

SegFix: Model-Agnostic Boundary Refinement for Segmentation

Yuhui Yuan^{1,2,4*}, Jingyi Xie^{3*}, Xilin Chen^{1,2}, and Jingdong Wang⁴

¹ Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS), Institute of Computing Technology, CAS

² University of Chinese Academy of Sciences

³ University of Science and Technology of China

⁴ Microsoft Research Asia

{yuhui.yuan, jingdw}@microsoft.com, hsfzxjy@mail.ustc.edu.cn,
xlchen@ict.ac.cn

Abstract. We present a model-agnostic post-processing scheme to improve the boundary quality for the segmentation result that is generated by any existing segmentation model. Motivated by the empirical observation that the label predictions of interior pixels are more reliable, we propose to replace the originally unreliable predictions of boundary pixels by the predictions of interior pixels. Our approach processes only the input image through two steps: (i) localize the boundary pixels and (ii) identify the corresponding interior pixel for each boundary pixel. We build the correspondence by learning a direction away from the boundary pixel to an interior pixel. Our method requires no prior information of the segmentation models and achieves nearly real-time speed. We empirically verify that our SegFix consistently reduces the boundary errors for segmentation results generated from various state-of-the-art models on Cityscapes, ADE20K and GTA5. Code is available at: <https://github.com/openseg-group/openseg.pytorch>.

Keywords: Semantic Segmentation; Instance Segmentation; Boundary Refinement; Model Agnostic

1 Introduction

The task of semantic segmentation is formatted as predicting the semantic category for each pixel in an image. Based on the pioneering fully convolutional network [46], previous studies have achieved great success as reflected by increasing the performance on various challenging semantic segmentation benchmarks [7,16,68].

Most of the existing works mainly addressed semantic segmentation through (i) increasing the resolution of feature maps [12,13,54], (ii) constructing more reliable context information [65,62,23,27,64] and (iii) exploiting boundary information [5,9,44,55]. In this work, we follow the 3rd line of work and focus on

* Equal contribution.

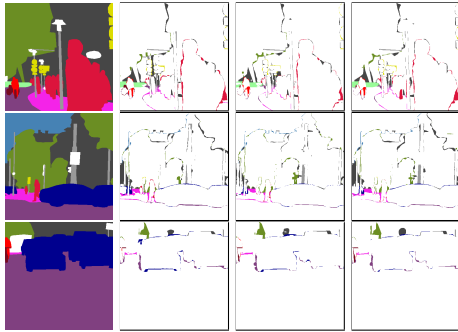


Fig. 1: **Qualitative analysis of the segmentation error maps.** The 1st column presents the ground-truth segmentation maps, and the 2nd/3rd/4th column presents the the error maps of DeepLabv3/HRNet/Gated-SCNN separately. These examples are cropped from Cityscapes val set. We can see that there exist many errors along the thin boundary for all three methods.

improving segmentation result on the pixels located within the thinning boundary⁵ via an effective model-agnostic boundary refinement mechanism.

Our work is mainly motivated by the observation that *most of the existing state-of-the-art segmentation models fail to deal well with the error predictions along the boundary*. We illustrate some examples of the segmentation error maps with DeepLabv3 [12], Gated-SCNN [55] and HRNet [54] in Figure 1. More specifically, we illustrate the statistics on the numbers of the error pixels *vs.* the distances to the object boundaries in Figure 2. We can observe that, for all three methods, the number of error pixels significantly decrease with larger distances to the boundary. In other words, predictions of the interior pixels are more reliable.

We propose a novel model-agnostic post-processing mechanism to reduce boundary errors by replacing labels of boundary pixels with the labels of corresponding interior pixels for a segmentation result. We estimate the pixel correspondences by processing the input image (without exploring the segmentation result) with two steps. The first step aims to localize the pixels along the object boundaries. We follow the contour detection methods [4,2,21] and simply use a convolutional network to predict a binary mask indicating the boundary pixels. In the second step, we learn a direction away from the boundary pixel to an interior pixel and identify the corresponding interior pixel by moving from the boundary pixel along the direction by a certain distance. Especially, our SegFix can reach nearly real-time speed with high resolution inputs.

Our SegFix is a general scheme that consistently improves the performance of various segmentation models across multiple benchmarks without any prior information. We evaluate the effectiveness of SegFix on multiple semantic segmentation benchmarks including Cityscapes, ADE20K and GTA5. We also extend SegFix to instance segmentation task on Cityscapes. According to the Cityscapes leaderboard, “HRNet + OCR + SegFix” and “PolyTransform + SegFix” achieve

⁵ In this paper, we treat the pixels with neighboring pixels belonging to different categories as the boundary pixels. We use the distance transform to generate the ground-truth boundary map with any given width in our implementation.

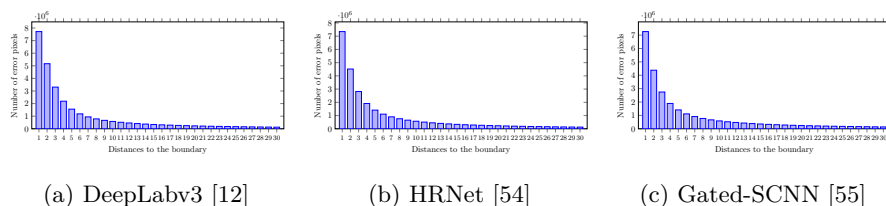


Fig. 2: **Histogram statistics of errors:** the number of error pixels *vs.* their (Euclidean) distances to the boundaries on Cityscapes `val` based on DeepLabv3/HRNet/Gated-SCNN. We can see that pixels with larger distance tend to be well-classified with higher probability and there exist many errors distributing within ~ 5 pixels width along the boundary.

84.5% and 41.2%, which rank the 1st and 2nd place on the semantic and instance segmentation track separately by the ECCV 2020 submission deadline.

2 Related Work

Distance/Direction Map for Segmentation: Some recent work [3,25,56] performed distance transform to compute distance maps for instance segmentation task. For example, [3,25] proposed to train the model to predict the truncated distance maps within each cropped instance. The other work [17,6,10,51] proposed to regularize the semantic or instance segmentation predictions with distance map or direction map in a multi-task mechanism. Compared with the above work, the key difference is that our approach does not perform any segmentation predictions and instead predicts the direction map from only the image, and then we refine the segmentation results of the existing approaches.

