

Active Boundary Loss for Semantic Segmentation

Chi Wang^{1,2}, Yunke Zhang^{1,2}, Miaomiao Cui², Peiran Ren²
 Yin Yang³, Xuansong Xie², Xian-Sheng Hua², Hujun Bao¹, Weiwei Xu^{1*}

¹ State Key Lab of CAD&CG, Zhejiang University

² Alibaba Inc

³ Clemson University

{wangchi1995, yunkezhang}@zju.edu.cn, {miaomiao.cmm, peiran.rpr}@alibaba-inc.com
 yin5@clemson.edu, xingtong.xxs@taobao.com, xiansheng.hxs@alibaba-inc.com, {bao, xww}@cad.zju.edu.cn

Abstract

This paper proposes a novel active boundary loss for semantic segmentation. It can progressively encourage the alignment between predicted boundaries and ground-truth boundaries during end-to-end training, which is not explicitly enforced in commonly used cross-entropy loss. Based on the predicted boundaries detected from the segmentation results using current network parameters, we formulate the boundary alignment problem as a differentiable direction vector prediction problem to guide the movement of predicted boundaries in each iteration. Our loss is model-agnostic and can be plugged in to the training of segmentation networks to improve the boundary details. Experimental results show that training with the active boundary loss can effectively improve the boundary F-score and mean Intersection-over-Union on challenging image and video object segmentation datasets.

Introduction

Semantic segmentation is a fine-grained, pixel-wise classification task that assigns each pixel a semantic class label to facilitate high-level image analysis and processing. Recently, the accuracy of semantic segmentation has been substantially improved with the introduction of fully convolutional networks (FCNs) (Long et al. 2015; Minaee et al. 2021). FCNs leverage convolutional layers and down-sampling operations to achieve a large receptive field. Although these operations can encode context information surrounding a pixel, they tend to propagate feature information throughout the image, leading to undesirable feature smoothing across object boundaries. Thus, the segmentation results might be blurred and lack fine object boundary details. To address this issue, boundary-aware information flow control and multi-task training methods have been proposed to improve the discriminative power of features belonging to different objects (Bertasius et al. 2016; Takikawa et al. 2019; Zhu et al. 2019). Alternatively, the segmentation errors at boundaries can be remedied by learning the correspondence between a boundary pixel and its corresponding interior pixel (Yuan et al. 2020b). Despite the empirical success of boundary-aware methods in improving the segmentation accuracy, there still exist a significant amount of seg-

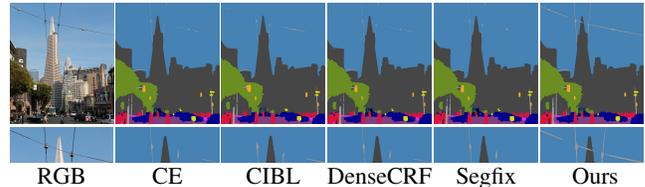


Figure 1: Segmentation results of an internet image. CE: Produced by DeeplabV3 (Chen et al. 2017a) trained with cross-entropy loss on Cityscapes (Cordts et al. 2016) dataset. CIBL: Produced by DeeplabV3 trained with cross-entropy loss + Lovász-softmax (Berman et al. 2018) + Boundary Loss (Kervadec et al. 2019) on Cityscapes dataset. DenseCRF: Refined results of column ‘CE’ by DenseCRF (Krähenbühl et al. 2011). Segfix: Refined results of column ‘CE’ by Segfix (Yuan et al. 2020b). Ours: Re-trained by adding our loss.

mentation errors at object boundaries, especially for small and thin objects. The mutual dependence between semantic segmentation and boundary detection should be further studied to improve the quality of segmentation results.

In this paper, we propose a novel active boundary loss (ABL) to progressively encourage the alignment between predicted boundaries (PDBs) and ground-truth boundaries (GTBs) during end-to-end training, in which the PDBs are semantic boundaries detected in the segmentation results of the current network. To facilitate end-to-end training, the loss is formulated as a differentiable direction vector prediction problem. Specifically, for a pixel on the PDBs, we first determine a direction pointing to the closest GTB pixel, and then move the PDB at this pixel towards the direction in a probabilistic manner. Moreover, we also propose to detach the gradient flow to suppress possible conflicts. Overall, the behavior of ABL is dynamic because the PDBs are changing with the updated network parameters during training. It can be viewed as a variant of classical active contour methods (Kass et al. 1988), since our method first determines the direction vectors in accordance with the PDBs in the current iteration and lets the PDBs move along the direction vectors to reach the GTBs.

Unlike the cross-entropy loss that only supervises pixel-level classification accuracy, ABL supervises the relationship between PDB and GTB pixels. It embeds boundary in-

*Corresponding author.

formation such that the network can pay attention to boundary pixels to improve the segmentation results. Moreover, Intersection-over-Union (IoU) loss pays more attention to overall regions of semantic classes but does not focus on the boundary matching. Thus, ABL can provide complementary information during the training of the network. As a result, it can be combined with other loss terms to further improve semantic boundary quality. In our work, we let the ABL work with the most commonly used cross-entropy loss and the lovasz-softmax loss (Berman et al. 2018), a surrogate IoU loss, to significantly improve the boundary details in image segmentation. The lovasz-softmax loss is introduced to regularize the training so that the ABL can be used even when the PDBs might be noisy and far from the GTBs. The advantage of ABL is that it is model-agnostic and can be plugged into the training of image segmentation networks to improve the boundary details. As illustrated in Fig. 1, it is beneficial to preserve the boundaries of thin objects that contain a small number of interior pixels.

