow
+PSP	78.6	ResNet101	4.9 \uparrow
+PSP+LD	78.9	ResNet50	5.2 \uparrow

(d) Ablation study on different backbones.

Method	mIoU(%)	Δb
FCN (Baseline)	73.7	-
FCN + CGNL	78.2	-
+CGNL+LD(depth-wise convolution)	79.6	1.4 \uparrow
+CGNL+LD(bilinear interpolation)	77.6	0.6 \downarrow
+CGNL+LD(average pooling)	76.5	1.7 \downarrow

(e) Ablation study on downsampling strategies for mask estimation in LD using ResNet50 as backbone, where the downsampling ratio is 8.

Table 1: Comparison results on Cityscapes validation set, where Δa denotes the performance difference comparing with baseline, and Δb denotes performance difference between using GALD module and the corresponding GA module. All methods are evaluated with single-scale crop test.

then doing GA while Parallel concatenates the output of LD and GA. Table. 1(c) reports the results of the three different arrangements, where all improve the baseline and GALD achieves best result. Fig. 4 shows the mask maps learned in LDGA and GALD, where mask maps learned by GALD are more focused on regions inside large objects then weight global features more in these regions, while mask maps from LDGA have no obvious focus on large objects since the LD module has not accessed to global feature yet.

Compared with stronger backbone To further prove the effectiveness of our method, we compare GALD using ResNet50 as backbone with a stronger backbone ResNet101 in Table 1(d). Our method achieves similar performance improvement comparing GA modules with stronger backbone which further prove the effectiveness of LD module.

Comparison with different downsampling strategies We also explore three different downsampling strategies for LD, including average pooling, bilinear interpolation and depth-wise stride convolution. Table 1(e) reports the comparison results, depth-wise stride convolution achieves the best result, while average pooling and bilinear interpolation even slightly degrades the performance, which shows that the learnable filters for each channel is important to refine the features from the GA module.

Visualization of GALD To further study the features at different stages, we add another two segmentation heads on features outputted from FCN and GA respectively, the model is fine tuned until converge to analyze segmentation ability of features from different stages. Figure 5 compares the segmentation results, segmentation based on GA resolves the ambiguities in FCN features but also tends to over smoothing regions of small patterns which are shown




Figure 4: Comparison of mask maps learned in different arrangements of GA and LD. The mask maps are calculated by the mean of \mathbf{M} along channel dimension. Best view in color.

Method	Backbone	mIoU(%)
SAC [15] †	ResNet101	78.1
AAF [16] †	ResNet101	79.1
BiSeNet [17] †	ResNet101	78.9
PSANet [18] †	ResNet101	80.1
DFN [19] †	ResNet101	79.3
DenseASPP [20] †	DenseNet161	80.6
Glore [21] †	ResNe50	79.5
Glore [21] †	ResNet101	80.9
DAnet [22] †	ResNet101	81.5
GALDNet †	ResNet50	80.8
GALDNet †	ResNet101	81.8

(a) Results on Cityscapes test server trained with fine-data.

Method	Backbone	mIoU(%)
PSP [23] ‡	ResNet101	81.2
DeepLabv3+ [24] ‡	Xception	82.1
DPC [25] ‡	Xception	82.6
Auto-DeepLab [26] ‡	-	82.1

GALDNet ‡	ResNet101	82.9
GALDNet(+Mapillary) ‡	ResNet101	83.3

(b) Results on Cityscapes test server trained with both fine and coarse data

Table 2: State-of-the-art comparison experiments on Cityscapes test set. † means training with only the train-fine dataset. ‡ means training with both the train-fine and coarse data

in red boxes. Segmentation of GALD keeps the global structure of GA while refines back the details.

4.3.2 Comparison with state-of-the-art

We further compare our results with other state-of-the-art methods in this section. We choose dilated ResNet50 and ResNet101 as backbone models. The results are summarized in Table 2. For fair comparison, we first compare methods trained with only fine annotation data in Table 2(a), and then compare the results with other methods using extra training data in Table 2(b). Following [23], multi-scale crop test is used for final test submission. As illustrated, our method surpasses all previous methods. In particular, our model based on a weak backbone ResNet50 can still achieve comparable performance, which is higher than most methods with stronger backbone. By using extra coarse annotation data for training, our method achieves 82.9% mIoU, which also surpasses the state-of-the-art methods. By further