Level Set for Segmentation: Many previous efforts [50,8,31] have used the level set approach to address the semantic segmentation problem before the era of deep learning. The most popular formulation of level set is the signed distance function, with all the zero values corresponding to predicted boundary positions. Recent work [1,14,56,33] extended the conventional level-set scheme to deep network for regularizing the boundaries of predicted segmentation map. Instead of representing the boundary with a level set function directly, we implicitly encode the relative distance information of the boundary pixels with a boundary map and a direction map.

DenseCRF for Segmentation: Previous work [11,58,67,41] improved their segmentation results with the DenseCRF [36]. Our approach is also a kind of general post processing scheme while being simpler and more efficient for usage. We empirically show that our approach not only outperforms but also is complementary with the DenseCRF.

Refinement for Segmentation: Extensive studies [22,24,29,37,38] have proposed various mechanisms to refine the segmentation maps from coarse to fine. Different from most of the existing refinement approaches that depend on the

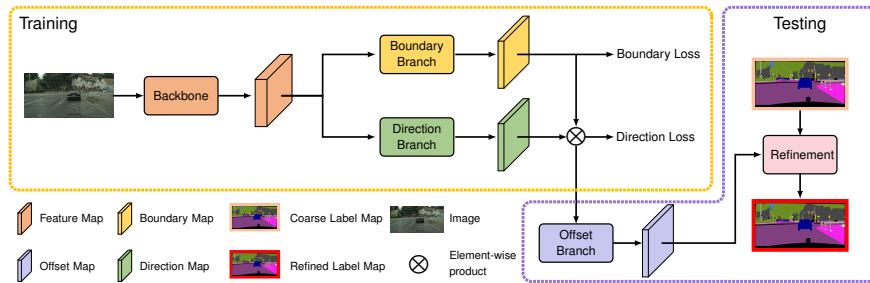


Fig. 3: **Illustrating the SegFix framework:** In the training stage, we first send the input image into a backbone to predict a feature map. Then we apply a boundary branch to predict a binary boundary map and a direction branch to predict a direction map and mask it with the binary boundary map. We apply boundary loss and direction loss on the predicted boundary map and direction map separately. In the testing stage, we first convert the direction map to offset map and then refine the segmentation results of any existing methods according to the offset map.

segmentation models, to the best of our knowledge, our approach is the first model-agnostic segmentation refinement mechanism that can be applied to refine the segmentation results of any approach without any prior information.

Boundary for Segmentation: Some previous efforts [1,59,45,60] focused on localizing semantic boundaries. Other studies [5,55,18,44,43,32,19,20] also exploited the boundary information to improve the segmentation. For example, BNF [5] introduced a global energy model to consider the pairwise pixel affinities based on the boundary predictions. Gated-SCNN [55] exploited the duality between the segmentation predictions and the boundary predictions with a two-branch mechanism and a regularizer.

These methods [5,18,55,32] are highly dependent on the segmentation models and require careful re-training or fine-tuning. Different from them, SegFix does not perform either segmentation prediction or feature propagation and we instead refine the segmentation maps with an offset map directly. In other words, we only need to train a single unified SegFix model once w/o any further fine-tuning the different segmentation models (across multiple different datasets). We also empirically verify that our approach is complementary with the above methods, e.g., Gated-SCNN [55] and Boundary-Aware Feature Propagation [18].

Guided Up-sampling Network: The recent work [47,48] performed a segmentation guided offset scheme to address boundary errors caused by the bi-linear up-sampling. The main difference is that they do not apply any explicit supervision on their offset maps and require re-training for different models, while we apply explicit semantic-aware supervision on the offset maps and our offset maps can be applied to various approaches directly without any re-training. We also empirically verify the advantages of our approach.

3 Approach

3.1 Framework

The overall pipeline of SegFix is illustrated in Figure 3. We first train a model to pick out boundary pixels (with the boundary maps) and estimate their corresponding interior pixels (with offsets derived from the direction maps) from only the image. We do not perform segmentation directly during training. We apply this model to generate offset maps from the images and use the offsets to get the corresponding pixels which should mostly be the more confident interior pixels, and thereby refine segmentation results from any segmentation model. We mainly describe SegFix scheme for semantic segmentation and we illustrate the details for instance segmentation in the Appendix E.

Training stage. Given an input image \mathbf{I} of shape $H \times W \times 3$, we first use a backbone network to extract a feature map \mathbf{X} , and then send \mathbf{X} in parallel to (1) the *boundary branch* to predict a binary map \mathbf{B} , with 1 for the boundary pixels and 0 for the interior pixels, and (2) the *direction branch* to predict a direction map \mathbf{D} with each element storing the direction pointing from the boundary pixel to the interior pixel. The direction map \mathbf{D} is then masked by the binary map \mathbf{B} to yield the input for the offset branch.

For model training, we use a binary cross-entropy loss as the boundary loss on \mathbf{B} and a categorical cross-entropy loss as the direction loss on \mathbf{D} separately.

Testing stage. Based on the predicted boundary map \mathbf{B} and direction map \mathbf{D} , we apply the *offset branch* to generate an offset map $\Delta\mathbf{Q}$. A coarse label map \mathbf{L} output by any semantic segmentation model will be refined as:

$$\tilde{\mathbf{L}}_{\mathbf{p}_i} = \mathbf{L}_{\mathbf{p}_i + \Delta\mathbf{q}_i}, \quad (1)$$

where $\tilde{\mathbf{L}}$ is refined label map, \mathbf{p}_i represents the coordinate of the boundary pixel i , $\Delta\mathbf{q}_i$ is the generated offset vector pointing to an interior pixel, which is indeed an element of $\Delta\mathbf{Q}$. $\mathbf{p}_i + \Delta\mathbf{q}_i$ is the position of the identified interior pixel.

Considering that there might be some “fake” interior pixels⁶ when the boundary is thick, we propose two different schemes as following: (i) re-scaling all the offsets by a factor, e.g., 2. (ii) iteratively applying the offsets (of the “fake” interior pixels) until finding an interior pixel. We choose (i) by default for simplicity as their performance is close.