We tested the ABL with state-of-the-art image segmentation networks, including CNN-based networks DeepLabV3 (Chen et al. 2017a), OCR network (Yuan et al. 2020a) and Transformer-based network SwinTransformer (Liu et al. 2021). We have also tested the ABL with STM (Oh et al. 2019), a video object segmentation (VOS) network, to show that our loss can be applied to improve VOS results as well. The forward inference stage of these networks remains the same during testing. The experimental results show that training with the ABL can effectively improve the boundary F-score and mean Intersection-over-Union (mIoU) on challenging segmentation datasets.

Related Work

FCN-based semantic segmentation. FCNs (Long et al. 2015) for semantic segmentation frequently utilize encoder-decoder structure to generate pixel-wise labelling results for high-resolution images. Successor methods (Ronneberger et al. 2015; Ding et al. 2018; Minaee et al. 2021) are dedicated to a better fusion of multi-scale features to enhance the accuracy of localization and handle small objects.

FCN-based methods have also been widely used in VOS, including propagation-based methods (Hu et al. 2017; Oh et al. 2019; Voigtlaender et al. 2019) and detection-based methods (Caelles et al. 2017; Li et al. 2017; Shin Yoon et al. 2017). The key challenge is how to leverage temporal coherence and learn discriminative features of target objects to handle occlusion, appearance change, and fast motion. Since our loss is model-agnostic, it can also be applied to VOS for the purpose of boundary refinement.

Boundary-aware semantic segmentation. One way to exploit boundary information in deep learning-based semantic segmentation is through multi-task training, in which additional branches are often inserted to detect semantic boundaries (Chen et al. 2020; Gong et al. 2018; Ruan et al. 2019; Su et al. 2019; Xu et al. 2018a; Takikawa et al. 2019; Zhu et al. 2021). A key challenge in these methods is how to efficiently fuse features from a boundary detection branch to improve semantic segmentation.

There are also works focusing on the control of information flow through boundaries (Bertasius et al. 2016; Ke et al. 2018; Bertasius et al. 2017; Ding et al. 2019; Chen et al. 2016; Borse et al. 2021). These methods usually learn pairwise pixel-level affinity to maintain the feature difference for pixels near semantic boundaries, while enhancing the similarity of features for interior pixels simultaneously.

The boundary details of the segmentation results can also be improved in post-refinement. DenseCRF (Krähenbühl et al. 2011) is often used to refine the segmentation results around boundaries. Segfix (Yuan et al. 2020b) trains a separate network to predict the correspondence between boundary and interior pixels. Thus, labels of interior pixels can be transferred to boundary pixels. Although these methods can efficiently refine most boundaries, they fail to model the relationship of pixels inside thin objects that contain a small number of interior pixels, which may downgrade the quality of slender object boundaries, as shown in Fig. 1. In contrast, the ABL encourages the alignment of PDBs and GTBs. Our experiment shows that it can handle such boundaries well.

The uniqueness of our ABL is that it allows propagating the GTB information with a distance transform for regulating the network behavior at the PDBs, while the network structure can remain the same. As a loss, ABL can save efforts in network design. Kervadec et al. (2019) proposed Boundary Loss (BL) for image segmentation, which is most related to our work. However, this loss is designed for unbalanced binary segmentation and actually a regional IoU loss. In our implementation, the ABL is coupled with an IoU loss in (Berman et al. 2018) to further refine the boundary details.

Active Boundary Loss

The ABL continuously monitors the changes on the PDBs in segmentation results to determine the plausible moving directions. Its computation is divided into two phases. First, for each pixel i on the PDBs, we determine its next candidate boundary pixel j closer to the GTBs in accordance with the relative location between the PDBs and GTBs. Second, we use the KL divergence as logits to encourage the increase in KL divergence between the class probability distribution of i and j . Meanwhile, this process reduces the KL divergence between i and the rest of its neighboring pixels. In this way, the PDBs can be gradually pushed towards the GTBs. Unfortunately, candidate boundary pixel conflicts might occur, severely degrading the performance of the ABL. Thus, we carefully reduce the conflicts through gradient flow control in the computation of ABL, which is crucial to its success.

The overall pipeline of ABL is illustrated in Fig. 2. Each phase and how to suppress the conflicts are detailed as follows. Hereafter, we use \mathbf{A}_i to denote the value stored at pixel i of a map \mathbf{A} .