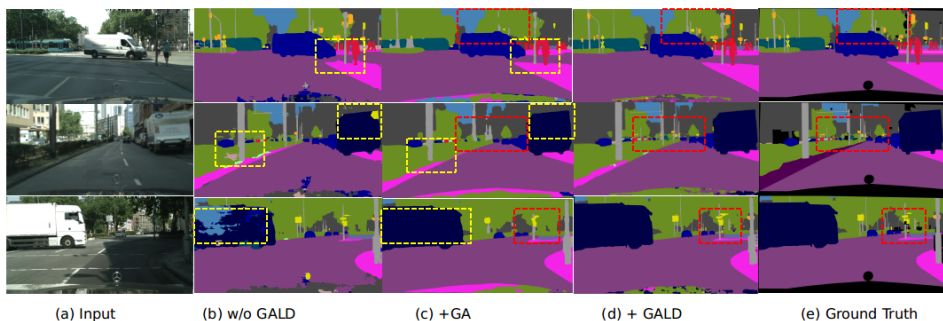


Figure 5: Visualization of different parts output results in one model.(a), input images; (b),results after FCN’s outputs; (c), results after GA module’s outputs; (d), results after GALD module’s outputs; (e), ground truth. Yellow boxes highlight regions that GA can handle global semantic consistency, while red boxes highlight regions that LD can recover more detailed information. Best view in color.

adding Mapillary [20] as training data, the proposed method achieves 83.3% mIoU based on ResNet101. To the best of our knowledge, this is the first single model using ResNet101 as backbone that surpasses 83% mIoU on Cityscapes test server. More detailed per-class results, visualization results and training settings can be referred in the supplementary material.

4.4 Results on Pascal VOC and COCO dataset

Pascal VOC: We perform experiments on the PASCAL VOC 2007 data set to evaluate the effect of GALD for object detection. We train all the models on the union set of VOC 2007 trainval and VOC 2012 trainval (07+12) for 14 epochs with weight decay of 0.0001 and momentum of 0.9. For comparison, experiments of non-local block [27] are also summarized and are denoted as NL. As results listed in Table 3(a), GALD consistently improves detection accuracy over the strong baseline Faster-RCNN using both ResNet50 and ResNet101 as backbone, which demonstrates the effectiveness of GALD for object detection.

COCO: To further verify the generality of GALD, we conduct the experiments on instance segmentation task on MS COCO based on the state-of-the-art method Mask R-CNN. Table 3(b) summarizes the AP of bounding box (AP-box) and AP of mask (AP-mask) evaluated on COCO minival. GALD improves the baseline by about 1% regardless the used backbone. Figure 6(b) compares the object detection and instance segmentation results of our method with baseline. With GALD, Mask R-CNN can find objects that are missed in baseline (e.g., the “light” in the third column), resolve ambiguity in region classification (e.g., the “bed” in the first column) and help to better estimate the spatial contents for objects (e.g., “bear” in last column).

Backbone	Detector	mAP@.5	Backbone	Detector	AP-box	AP-mask
ResNet50	Faster-RCNN	80.6	ResNet50	Mask-RCNN	38.2	34.8
ResNet50	+ NL	81.3 (0.7 \uparrow)	ResNet50	+ NL	39.0 (0.8 \uparrow)	35.3 (0.8 \uparrow)
ResNet50	+CGNL	81.1 (0.5 \uparrow)	ResNet50	+ CGNL	38.9 (0.7 \uparrow)	35.4 (0.6 \uparrow)
ResNet50	+ GALD	81.5 (0.9 \uparrow)	ResNet50	+ GALD	39.2 (1.0 \uparrow)	35.6 (1.1 \uparrow)
ResNet101	Faster-RCNN	80.7	ResNet101	Mask-RCNN	40.2	36.3
ResNet101	+ NL	82.3 (1.6 \uparrow)	ResNet101	+ NL	40.9 (0.7 \uparrow)	37.2 (0.9 \uparrow)
ResNet101	+ GALD	83.0 (2.3 \uparrow)	ResNet101	+ GALD	41.1 (0.9 \uparrow)	37.8 (1.5 \uparrow)

(a) Object detection results on VOC 2007 test set measured by mAP(%), Faster-RCNN with FPN serves as the baseline.

(b) Results of object detection and instance segmentation on COCO dataset. Our method can improve Mask-RCNN baseline by around 1% across different backbones.

Table 3: Results on Pascal VOC dataset (a) and MS COCO dataset (b).



Figure 6: Comparison of object detection and instance segmentation results on MS COCO.