During testing stage, we only need to generate the offset maps on test set *for once*, and could apply the same offset maps to refine the segmentation results from any existing segmentation model without requiring any prior information. In general, our approach is agnostic to any existing segmentation models.

⁶ We use “fake” interior pixels to represent pixels (after offsets) that still lie on the boundary when the boundary is thick. Notably, we identify an pixel as interior pixel/boundary pixel if its value in the predicted boundary map \mathbf{B} is 0/1.

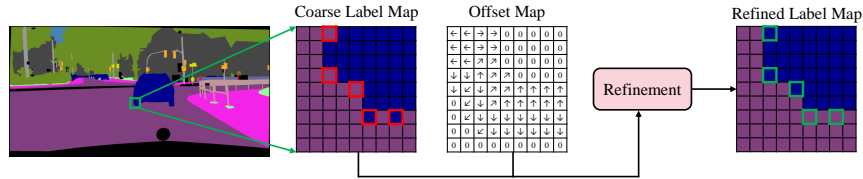


Fig. 4: **Illustrating the refinement mechanism of our approach:** we refine the coarse label map based on the offset map by replacing the labels of boundary pixels with the labels of (more) reliable interior pixels. We represent different offset vectors with different arrows. We mark the error positions in the coarse label map with \square and the corresponding corrected positions in the refined label map with \square . For example, the top-most error pixel (class road) in the coarse label map is associated with a direction \rightarrow . We use the label (class car) of the updated position based on offset $(1, 0)$ as the refined label. Only one-step shifting based on the offset map already refines several boundary errors.

3.2 Network Architecture

Backbone. We adopt the recently proposed high resolution network (HRNet) [54] as backbone, due to its strengths at maintaining high resolution feature maps and our need to apply full-resolution boundary maps and direction maps to refine full-resolution coarse label maps. Besides, we also modify HRNet through applying a 4×4 deconvolution with stride 2 on the final output feature map of HRNet to increase the resolution by $2\times$, which is similar to [15], called Higher-HRNet. We directly perform the boundary branch and the direction branch on the output feature map with the highest resolution. The resolution is $\frac{H}{s} \times \frac{W}{s} \times D$, where $s = 4$ for HRNet and $s = 2$ for Higher-HRNet. We empirically verify that our approach consistently improves the coarse segmentation results for all variations of our backbone choices in Section 4.2, e.g., HRNet-W18 and HRNet-W32.

Boundary branch/loss. We implement the boundary branch as 1×1 Conv \rightarrow BN \rightarrow ReLU with 256 output channels. We then apply a linear classifier (1×1 Conv) and up-sample the prediction to generate the final boundary map \mathbf{B} of size $H \times W \times 1$. Each element of \mathbf{B} records the probability of the pixel belonging to the boundary. We use binary cross-entropy loss as the boundary loss.

Direction branch/loss. Different from the previous approaches [1,3] that perform regression on continuous directions in $[0^\circ, 360^\circ)$ as the ground-truth, our approach directly predicts discrete directions by evenly dividing the entire direction range to m partitions (or categories) as our ground-truth ($m = 8$ by default). In fact, we empirically find that our discrete categorization scheme outperforms the regression scheme, e.g., mean squared loss in the angular domain [3], measured by the final segmentation performance improvements. We illustrate more details for the discrete direction map in Section 3.3.

We implement the direction branch as 1×1 Conv \rightarrow BN \rightarrow ReLU with 256 output channels. We further apply a linear classifier (1×1 Conv) and up-sample the classifier prediction to generate the final direction map \mathbf{D} of size $H \times W \times m$.

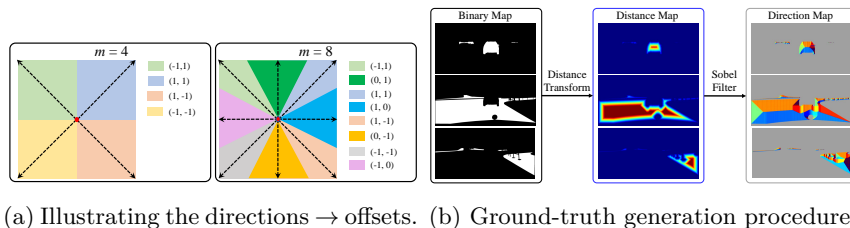


Fig. 5: (a) We divide the entire direction value range $[0^\circ, 360^\circ)$ to m partitions or categories (marked with different colors), For example, when $m = 4$, we have $[0^\circ, 90^\circ)$, $[90^\circ, 180^\circ)$, $[180^\circ, 270^\circ)$ and $[270^\circ, 360^\circ)$ correspond to 4 different categories separately. The above 4 direction categories correspond to offsets $(1, 1)$, $(-1, 1)$, $(-1, -1)$ and $(1, -1)$ respectively. The situation for $m = 8$ is similar. (b) Binary maps \rightarrow Distance maps \rightarrow Direction maps. The ground-truth binary maps are of category car, road and side-walk. We first apply distance transform on each binary map to compute the ground-truth distance maps. Then we use Sobel filter on the distance maps to compute the ground-truth direction maps. We choose different colors to represent different distance values or the direction values.

We mask the direction map \mathbf{D} by multiplying by the (binarized) boundary map \mathbf{B} to ensure that we only apply direction loss on the pixels identified as boundary by the boundary branch. We use the standard category-wise cross-entropy loss to supervise the discrete directions in this branch.

Offset branch. The offset branch is used to convert the predicted direction map \mathbf{D} to the offset map $\Delta\mathbf{Q}$ of size $H \times W \times 2$. We illustrate the mapping mechanism in Figure 5 (a). For example, the “upright” direction category (corresponds to the value within range $[0^\circ, 90^\circ)$) will be mapped to offset $(1, 1)$ when $m = 4$. Last, we generate the refined label map through shifting the coarse label map with the grid-sample scheme [30]. The process is shown in Figure 4.

3.3 Ground-truth generation and analysis

There may exist many different mechanisms to generate ground-truth for the boundary maps and the direction maps. In this work, we mainly exploit the conventional distance transform [34] to generate ground-truth for both semantic segmentation task and the instance segmentation task.

We start from the ground-truth segmentation label to generate the ground-truth distance map, followed by boundary map and direction map. Figure 5 (b) illustrates the overall procedure.