Phase I. This phase starts with detecting the PDBs using the class probability map $\mathbf{P} \in \mathbb{R}^{C \times H \times W}$ output by the current network, where C denotes the number of semantic classes, and the image resolution is $H \times W$. Specifically, we compute a boundary map \mathbf{B} through the computation of KL-divergence to indicate the locations of PDBs. For a pixel i in \mathbf{B} , its value \mathbf{B}_i is computed as follows:

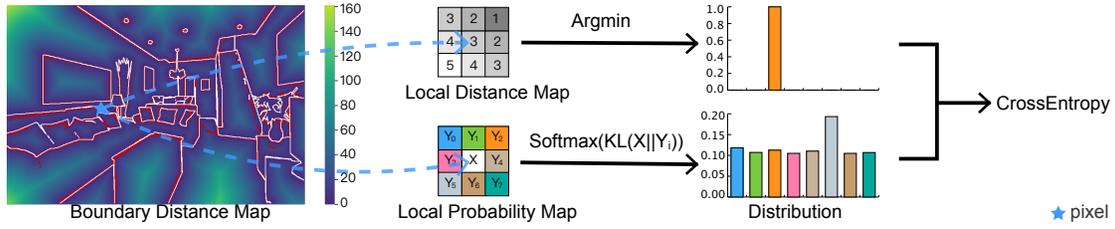


Figure 2: Pipeline of ABL. Boundary distance map is obtained via the distance transform of GTBs, taken an image in the ADE20K (Zhou et al. 2017) dataset as an example. The overlaid white and red lines in the boundary distance map indicate the GTBs and PDBs, respectively. Local distance map: the number indicates the closest distance to the GTBs. Local probability map: \mathbf{X} and $\mathbf{Y}_i, i \in \{0, 1, \dots, 7\}$ denote the class probability distribution for these pixels.

$$\mathbf{B}_i = \begin{cases} 1 & \text{if } \exists \text{KL}(\mathbf{P}_i, \mathbf{P}_j) > \epsilon, j \in \mathcal{N}_2(i); \\ 0 & \text{Otherwise,} \end{cases} \quad (1)$$

where 1 indicates the existence of PDBs, and \mathbf{P}_i is the C -dimensional vector extracted from the probability map at pixel i . $\mathcal{N}_2(i)$ indicates the 2-neighborhood of pixel i . Specifically, the offsets of pixels in $\mathcal{N}_2(i)$ to pixel i are $\{\{1, 0\}, \{0, 1\}\}$. Since it's difficult to define a perfect fixed threshold ϵ to detect PDBs, we choose an adaptive threshold to ensure that the number of boundary pixels in \mathbf{B} is less than 1% of the total pixels of the input image, where 1% is a ratio to approximate the number of boundary pixels in an image. Empirically, we observe that setting ϵ in this adaptive way can largely avoid the emergence of excessive misleading pixels in \mathbf{B} far from the GTBs, especially in the early training period. Controlling boundary pixel number also helps to save the computational cost of ABL.

Subsequently, for a pixel i on PDBs, its next candidate boundary pixel j is selected as its neighboring pixel with the smallest distance value computed by the distance transform¹ of the GTBs. The GTBs are also determined using Eq. 1, but the KL divergence is replaced by checking whether the ground-truth class labels are equal between pixel i and $j \in \mathcal{N}_2(i)$. To represent the coordinate of pixel j in the computation of ABL, we convert it into an offset to pixel i and then encode it as a one-hot vector. Specifically, we compute a target direction map $\mathbf{D}^g \in \{0, 1\}^{8 \times H \times W}$, where the one-hot vector for a pixel i stored at \mathbf{D}_i^g is 8D, because we use 8-neighborhood in this operation. The formula to compute \mathbf{D}_i^g can be written as:

$$\mathbf{D}_i^g = \Phi(\arg \min_j \mathbf{M}_{i+\Delta_j}), \quad j \in \{0, 1, \dots, 7\}, \quad (2)$$

where Δ_j represents the j^{th} element in the set of directions $\Delta = \{\{1, 0\}, \{-1, 0\}, \{0, -1\}, \{0, 1\}, \{-1, 1\}, \{1, 1\}, \{-1, -1\}, \{1, -1\}\}$, and \mathbf{M} is the result of distance transform of GTBs. The pixel $i + \Delta_j$ with the smallest distance is selected as the next candidate boundary pixel. The function Φ converts index j into a one-hot vector. For instance, if $j = 1$, $\Phi(j)$ should be $\{0, 1, 0, 0, 0, 0, 0, 0\}$, which is similar to the direction representation used in Segfix. In implementation, we dilate \mathbf{B} with 1 pixel and perform this operation for all the pixels in dilated \mathbf{B} to accelerate the movement of the PDBs, since more pixels are covered.

¹ `scipy.ndimage.morphology.distance_transform_edt` is used in the implementation of distance transform.

Phase II. The 8D vector \mathbf{D}_i^g computed in Eq. 2 is set to be the target distribution in the cross-entropy loss. We aim to increase the KL divergence between the class probability distribution of i and j , and simultaneously reduce the KL divergence between i and the rest of its neighboring pixels. An 8D vector using the KL divergence between pixel i and its neighboring pixel j as logits, denoted by \mathbf{D}_i^p , is then computed as follows:

$$\mathbf{D}_i^p = \left\{ \frac{e^{\text{KL}(\mathbf{P}_i, \mathbf{P}_{i+\Delta_k})}}{\sum_{m=0}^7 e^{\text{KL}(\mathbf{P}_i, \mathbf{P}_{i+\Delta_m})}}, k \in \{0, 1, \dots, 7\} \right\}, \quad (3)$$

where KL indicates the function to compute the KL divergence using \mathbf{P}_i and $\mathbf{P}_{i+\Delta_k}$.