5 Conclusion

In this paper, we propose GALD to adaptively distribute global information to each position for scene understanding tasks. In contrast to existing methods that assign global information uniformly to each position and cause the problem of blurring, GALD learns a set of mask maps to distribute global information adaptively according pattern distributions over the image. GALD benefits from both the GA module for ambiguity resolving and LD module for detail refinement. Extensive experiments verify the universality of GALD in improving the performance of semantic segmentation, object detection and instance segmentation. In the future, we will study the effectiveness of GALD for more vision tasks where both global and local information are important such as depth estimation.

Acknowledgments

We gratefully acknowledge the support of DeepMotion AI Research for providing the computing resources in carrying out this research. LZ is supported by EPSRC Programme Grant Seebibyte EP/M013774/1.

References

- [1] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. mmdetection. <https://github.com/open-mmlab/mmdetection>, 2018.
- [2] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected CRFs. *ICLR*, 2015.
- [3] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Re-thinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [4] Liang-Chieh Chen, Maxwell Collins, Yukun Zhu, George Papandreou, Barret Zoph, Florian Schroff, Hartwig Adam, and Jon Shlens. Searching for efficient multi-scale architectures for dense image prediction. In *NIPS*. 2018.
- [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. *PAMI*, 2018.
- [6] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018.
- [7] Yunpeng Chen, Yannis Kalantidis, Jianshu Li, Shuicheng Yan, and Jiashi Feng. A²-nets: Double attention networks. In *NIPS*. 2018.
- [8] Yunpeng Chen, Marcus Rohrbach, Zhicheng Yan, Shuicheng Yan, Jiashi Feng, and Yannis Kalantidis. Graph-based global reasoning networks. *arXiv preprint arXiv:1811.12814*, 2018.
- [9] 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 *CVPR*, 2016.
- [10] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010.
- [11] Jun Fu, Jing Liu, Haijie Tian, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. *arXiv preprint arXiv:1809.02983*, 2018.
- [12] Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Ccnet: Criss-cross attention for semantic segmentation. *arXiv preprint arXiv:1811.11721*, 2018.
- [13] Tsung-Wei Ke, Jyh-Jing Hwang, Ziwei Liu, and Stella X. Yu. Adaptive affinity fields for semantic segmentation. In *ECCV*, 2018.
- [14] Yin Li and Abhinav Gupta. Beyond grids: Learning graph representations for visual recognition. In *NIPS*. 2018.

- [15] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*. Springer, 2014.
- [16] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.
- [17] Chenxi Liu, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, Wei Hua, Alan Yuille, and Li Fei-Fei. Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation. *arXiv preprint arXiv:1901.02985*, 2019.
- [18] Wei Liu, Andrew Rabinovich, and Alexander C Berg. Parsenet: Looking wider to see better. *arXiv preprint arXiv:1506.04579*, 2015.
- [19] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [20] Wenjie Luo, Yujia Li, Raquel Urtasun, and Richard Zemel. Understanding the effective receptive field in deep convolutional neural networks. In *Advances in neural information processing systems*, pages 4898–4906, 2016.
- [21] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, 2017.
- [22] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [23] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- [24] Zifeng Wu, Chunhua Shen, and Anton van den Hengel. High-performance semantic segmentation using very deep fully convolutional networks. *arXiv preprint arXiv:1604.04339*, 2016.
- [25] Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. Denseaspp for semantic segmentation in street scenes. In *CVPR*, 2018.
- [26] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *ECCV*, 2018.
- [27] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Learning a discriminative feature network for semantic segmentation. In *CVPR*, 2018.
- [28] Yuhui Yuan and Jingdong Wang. Ocnet: Object context network for scene parsing. *arXiv preprint arXiv:1809.00916*, 2018.
- [29] Kaiyu Yue, Ming Sun, Yuchen Yuan, Feng Zhou, Errui Ding, and Fuxin Xu. Compact generalized non-local network. In *NIPS*.
- [30] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaogang Wang, Amrith Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *CVPR*, June 2018.

- [31] Hang Zhang, Han Zhang, Chenguang Wang, and Junyuan Xie. Co-occurrent features in semantic segmentation. In *CVPR*, June 2019.
- [32] Rui Zhang, Sheng Tang, Yongdong Zhang, Jintao Li, and Shuicheng Yan. Scale-adaptive convolutions for scene parsing. In *ICCV*, 2017.
- [33] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017.
- [34] Hengshuang Zhao, Yi Zhang, Shu Liu, Jianping Shi, Chen Change Loy, Dahua Lin, and Jiaya Jia. Psanet: Point-wise spatial attention network for scene parsing. In *ECCV*, 2018.
- [35] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Object detectors emerge in deep scene cnns. *arXiv preprint arXiv:1412.6856*, 2014.