Distance map. For each pixel, our distance map records its minimum (Euclidean) distance to the pixels belonging to other object category. We illustrate how to compute the distance map as below.

First, we decompose the ground-truth label into K binary maps associated with different semantic categories, e.g., car, road, sidewalk. The k^{th} binary map records the pixels belonging to the k^{th} semantic category as 1 and 0 otherwise.

metric	method	w/o SegFix	w/ SegFix	
			$m = 4$	$m = 8$
mIoU	DeepLabv3 (Our impl.)	79.5	82.6 (+3.1)	82.4 (+2.9)
	HRNet-W48 (Our impl.)	81.1	84.1 (+3.0)	84.1 (+3.0)
	Gated-SCNN (Our impl.)	81.0	84.2 (+3.2)	84.1 (+3.1)
F-score	DeepLabv3 (Our impl.)	56.6	68.6 (+12.0)	68.4 (+11.8)
	HRNet-W48 (Our impl.)	62.4	73.8 (+11.4)	73.8 (+11.4)
	Gated-SCNN (Our impl.)	61.4	72.3 (+10.9)	72.3 (+10.9)

Table 1: **Improvements with ground-truth boundary offset on Cityscapes val.** We report both the segmentation performance mIoU and the boundary performance F-score (1px width).

Second, we perform distance transform [34]⁷ on each binary map independently to compute the distance map. The element of k^{th} distance map encodes the distance from a pixel belonging to k^{th} category to the nearest pixel belonging to other categories. Such distance can be treated as the distance to the object boundary. We compute a fused distance map through aggregating all the K distance maps.

Note that the values in our distance map are (always positive) different from the conventional signed distances that represent the interior/exterior pixels with positive/negative distances separately.

Boundary map. As the fused distance map represents the distances to the object boundary, we can construct the ground-truth boundary map through setting all the pixels with distance value smaller than a threshold γ as boundary⁸. We empirically choose small γ value, e.g., $\gamma = 5$, as we are mainly focused on the thin boundary refinement.

Direction map. We perform the Sobel filter (with kernel size 9×9) on the K distance maps independently to compute the corresponding K direction maps respectively. The Sobel filter based direction is in the range $[0^\circ, 360^\circ)$, and each direction points to the interior pixel (within the neighborhood) that is furthest away from the object boundary. We divide the entire direction range to m categories (or partitions) and then assign the direction of each pixel to the corresponding category. We illustrate two kinds of partitions in Figure 5 (a) and we choose the $m = 8$ partition by default. We apply the evenly divided direction map as our ground-truth for training. Besides, we also visualize some examples of direction map in Figure 5 (b).

Empirical Analysis. We apply the generated ground-truth on the segmentation results of three state-of-the-art methods including DeepLabv3 [12], HRNet [54] and Gated-SCNN [55] to investigate the potential of our approach. Specifically, we first project the ground-truth direction map to offset map and then refine the segmentation results on Cityscapes val based on our generated ground-truth offset map. Table 1 summarizes the related results. We can see that our approach significantly improves both the overall mIoU and the boundary F-

⁷ We use `scipy.ndimage.morphology.distance_transform_edt` in implementation.

⁸ We define the boundary pixels and interior pixels based on their distance values.

score. For example, our approach ($m = 8$) improves the mIoU of Gated-SCNN by 3.1%. We may achieve higher performance through re-scaling the offsets for different pixels adaptively, which is not the focus of this work.

Discussion. The key condition for ensuring the effectiveness of our approach is that *segmentation predictions of the interior pixels are more reliable empirically*. Given accurate boundary maps and direction maps, we could always improve the segmentation performance in expectation. In other words, the segmentation performance ceiling of our approach is also determined by the interior pixels’ prediction accuracy.

4 Experiments: Semantic Segmentation

4.1 Datasets & Implementation Details

Cityscapes [16] is a real-world dataset that consists of 2,975/500/1,525 images with resolution 2048×1024 for training/validation/testing respectively, which contains 19/8 semantic categories for semantic/instance segmentation task.

ADE20K [68] is a very challenging benchmark consisting of around 20,000/2,000 images for training/validation respectively. The dataset contains 150 fine-grained semantic categories.

GTA5 [52] is a synthetic dataset that consists of 12,402/6,347/6,155 images with resolution 1914×1052 for training/validation/testing respectively. The dataset contains 19 semantic categories which are compatible with Cityscapes.

Implementation details. We perform the same training and testing settings on Cityscapes and GTA5 benchmarks as follow. We set the initial learning rate as 0.04, weight decay as 0.0005, crop size as 512×512 and batch size as 16, and train for 80K iterations. For the ADE20K benchmark, we set the initial learning as 0.02 and all the other settings are kept the same as on Cityscapes. We use “poly” learning rate policy with power = 0.9. For data augmentation, we all apply random flipping horizontally, random cropping and random brightness jittering within the range of $[-10, 10]$. Besides, we all apply syncBN [53] across multiple GPUs to stabilize the training. We simply set the loss weight as 1.0 for both the boundary loss and direction loss without tuning.

Notably, our approach does not require extra training or fine-tuning any semantic segmentation models. We only need to predict the boundary mask and the direction map for all the test images *in advance* and refine the segmentation results of any existing approaches accordingly.

Evaluation metrics. We use two different metrics including: *mask F-score* and top-1 *direction accuracy* to evaluate the performance of our approach during the training stage. Mask F-score is performed on the predicted binary boundary map and direction accuracy is performed on the predicted direction map. Especially, we only measure the direction accuracy within the regions identified as boundary by the boundary branch.