For those pixels on the PDBs, the ABL is computed as the weighted cross-entropy loss:

$$\text{ABL} = \frac{1}{N_b} \sum_i \Lambda(\mathbf{M}_i) \text{CE}(\mathbf{D}_i^p, \mathbf{D}_i^g). \quad (4)$$

The weight function Λ is computed as $\Lambda(x) = \frac{\min(x, \theta)}{\theta}$, where N_b is the number of pixels on the PDBs and θ is a hyper-parameter set to 20. The closest distance to the GTBs at pixel i is used as a weight to penalize its deviation from the GTBs. If \mathbf{M}_i is 0, indicating that the pixel is already on the GTBs, this pixel will be discarded in the ABL.

Conflict suppression. Determining pixels on the PDBs using KL divergence might lead to the conflict case, as shown in Fig. 3. In this case, pixels V_1 and V_2 are deemed to be on a PDB (indicated by the red curve) because the KL divergence values computed for (V_1, W_1) and (V_2, V_3) are larger than the threshold. However, the GTB (indicated by the green curve) leads to the conflict when computing the ABL for V_1 and V_2 because the GTB is to the right of V_1 and V_2 . Thus, for pixel V_1 , we need to *increase* $\text{KL}(\mathbf{P}_{V_1}, \mathbf{P}_{V_2})$ because V_2 is the closest to the GTB and it should be the next candidate pixel in the neighborhood of V_1 . In contrast, for pixel V_2 , we need to *decrease* $\text{KL}(\mathbf{P}_{V_2}, \mathbf{P}_{V_1})$ because pixel V_3 is the next candidate boundary pixel for V_2 rather than V_1 . Thus, the gradients of the ABL computed for \mathbf{P}_{V_1} and \mathbf{P}_{V_2} might contradict with each other.

While it might be possible to design a global search algorithm to remove such kind of conflicts, it will significantly slow down the training. Thus, we choose to suppress the conflicts through the easy-to-implement detaching operation in Pytorch. Specifically, through the detaching operation, the gradient of ABL is computed only for the pixels on

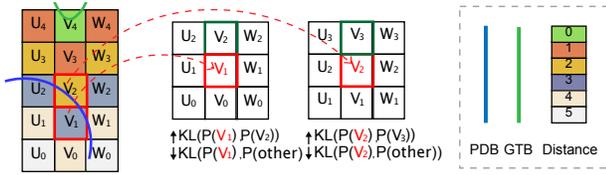


Figure 3: An example of conflict. V_4 : pixel on a GTB. \uparrow : increase. \downarrow : decrease. The KL divergence between V_1 and V_2 is required to increase for V_1 but to decrease for V_2 , resulting in contradictory gradients for V_1 and V_2 .

the PDBs, but not for its neighboring pixels. This process indicates that, for a 3×3 patch, we focus on the adjustment of the class probability distribution of pixels on the PDBs only so as to move the PDBs towards the GTBs. As a result, the conflicting gradient flow from $\text{KL}(\mathbf{P}_{V_2}, \mathbf{P}_{V_1})$ to \mathbf{P}_{V_1} is blocked in this case, and vice versa. Empirically, we found the mIoU drops around 3% without the detaching operation.

Furthermore, we use label smoothing (Szegedy et al. 2016) to regularize the ABL by setting the largest probability of the one-hot target probability distribution to 0.8 and the rest to 0.2/7 (the parameters, 0.8 and 0.2/7 are determined through experiments). This process can avoid overconfident decisions of network parameter updating, especially when there exist several pixels with the same distance value in the neighborhood of pixels on the PDBs. The detaching operation is also beneficial in this case to avoid conflicts in the gradient flow.

Training Loss

The training loss \mathcal{L}_t we mainly used to train a semantic segmentation network consists of three terms:

$$\mathcal{L}_t = \text{CE} + \text{IoU} + w_a \text{ABL}, \quad (5)$$

where **CE** is the most commonly used cross-entropy (CE) loss, which focuses on the per-pixel classification. The combination of lovász-softmax loss, namely **IoU**, and our **ABL** are two loss terms that are added to improve the boundary details, and w_a is a weight. The lovász-softmax loss is expressed as follows (Berman et al. 2018):

$$\text{IoU} = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \overline{\Delta_{J_c}}(\mathbf{m}(c)), \quad (6)$$

where \mathcal{C} is the number of classes, and $\mathbf{m}(c)$ is the vector of prediction errors for class $c \in \mathcal{C}$. $\overline{\Delta_{J_c}}$ indicates the lovász extension of the Jaccard loss Δ_{J_c} .

The reason for introducing the lovász-softmax loss is twofold: 1) This loss tends to prevent small objects from being ignored in segmentation such that the **ABL** can be used to improve their boundary details, since the **ABL** relies on the existence of predicted boundaries as the beginning step of its computation. 2) It can balance with the noisy predicted boundary pixels, especially at the early training period. The improvement of **ABL** over **CE** plus **IoU** is verified in the Experiments section.

Experiments

We implemented the ABL on a GPU server (2 Intel Xeon Gold 6148 CPUs, 512GB memory) with 4 Nvidia Tesla V100 GPUs. In this section, we report ablation studies, quantitative and qualitative results obtained from the evaluation of the ABL in image segmentation experiments and a test of fine-tuning VOS network.

Baselines. We use the OCR network (Yuan et al. 2020a) [backbone: HRNetV2-W48 (Wang et al. 2019)], DeeplabV3 (Chen et al. 2017a) [backbone: ResNet-50 (He et al. 2016)], and UperNet (Xiao et al. 2018) [backbone: SwinTransformer (Liu et al. 2021)] as the baseline models for the task of semantic image segmentation. To verify that our ABL can be applied to the task of video object segmentation, we use STM (Oh et al. 2019) as the baseline, since its pre-trained model is publicly available.