To verify the effectiveness of our approach for semantic segmentation, we follow the recent Gated-SCNN [55] and perform two quantitative measures including: *class-wise mIoU* to measure the overall segmentation performance on

backbone	#param (M)	runtime (ms)	mask F-score	direction accuracy	mIoU Δ	F-score Δ
HRNet-W18	9.6	16	71.44	64.44	+0.8	+3.7
HRNet-W32	29.4	20	72.24	65.10	+0.9	+3.9
Higher-HRNet	47.3	69	73.67	66.87	+1.0	+4.4

Table 2: **Influence of backbones.** The runtime is tested with an input image of resolution 2048×1024 on a single V100 GPU (PyTorch1.4 + TensorRT). SegFix reaches real-time speed with light-weight backbone, e.g., HRNet-W18 or HRNet-W32.

	boundary width				# directions		
	$\gamma = 3$	$\gamma = 5$	$\gamma = 10$	$\gamma = \infty$	$m = 4$	$m = 8$	$m = 16$
mIoU Δ	+0.94	+0.96	+0.95	+0.84	+0.97	+0.96	+0.96
F-score Δ	+4.1	+4.2	+4.1	+3.6	+4.1	+4.2	+4.2

Table 3: **Influence of the boundary width and direction number.** SegFix is robust to boundary width and direction number. We choose $\gamma = 5$ and $m = 8$ according to their F-scores.

regions; *boundary F-score* to measure the boundary quality of predicted mask with a small slack in distance. In our experiments, we measure the boundary F-score using thresholds 0.0003, 0.0006 and 0.0009 corresponding to 1, 2 and 3 pixels respectively. We mainly report the performance with threshold as 0.0003 for most of our ablation experiments.

4.2 Ablation Experiments

We conduct a group of ablations to analyze the influence of various factors within SegFix. We report the improvements over the segmentation baseline DeepLabv3 (mIoU/F-score is 79.5%/56.6%) if not specified.

Backbone. We study the performance of our SegFix based on three different backbones with increasing complexities, i.e., HRNet-W18, HRNet-W32 and Higher-HRNet. We apply the same training/testing settings for all three backbones. According to the comparisons in Table 2, our SegFix consistently improves both the segmentation performance and the boundary quality with different backbone choices. We choose Higher-HRNet in the following experiments if not specified as it performs best. Besides, we also report their running time in Table 2.

Boundary branch. We verify that SegFix is robust to the choice of hyperparameter γ within the boundary branch and illustrate some qualitative results.

□ **boundary width:** Table 3 shows the performance improvements based on boundary width with different widths. We choose different γ values to control the boundary width, where smaller γ leads to thinner boundaries. We also report the performance with $\gamma = \infty$, which means all pixels is identified as boundary. We find their improvements are close and we choose $\gamma = 5$ by default.

□ **qualitative results:** We show the qualitative results with our boundary branch in the Appendix G. We find that the predicted boundaries are of high quality. Besides, we also compute the F-scores between the boundary computed from the segmentation map of the existing approaches, e.g., Gated-SCNN and




Fig. 6: **Qualitative results of our direction branch predictions.** The 1st and 3rd columns represent the ground-truth segmentation map. The 2nd and 4th columns illustrate the predicted directions with the segmentation map of HRNet as the background. We mark the directions that fix errors with **blue arrow** and directions that lead to extra errors with **red arrow**. Our predicted directions addresses boundary errors for various object categories such as bicycle, traffic light and traffic sign. (Better viewed zoom in)

HRNet, and the predicted boundary from our boundary branch. The F-scores are around 70%, which (in some degree) means that their boundary maps are well aligned and ensures that more accurate direction predictions bring larger performance gains.

Direction branch. We analyze the influence of the direction number m and then present some qualitative results of our predicted directions.

□ **direction number:** We choose different direction numbers to perform different direction partitions and control the generated offset maps that are used to refine the coarse label map. We conduct the experiments with $m = 4$, $m = 8$ and $m = 16$. According to the reported results on the right 3 columns in Table 3, we find different direction numbers all lead to significant improvements and we choose $m = 8$ if not specified as our SegFix is less sensitive to the choice of m .

□ **qualitative results:** In Figure 6, we show some examples to illustrate that our predicted boundary directions improve the errors. Overall, the improved pixels (marked with **blue arrow**) are mainly distributed along the very thin boundary.

Comparison with GUM. We compare SegFix with the previous model-dependent guided up-sampling mechanism [47,48] based on DeepLabv3 as the baseline. We report the related results in Table 4. It can be seen that our approach significantly outperforms GUM measured by both mIoU and F-score. We achieve higher performance through combining GUM with our approach, which achieves 5.0% improvements on F-score compared to the baseline.

Comparison with DenseCRF. We compare our approach with the conventional well-verified DenseCRF [36] based on the DeepLabv3 as our baseline. We fine-tune the hyper-parameters of DenseCRF and set them empirically following [11]. According to Table 5, our approach not only outperforms DenseCRF but also is complementary with DenseCRF. The possible reasons for the limited mIoU improvements of DenseCRF might be that it brings more extra errors on the interior pixels.

	baseline	GUM (Our impl.)	SegFix	GUM+SegFix
mIoU	79.5	79.8 (+0.3)	80.5 (+1.0)	80.6 (+1.1)
F-score	56.6	57.7 (+1.1)	60.9 (+4.3)	61.6 (+5.0)

Table 4: **Comparison with GUM** [47]. SegFix not only outperforms GUM but also is complementary with GUM.

	baseline	DenseCRF	SegFix	DenseCRF+SegFix
mIoU	79.5	79.7 (+0.2)	80.5 (+1.0)	80.5 (+1.0)
F-score	56.6	60.9 (+4.3)	61.0 (+4.4)	64.1 (+7.5)

Table 5: **Comparison with DenseCRF** [36]. SegFix achieves comparable F-score improvements and much larger mIoU gains.

	Gated-SCNN	Gated-SCNN+SegFix
mIoU	81.0	81.5 (+0.5)
F-score	61.4	63.1 (+1.7)

Table 6: **Comparison with Gated-SCNN** [55] The result of Gated-SCNN is based on multi-scale testing.

	ADE20K		GTA5	
	baseline	+SegFix	baseline	+SegFix
mIoU	44.8	45.4(+0.6)	77.8	80.6(+2.8)
F-score	16.4	19.3(+2.9)	50.2	61.7(+11.5)

Table 7: **DeepLabv3 with SegFix on ADE20K and GTA5**. We all choose DeepLabv3 as the baseline.

Application to Gated-SCNN. Considering that Gated-SCNN [55] introduced multiple components to improve the performance, it is hard to compare our approach with Gated-SCNN fairly to a large extent. To verify the effectiveness of our approach to some extent, we first take the open-sourced Gated-SCNN (multi-scale testing) segmentation results on Cityscapes validation set as the coarse segmentation maps, then we apply the SegFix offset maps to refine the results. We report the results in Table 6 and SegFix improves the boundary F-score by 1.7%, suggesting that SegFix is complementary with the strong baseline that also focuses on improving the segmentation boundary quality. Besides, we also report the detailed category-wise improvements measured by both mIoU and boundary F-score in the Table 2 and Table 3 of Appendix.