Dataset. We evaluate our loss mainly on the image segmentation dataset Cityscapes (Cordts et al. 2016) and ADE20K (Zhou et al. 2017). These two datasets provide densely annotated images that are important for the training of our method to align semantic boundaries. Cityscapes dataset contains high-quality dense annotations of 5000 images with 19 object classes, and ADE20K is a more challenging dataset with 150 object classes. There are 20210/2000/3000 images for the training/validation/testing set in ADE20K, respectively. Following the training protocol of (Yuan et al. 2020a), we use random crop, scaling (from 0.5 to 2), left-right flipping and brightness jittering between -10 and 10 degrees in data augmentation. In multi-scale inference, we apply scales $\{0.5, 0.75, 1.0, 1.25, 1.5, 1.75, 2.0\}$ and $\{0.5, 0.75, 1.0, 1.25, 1.5, 1.75\}$ as well as their mirrors.

Training parameters. We use stochastic gradient descent as the optimizer and utilize a “poly” learning rate policy similar to Chen et al. (2017b) in the training. Hence, the initial learning rate is multiplied by $(1 - \frac{iter}{maxiter})^{power}$ with power = 0.9. Sync Batch Normalization (Zhang et al. 2018) is used in all our experiments to improve stability. The detailed training and testing settings are as follows:

- ADE20K: the parameters are set as follows: initial learning rate = 0.02, weight decay = 0.0001, crop size = 520×520 , batch size = 16, and 150k training iterations, which are the same as the setting in Yuan et al. (2020a).
- Cityscapes: the parameters are set as follows: initial learning rate = 0.01 or 0.04, crop size = 512×1024 (used in OCR) or 769×769 (used in DeeplabV3), weight decay = 0.0005, batch size = 8, and 80k training iterations. The parameter setting is the same as the setting in Yuan et al. (2020b), but we do not use coarse data in experiments.

Evaluation Metrics. Three metrics, *i.e.* pixel accuracy (pix-Acc), mean Intersection-over-Union (mIoU), and boundary F-score (Perazzi et al. 2016a; Yuan et al. 2020b), are used to demonstrate the performance of the ABL. The first two metrics are used to evaluate the pixel-level and region-level accuracy of a segmentation result, respectively. Boundary F-scores are used to measure the quality of boundary alignment and computed within the area of the dilated GTBs. The dilation parameters are set to 1, 3, 5 pixels in our implementation. To better preserve boundary details in the evaluation,

loss	mIoU	pixAcc
CE	79.5	96.3
CE + ABL(20%)	79.6	96.3
CE + IoU	80.2	96.3
CE + IoU + BL	80.2	96.4
CE + IABL	80.5	96.4
CE + IABL w/o detach	76.0	95.6

Table 1: The influence of loss terms on Cityscapes dataset. These experiments are conducted using DeeplabV3 network and single-scale inference. ABL(20%): addin additional ABL in the last 20% epochs.

Method	loss	mIoU		pixAcc	
		SS	MS	SS	MS
OCR (2020a) [HRNetV2 -W48 (2019)]	CE	44.51	45.66	81.66	82.20
	CE + IoU	44.73	46.54	81.52	82.29
	CE + IABL	45.38	46.88	81.63	82.43
	CE + IFKL	45.11	45.96	81.61	82.28
UperNet (2018) [Swin-T (2021)]	CE	44.51	45.81	81.09	81.96
	CE + IoU	45.39	47.15	81.20	82.22
	CE + IABL	45.98	47.58	81.34	82.39
	CE + IoU + BL	45.72	47.50	81.33	82.39

Table 2: The influence of loss terms on ADE20K dataset. SS: single-scale inference. MS: multi-scale inference.

we do not use resize operation in the testing.

Combination of loss terms. To ease the description of the ablation study, we denote different combinations of loss terms used in the training as follows: CE = cross-entropy; CE+IoU = cross-entropy + lovász-softmax; CE+IABL = cross-entropy + lovász-softmax + ABL. w_a is set to 1.0 for ADE20K dataset but 1.5 for Cityscapes dataset, since training images’ resolution is much larger for Cityscapes.

In addition, we rely on the KL divergence of the class probability distributions of adjacent pixels, which can be viewed as the pair-wise term used in condition random field (Lafferty et al. 2001). Hence, it is necessary to verify how simply enforcing the KL divergence loss at each edge of an image works in the image segmentation, *i.e.* enforcing the loss for each edge between a pair of adjacent pixels, not only at semantic boundaries. To this end, we define a full KL-divergence (FKL) loss as follows:

$$\text{FKL} = \frac{1}{N_e} \sum_e \text{BCE}\left(\frac{1}{1+e^{\text{KL}(\mathbf{P}_{e_i}, \mathbf{P}_{e_j})}}, (\mathbf{G}_{e_i} \neq \mathbf{G}_{e_j})\right), \quad (7)$$

where e denotes an image edge that connects a pair of pixels e_i and e_j , N_e is the total number of edges in an image, and $(\mathbf{G}_{e_i} \neq \mathbf{G}_{e_j})$ returns 1 if the ground-truth label of pixel e_i is not equal to the label of e_j , otherwise 0. If the **FKL** loss is used with cross-entropy and lovász-softmax loss in the training, we denote this combination as CE+IFKL.

Ablation Studies

Loss terms. We first test the influence of loss terms on the Cityscapes validation dataset by re-training the DeepLabV3 network and show the results in Tab. 1. Since the gradient

Method	Backbone	OS	mIoU	pixAcc
CPN (2020)	ResNet-101+CPL	8×	46.27	81.85
PyConv (2020)	PyConvResNet-101	8×	44.58	81.77
DNL (2020)	HRNetV2-W48	4×	45.82	-
OCR (2020a)	HRNetV2-W48	4×	45.66	82.20
OCR+IABL	HRNetV2-W48	4×	46.88	82.43
UperNet (2018)	Swin-B (2021)	4×	51.66	84.06
UperNet+IABL	Swin-B	4×	52.40	84.11

Table 3: Results on ADE20K validation set. OS: Output stride. All results are obtained using multi-scale inference.

Method	Backbone	OS	mIoU
GSCNN (2019)	WideResNet-38+ASPP	8×	80.8
DANet (2019a)	ResNet-101+MG	8×	81.5
ACNet (2019b)	ResNet-101+MG	4×	82.0
OCR (2020a)	HRNetV2-W48	4×	82.2
OCR+IABL	HRNetV2-W48	4×	82.9

Table 4: Results on Cityscapes validation set. OS: Output stride. All results are obtained using multi-scale inference.

of ABL is not useful when PDBs are far from the GTBs, adding ABL at the beginning of the training does not improve network performance. Thus, we start to add ABL at the last 20% epochs to verify its effect, but only obtain a 0.1% improvement over mIoU. Then, we re-train the network with CE+IoU and CE+IABL. It shows that adding ABL to CE+IoU in training can increase the mIoU by 0.3%, and the combination of IoU loss and ABL, *i.e.* CE+IABL contribute 1% improvement on mIoU in this study. Although the ABL does not contribute most to mIoU in this case, we do see the obvious improvement of boundary alignment in qualitative comparisons. In Tab. 2, we test the contribution of each loss term on ADE20K dataset by re-training the OCR network. While the mIoU and pixel accuracy can both be improved after adding IoU loss and ABL, CE+IABL contributes most of the improvement to mIoU by around 0.65% over CE+IoU in the single-scale inference, which verifies the contribution of the proposed ABL in this experiment. We argue that the ABL can contribute more to a dataset with a large number of semantic classes and hence more GTBs. For instance, ADE20K has 150 classes, while Cityscapes only has 19 classes. More GTBs give the ABL more space to adjust the network’s behavior.

Detaching operation. We verify the effectiveness of the detaching operation in Tab. 1. Significant drops of pixel accuracy and mIoU can be observed when training without the aforementioned detaching operation to suppress conflicts. Hence, it is important to control the gradient flow when there exist contradictory targets for KL divergence between two neighboring pixels.

FKL loss. In Tabs. 2 and 7, it can be seen that the combination of IoU loss and FKL, denoted by IFKL in the 3rd row, can also improve the pixel accuracy and mIoU quantitatively. However, CE+IFKL does *not* perform as well as

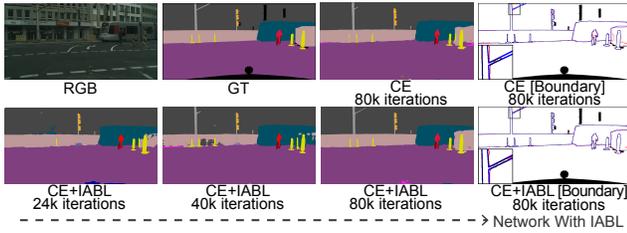


Figure 4: Progressive refinement of boundary details in the training. Dataset: Cityscapes. Network: DeepLabV3. The input image is taken from the Cityscapes training set as an example. The GTBs are in blue and PDBs in red.

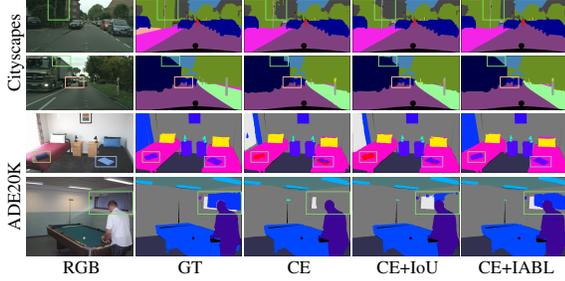


Figure 5: Qualitative results taken from Cityscapes and ADE20K validation set.

CE+IABL. We speculate that it is because FKL treats every pixel equally, while the ABL pays more attention to the pixels on the PDB. Such design allows the network to adjust its behavior in a progressive way, avoiding over-confident decisions when updating the network parameters.

The degree of ABL’s dependence on IoU loss. To evaluate ABL’s contribution further, we design an *IoU weight decay experiment*, which linearly decreases the weight of IoU loss from 1 to 0 during training but increase the weight of ABL from 0 to 1. It achieves mIoU 75.65% on the Cityscapes validation set with the FCN [backbone: HRNetV2-W18s], comparable to the mIoU 75.59% trained with CE+IoU+ABL without weight decay. It can be seen that the decreased IoU weight does not lead to the downgrade of segmentation performance. Moreover, we do observe that ABL can refine the semantic boundary for thin structures and complex boundaries, as shown in Tab. 6 and Figs. 4–6.