4.3 Application to State-of-the-art

We generate the boundary maps and the direction maps in advance and apply them to the segmentation results of various state-of-the-art approaches without extra training or fine-tuning.

Cityscapes val: We first apply our approach on various state-of-the-art approaches (on Cityscapes `val`) including DeepLabv3, Gated-SCNN and HRNet. We report the category-wise mIoU improvements in Table 2 of Appendix. It can be seen that our approach significantly improves the segmentation quality along the boundaries of all the evaluated approaches. We provides some qualitative examples of the improvements with our approach along the thin boundaries based on both DeepLabv3 and HRNet in the Figure 1 of Appendix.

Cityscapes test: We further apply our approach on several recent state-of-the-art methods on Cityscapes `test` including PSANet [66], DANet [23], BFP [18], HRNet [54], Gated-SCNN [55], VPLR [69] and HRNet + OCR [61]. We directly

method	road	sidewalk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrian	sky	person	rider	car	truck	bus	train	motorcycle	bicycle	mean
PSANet	98.7	87.0	93.5	58.9	62.5	67.8	76.0	80.0	93.7	72.6	95.4	86.9	73.0	96.2	79.3	91.2	84.9	71.1	77.9	81.4
+ SegFix	98.7	87.4	93.7	59.3	62.8	69.5	77.6	81.4	93.9	73.0	95.6	88.0	73.9	96.5	79.6	91.5	85.1	71.8	78.6	82.0
DANet	98.6	86.1	93.5	56.1	63.3	69.7	77.3	81.3	93.9	72.9	95.7	87.3	72.9	96.2	76.8	89.4	86.5	72.2	78.2	81.5
+ SegFix	98.7	86.6	93.7	56.5	63.5	71.4	78.7	82.4	94.1	73.2	95.9	88.2	73.7	96.5	77.0	89.7	86.8	72.8	78.8	82.0
BFP	98.7	87.0	93.5	59.8	63.4	68.9	76.8	80.9	93.7	72.8	95.5	87.0	72.1	96.0	77.6	89.0	86.9	69.2	77.6	81.4
+ SegFix	98.7	87.5	93.7	60.2	63.7	71.1	78.4	82.4	94.0	73.2	95.7	88.1	72.9	96.3	77.8	89.3	87.2	69.9	78.4	82.0
HRNet	98.8	87.5	93.7	55.6	62.3	71.8	79.3	81.8	94.0	73.1	95.8	88.5	76.1	96.3	72.2	86.5	84.7	73.8	79.4	81.8
+ SegFix	98.8	87.9	93.9	56.0	62.5	73.6	80.7	83.2	94.1	73.4	95.9	89.3	76.7	96.6	72.4	86.7	85.0	74.3	80.2	82.2
VPLR	98.8	87.8	94.2	64.1	65.0	72.4	79.0	82.8	94.2	74.0	96.1	88.2	75.4	96.3	78.8	94.0	91.6	73.8	79.0	83.5
+ SegFix	98.8	88.0	94.3	64.4	65.3	73.3	80.0	83.5	94.3	74.3	96.2	89.0	76.2	96.7	79.0	94.2	92.0	74.4	79.7	83.9
HRNet + OCR	98.9	88.3	94.3	66.8	66.6	73.6	80.3	83.7	94.3	74.4	96.0	88.7	75.4	96.6	82.5	94.0	90.8	73.8	79.7	84.2
+ SegFix	98.9	88.3	94.4	68.0	67.8	73.6	80.6	83.9	94.4	74.5	96.1	89.2	75.9	96.8	83.6	94.2	91.3	74.0	80.1	84.5

Table 8: **Results on Cityscapes Semantic Segmentation task.** Category-wise improvements of SegFix based on various state-of-the-art methods on Cityscapes **test**. Notably, “HRNet + OCR + SegFix” ranks the first place on the Cityscapes semantic segmentation leaderboard by the ECCV 2020 submission deadline.

method	person	rider	car	truck	bus	train	motorcycle	bicycle	mean (%)
Mask-RCNN	36.0	28.8	51.6	30.0	38.7	27.3	23.9	19.4	32.0
+ SegFix	37.9	30.3	54.1	31.0	40.0	27.9	25.1	20.5	33.3 (+1.3)
PointRend	36.6	29.7	53.7	29.9	40.4	33.3	23.6	19.6	33.3
+ SegFix	38.7	31.1	56.2	31.1	41.6	34.1	24.6	20.7	34.8 (+1.5)
PANet	41.5	33.6	58.2	31.8	45.3	28.7	28.2	24.1	36.4
+ SegFix	43.3	34.9	60.4	32.9	47.0	30.1	29.1	24.7	37.8 (+1.4)
PolyTransform	42.4	34.8	58.5	39.8	50.0	41.3	30.9	23.4	40.1
+ SegFix	44.3	35.9	60.5	40.5	51.2	41.6	31.7	24.1	41.2 (+1.1)

Table 9: **Results on Cityscapes Instance Segmentation task.** Our SegFix significantly improves the mask AP of Mask-RCNN [26], PointRend [35], PANet [42] and PolyTransform [40] on Cityscapes **test** (w/ COCO pre-training). Notably, “PolyTransform + SegFix” ranks the second place on the Cityscapes instance segmentation leaderboard by the ECCV 2020 submission deadline.

apply the same model that are trained with only the 2,975 training images without any other tricks, e.g., training with validation set or Mapillary Vistas [49], online hard example mining.

Notably, the state-of-the-art methods have applied various advanced techniques, e.g., multi-scale testing, multi-grid, performing boundary supervision or utilizing extra training data such as Mapillary Vistas or Cityscapes video, to improve their results. In Table 8, our model-agnostic boundary refinement scheme consistently improves all the evaluated approaches. For example, with our SegFix, “HRNet + OCR” achieves 84.5% on Cityscapes **test**. The improvements of our SegFix is in fact already significant considering the baseline is already very strong and the performance gap between top ranking methods is just around 0.1% ~ 0.3%. We believe that lots of other advanced approaches might also benefit from our approach.