Quantitative Evaluation

Results on ADE20K and Cityscapes validation sets. In Tabs. 3 and 4, we show that training with IABL along with cross-entropy loss can improve the pixel accuracy and mIoU over state-of-the-art image segmentation networks. For the ADE20K dataset, training the OCR network with additional IABL improves the mIoU and pixel accuracy by 1.22% and 0.23% over that trained with cross-entropy only on the validation set. Not only effective on CNN-based network, additional IABL supervision on Transformer-based network UperNet [backbone: SwinTransformer-B] also brings an improvement in the mIoU by 0.74% on the validation set,

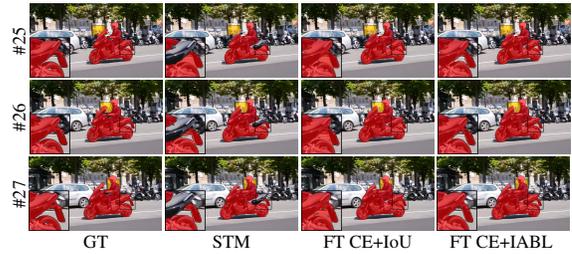


Figure 6: Qualitative results taken from DAVIS-2016 validation set. VOS Network: STM. FT: STM fine-tuning. Video: scooter-black (43 frames in total). #N: frame number.

which ranks the first place in the Tab. 3. Similarly, for the Cityscape dataset, training the OCR network with IABL improves the mIoU by 0.7% on the validation set.

Comparison with Segfix. We compare our method with Segfix (Yuan et al. 2020b) on the Cityscapes validation set by using mIoU and boundary F-score metrics, since both methods focus on improving boundary details in semantic segmentation. In Tab. 5, our method achieves comparable performance when using DeepLabV3 as the segmentation network, but improves the mIoU over Segfix by 0.3% when using the OCR network. In Tab. 6, we show the class-wise boundary F-scores of Segfix and our method. The scores are computed using the GTB dilation parameters 1, 3, and 5 pixels. While Segfix outperforms in the cases of 1 pixel, our method achieves a higher score in the cases of 3, 5 pixels.

Segfix is an elegant boundary refinement solution that propagates the interior labels to class boundaries. However, the propagation operation might downgrade the segmentation performance for thin objects that contain a small number of interior pixels. In contrast, the ABL is an end-to-end training loss that encourages the alignment of PDBs and GTBs, which achieves better mIoU and boundary F-scores, even for thin structures. Taking the class of traffic light as an example (Tab 6, 9th column), our method achieves a consistent improvement of the boundary F-score over Segfix in all parameter settings, which shows that our method can handle boundaries of thin objects well. Since Segfix is a post-processing method, it can also be used to improve the segmentation results of our method.

Comparison with Boundary Loss. We train ABL + generalized Dice loss (GDL) (Sudre et al. 2017) on the white matter hyperintensities (WMH) dataset with the same network architecture and training parameters as (Kervadec et al. 2019) for a fair comparison. We also use the same loss conjunction method: $\text{Loss} = \alpha * \text{GDL} + (1 - \alpha) * \text{ABL}$, where α linearly decreases from 1 to 0.01. In Tab. 8, we show that training with ABL + GDL achieves higher dice similarity coefficient (DSC) and smaller Hausdorff distance (HD) than Boundary Loss (BL) + GDL. Moreover, we extend BL to a multiple-class loss and make a further comparison on Cityscapes validation set. In Tabs. 1 and 2, IoU+ABL achieves higher mIoU than IoU+BL.

The motivation of BL is to minimize the distance between GTBs and PDBs. With the geo-cuts optimization techniques

method	road	walk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrain	sky	person	rider	car	truck	bus	train	motorcycle	bicycle	mean
DLV3	98.4	86.5	93.1	63.9	62.6	66.1	72.2	80.0	92.8	66.3	95.0	83.3	65.5	95.3	74.5	89.0	80.0	67.4	78.4	79.5
+Segfix	98.5	87.1	93.5	64.6	63.1	69.0	74.9	82.4	93.2	66.7	95.3	84.9	66.9	95.8	75.0	89.6	80.7	68.4	79.7	80.5 (\uparrow 1.0)
+IABL	98.1	84.9	93.1	59.9	63.1	68.3	75.4	82.3	92.7	64.6	95.0	84.6	69.4	95.6	79.9	87.8	83.3	70.6	80.2	80.5 (\uparrow 1.0)
OCR	98.4	86.8	93.3	62.2	66.0	70.0	73.9	82.0	93.0	67.2	95.0	84.3	66.3	95.6	82.9	91.7	85.0	67.7	79.2	81.1
+Segfix	98.5	87.3	93.5	62.6	66.4	71.4	75.7	83.3	93.3	67.6	95.3	85.2	67.2	96.0	83.3	92.2	85.5	68.6	80.0	81.7 (\uparrow 0.6)
+IABL	98.3	86.7	93.6	63.9	64.9	70.2	76.9	83.2	93.2	69.0	95.3	85.6	70.7	96.0	80.0	93.3	86.6	69.3	80.8	82.0 (\uparrow 0.9)