We also apply the SegFix scheme on two other challenging semantic segmentation benchmarks including ADE20K and GTA5. Table 7 reports the results and SegFix achieves significant performance improvements along the boundary on both benchmarks. e.g., the boundary F-score of DeepLabv3 gains 2.9%/11.5%

on ADE20K `val`/GTA5 `test` separately. Besides, we propose a unified SegFix model and compare our SegFix to model ensemble in Appendix C and D.

5 Experiments: Instance Segmentation

In Table 9, we illustrate the results of SegFix on Cityscapes instance segmentation task. We can find that the SegFix consistently improves the mean AP scores over Mask-RCNN [26], PANet [42], PointRend [35] and PolyTransform [40]. For example, with SegFix scheme, PANet gains 1.4% points on the Cityscapes `test` set. We also apply our SegFix on the very recent PointRend and PolyTransform. Our SegFix consistently improves the performance of PointRend and PolyTransform by 1.5% and 1.1% separately, which further verifies the effectiveness of our method.

We use the public available checkpoints from Detectron2⁹ and PANet¹⁰ to generate the predictions of Mask-RCNN, PointRend and PANet. Besides, we use the segmentation results of PolyTransform directly. More training/testing details of SegFix on Cityscapes instance segmentation task are illustrated in the Appendix E. We believe that SegFix can be used to improve various other state-of-the-art instance segmentation methods directly w/o any prior requirements.

Notably, the improvements on the instance segmentation tasks (+1.1% ~ 1.5%) are more significant than the improvements on semantic segmentation task (+0.3% ~ 0.5%). We guess the main reason is that the instance segmentation evaluation (on Cityscapes) only considers 8 object categories without including the stuff categories. The performance of stuff categories is less sensitive to the boundary errors due to that their area is (typically) larger than the area of object categories. In summary, our SegFix achieves larger improvements on object categories than stuff categories.

6 Conclusion

In this paper, we have proposed a novel model-agnostic approach to refine the segmentation maps predicted by an unknown segmentation model. The insight is that the predictions of the interior pixels are more reliable. We propose to replace the predictions of the boundary pixels using the predictions of the corresponding interior pixels. The correspondence is learnt only from the input image. The main advantage of our method is that SegFix generalizes well on various strong segmentation models. Empirical results show that the effectiveness of our approach for both semantic segmentation and instance segmentation tasks. We hope our SegFix scheme can become a strong baseline for more accurate segmentation results along the boundary.

Acknowledgement: This work is partially supported by Natural Science Foundation of China under contract No. 61390511, and Frontier Science Key Research Project CAS No. QYZDJ-SSW-JSC009.

⁹ Detectron2: <https://github.com/facebookresearch/detectron2>

¹⁰ PANet: <https://github.com/ShuLiu1993/PANet>

References

1. Acuna, D., Kar, A., Fidler, S.: Devil is in the edges: Learning semantic boundaries from noisy annotations. In: CVPR (2019)
2. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. PAMI (2010)
3. Bai, M., Urtasun, R.: Deep watershed transform for instance segmentation. In: CVPR (2017)
4. Bertasius, G., Shi, J., Torresani, L.: High-for-low and low-for-high: Efficient boundary detection from deep object features and its applications to high-level vision. In: ICCV (2015)
5. Bertasius, G., Shi, J., Torresani, L.: Semantic segmentation with boundary neural fields. In: CVPR (2016)
6. Bischke, B., Helber, P., Folz, J., Borth, D., Dengel, A.: Multi-task learning for segmentation of building footprints with deep neural networks. In: ICIIP (2019)
7. Caesar, H., Uijlings, J., Ferrari, V.: Coco-stuff: Thing and stuff classes in context. In: CVPR (2018)
8. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic active contours. IJCV (1997)
9. Chen, L.C., Barron, J.T., Papandreou, G., Murphy, K., Yuille, A.L.: Semantic image segmentation with task-specific edge detection using cnns and a discriminatively trained domain transform. In: CVPR (2016)
10. Chen, L.C., Hermans, A., Papandreou, G., Schroff, F., Wang, P., Adam, H.: Masklab: Instance segmentation by refining object detection with semantic and direction features. In: CVPR (2018)
11. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. PAMI (2017)
12. Chen, L.C., Papandreou, G., Schroff, F., Adam, H.: Rethinking atrous convolution for semantic image segmentation. arXiv:1706.05587 (2017)
13. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: ECCV (2018)
14. Chen, X., Williams, B.M., Vallabhaneni, S.R., Czanner, G., Williams, R., Zheng, Y.: Learning active contour models for medical image segmentation. In: CVPR (2019)
15. Cheng, B., Xiao, B., Wang, J., Shi, H., Huang, T.S., Zhang, L.: Bottom-up higher-resolution networks for multi-person pose estimation. arXiv preprint arXiv:1908.10357 (2019)
16. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: CVPR (2016)
17. Dangi, S., Yaniv, Z., Linte, C.: A distance map regularized cnn for cardiac cine mr image segmentation. arXiv:1901.01238 (2019)
18. Ding, H., Jiang, X., Liu, A.Q., Thalmann, N.M., Wang, G.: Boundary-aware feature propagation for scene segmentation. ICCV (2019)
19. Ding, H., Jiang, X., Shuai, B., Liu, A.Q., Wang, G.: Semantic correlation promoted shape-variant context for segmentation. In: CVPR (2019)
20. Ding, H., Jiang, X., Shuai, B., Qun Liu, A., Wang, G.: Context contrasted feature and gated multi-scale aggregation for scene segmentation. In: CVPR (2018)
21. Dollár, P., Zitnick, C.L.: Fast edge detection using structured forests. ArXiv (2014)

22. Fieraru, M., Khoreva, A., Pishchulin, L., Schiele, B.: Learning to refine human pose estimation. In: CVPRW (2018)
23. Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., Lu, H.: Dual attention network for scene segmentation. In: CVPR (2019)
24. Gidaris, S., Komodakis, N.: Detect, replace, refine: Deep structured prediction for pixel wise labeling. In: CVPR (2017)
25. Hayder, Z., He, X., Salzmann, M.: Boundary-aware instance segmentation. In: CVPR (2017)
26. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV (2017)
27. Huang, L., Yuan, Y., Guo, J., Zhang, C., Chen, X., Wang, J.: Interlaced sparse self-attention for semantic segmentation. arXiv preprint arXiv:1907.12273 (2019)
28. Huang, Z., Wang, X., Huang, L., Huang, C., Wei, Y., Liu, W.: Ccnet: Criss-cross attention for semantic segmentation. In: ICCV (2019)
29. Islam, M.A., Naha, S., Rochan, M., Bruce, N., Wang, Y.: Label refinement network for coarse-to-fine semantic segmentation. arXiv:1703.00551 (2017)
30. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: NIPS (2015)
31. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. IJCV (1988)
32. Ke, T.W., Hwang, J.J., Liu, Z., Yu, S.X.: Adaptive affinity fields for semantic segmentation. In: ECCV (2018)
33. Kim, Y., Kim, S., Kim, T., Kim, C.: Cnn-based semantic segmentation using level set loss. In: WACV (2019)
34. Kimmel, R., Kiryati, N., Bruckstein, A.M.: Sub-pixel distance maps and weighted distance transforms. JMIV (1996)
35. Kirillov, A., Wu, Y., He, K., Girshick, R.: Pointrend: Image segmentation as rendering. arXiv:1912.08193 (2019)
36. Krähenbühl, P., Koltun, V.: Efficient inference in fully connected crfs with gaussian edge potentials. In: NIPS (2011)
37. Kuo, W., Angelova, A., Malik, J., Lin, T.Y.: Shapemask: Learning to segment novel objects by refining shape priors. ICCV (2019)
38. Li, K., Hariharan, B., Malik, J.: Iterative instance segmentation. In: CVPR (2016)
39. Li, X., Zhong, Z., Wu, J., Yang, Y., Lin, Z., Liu, H.: Expectation-maximization attention networks for semantic segmentation. In: ICCV (2019)
40. Liang, J., Homayounfar, N., Ma, W.C., Xiong, Y., Hu, R., Urtasun, R.: Poly-transform: Deep polygon transformer for instance segmentation. arXiv:1912.02801 (2019)
41. Lin, G., Milan, A., Shen, C., Reid, I.: Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In: CVPR (2017)
42. Liu, S., Qi, L., Qin, H., Shi, J., Jia, J.: Path aggregation network for instance segmentation. In: CVPR (2018)
43. Liu, S., De Mello, S., Gu, J., Zhong, G., Yang, M.H., Kautz, J.: Learning affinity via spatial propagation networks. In: NIPS (2017)
44. Liu, T., Ruan, T., Huang, Z., Wei, Y., Wei, S., Zhao, Y., Huang, T.: Devil in the details: Towards accurate single and multiple human parsing. arXiv:1809.05996 (2018)
45. Liu, Y., Cheng, M.M., Hu, X., Wang, K., Bai, X.: Richer convolutional features for edge detection. In: CVPR (2017)
46. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR (2015)
47. Mazzini, D.: Guided upsampling network for real-time semantic segmentation. arXiv preprint arXiv:1807.07466 (2018)

48. Mazzini, D., Schettini, R.: Spatial sampling network for fast scene understanding. In: CVPRW (2019)
49. Neuhold, G., Ollmann, T., Rota Bulò, S., Kotschieder, P.: The mapillary vistas dataset for semantic understanding of street scenes. In: ICCV (2017)
50. Osher, S., Sethian, J.A.: Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics* (1988)
51. Papandreou, G., Zhu, T., Chen, L.C., Gidaris, S., Tompson, J., Murphy, K.: Personlab: Person pose estimation and instance segmentation with a bottom-up, part-based, geometric embedding model. In: ECCV (2018)
52. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: ECCV (2016)
53. Rota Bulò, S., Porzi, L., Kotschieder, P.: In-place activated batchnorm for memory-optimized training of dnns. In: CVPR (2018)
54. Sun, K., Zhao, Y., Jiang, B., Cheng, T., Xiao, B., Liu, D., Mu, Y., Wang, X., Liu, W., Wang, J.: High-resolution representations for labeling pixels and regions. arXiv:1904.04514 (2019)
55. Takikawa, T., Acuna, D., Jampani, V., Fidler, S.: Gated-scnn: Gated shape cnns for semantic segmentation. ICCV (2019)
56. Wang, Z., Acuna, D., Ling, H., Kar, A., Fidler, S.: Object instance annotation with deep extreme level set evolution. In: CVPR (2019)
57. Wu, Y., Kirillov, A., Massa, F., Lo, W.Y., Girshick, R.: Detectron2. <https://github.com/facebookresearch/detectron2> (2019)
58. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. arXiv preprint arXiv:1511.07122 (2015)
59. Yu, Z., Feng, C., Liu, M.Y., Ramalingam, S.: Casenet: Deep category-aware semantic edge detection. In: CVPR (2017)
60. Yu, Z., Liu, W., Zou, Y., Feng, C., Ramalingam, S., Vijaya Kumar, B., Kautz, J.: Simultaneous edge alignment and learning. In: ECCV (2018)
61. Yuan, Y., Chen, X., Wang, J.: Object-contextual representations for semantic segmentation. arXiv preprint arXiv:1909.11065 (2019)
62. Yuan, Y., Wang, J.: Ocnets: Object context network for scene parsing. arXiv:1809.00916 (2018)
63. Zhang, H., Dana, K., Shi, J., Zhang, Z., Wang, X., Tyagi, A., Agrawal, A.: Context encoding for semantic segmentation. In: CVPR (2018)
64. Zhang, H., Zhang, H., Wang, C., Xie, J.: Co-occurrent features in semantic segmentation. In: CVPR (2019)
65. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network (2017)
66. Zhao, H., Zhang, Y., Liu, S., Shi, J., Loy, C.C., Lin, D., Jia, J.: PSANet: Point-wise spatial attention network for scene parsing. In: ECCV (2018)
67. Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, V., Su, Z., Du, D., Huang, C., Torr, P.H.S.: Conditional random fields as recurrent neural networks. In: ICCV (2015)
68. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralla, A.: Scene parsing through ade20k dataset. In: CVPR (2017)
69. Zhu, Y., Sapra, K., Reda, F.A., Shih, K.J., Newsam, S., Tao, A., Catanzaro, B.: Improving semantic segmentation via video propagation and label relaxation. In: CVPR (2019)