Table 5: Class-wise mIoU results obtained using single-scale inference on the Cityscapes validation set. DLV3: DeepLabV3. +Segfix: Segfix is used to refine the baseline output. +IABL: the network is trained with additional loss IABL.

	scale	road	walk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrain	sky	person	rider	car	truck	bus	train	motorcycle	bicycle	mean
1px	OCR	74.1	50.2	57.2	56.8	54.6	61.0	64.8	60.6	55.6	51.3	65.7	56.3	65.9	64.7	84.3	89.8	96.8	77.3	56.1	65.4
	+Segfix	76.0	52.6	59.3	58.1	55.5	64.2	67.8	64.1	57.7	52.9	67.1	59.3	67.6	68.7	84.7	89.8	97.0	78.4	58.6	67.3 (\uparrow 1.9)
	+IABL	74.6	51.0	57.4	53.7	50.8	62.1	74.7	65.1	55.7	52.3	66.8	60.1	68.9	67.0	84.6	90.3	96.4	79.5	61.5	67.0 (\uparrow 1.6)
3px	OCR	86.5	70.1	75.7	62.5	60.1	79.6	77.8	78.9	76.4	58.6	82.9	73.9	76.7	84.6	86.5	92.8	97.4	80.4	71.3	77.5
	+Segfix	87.2	71.0	76.4	63.0	60.7	79.7	78.5	79.3	77.3	60.0	83.5	74.5	77.4	85.9	86.6	92.3	97.5	81.1	72.3	78.1 (\uparrow 0.6)
	+IABL	86.2	70.1	75.4	59.4	56.3	80.7	86.9	81.9	76.5	60.1	84.1	77.6	80.4	86.1	87.2	93.2	97.0	83.1	76.9	78.9 (\uparrow 1.4)
5px	OCR	90.3	76.4	82.3	65.0	62.9	82.8	81.7	82.6	84.4	62.1	88.1	78.9	80.8	89.5	87.5	93.7	97.7	81.9	77.7	81.4
	+Segfix	90.6	76.9	82.5	65.1	63.0	82.6	81.8	82.3	84.6	63.1	88.2	78.7	81.1	90.1	87.3	93.1	97.7	82.3	77.8	81.5 (\uparrow 0.1)
	+IABL	89.8	76.4	81.8	61.7	58.6	83.8	89.9	84.9	84.3	63.3	89.1	82.2	84.1	90.6	88.1	94.0	97.2	84.6	82.5	82.5 (\uparrow 1.1)

Table 6: Class-wise Boundary F-score results obtained using multi-scale inference on the Cityscapes validation set.

Fine-tuning	No	YES	YES	YES	YES
Loss	CE	CE	CE+IoU	CE+IFKL	CE+IABL
\mathcal{J} -mean	88.67	88.81	89.08	89.08	89.29
\mathcal{F} -mean	89.86	90.25	90.66	90.63	90.82

Table 7: VOS results on DAVIS-2016 validation set.

(Boykov et al. 2006), this problem is converted to minimize the regional integral. This behavior will weaken the influence of pixels near GTBs since the distance weights there are much smaller, and the ratio of these pixels is small compared to the image size. In contrast, ABL focuses on PDB pixels, which can achieve better alignment.

VOS results. We fine-tune the state-of-the-art VOS network STM (Oh et al. 2019) with our loss to verify that our method can also be applied to VOS. Specifically, the STM is fine-tuned for 1k iterations with batch size 4 on both DAVIS-2016 (Perazzi et al. 2016b) and YouTube-VOS (Xu et al. 2018b) training data. The learning rate is set to $5e-8$, and the weight w_a is set to 5.0. In Tab. 7, it can be seen that fine-tuning with addition IABL can improve region similarity metric \mathcal{J} -mean and contour accuracy \mathcal{F} -mean by around 0.7% and 1%, respectively, when testing on Davis-2016 validation set. The definition of \mathcal{J} -mean and \mathcal{F} -mean can be found in the DAVIS-2016 dataset. Similar to image segmentation, training with CE+IABL can improve over CE+IoU loss, which also verifies ABL’s contribution in VOS. However, adding FKL loss does not show superior performance.

Qualitative Results

Fig. 4 illustrates the progressive refinement of boundary details when using IABL as the additional training loss.

	Loss	GDL	GDL+BL	GDL+ABL
DSC		0.727	0.748	0.768
HD(mm)		1.045	0.987	0.980

Table 8: BL vs. ABL on WMH validation set.

This result is obtained when training DeepLabV3 on the Cityscapes dataset. It can be seen that the PDBs (red lines) of the traffic light and other objects are pushed toward the GTBs (blue lines). In Figs. 1 and 5, we show how adding loss terms influences the quality of semantic boundaries. The results show that the proposed ABL can greatly improve the semantic boundary details. Fig. 6 illustrates the improved boundary details when fine-tuning STM with additional IABL. It also shows that fine-tuning with CE+IABL can further improve the boundary details over CE+IoU, such as the tail of motorcycle.

Conclusion

In this work, we proposed an active boundary loss to be used in the end-to-end training of segmentation networks. Its advantage is that it allows the propagation of the ground-truth boundary information using a distance transform so as to regulate the network behavior at predicted boundaries. We have demonstrated that integrating the ABL into the network training can substantially improve the boundary details in semantic segmentation.

In the future, it would be interesting to investigate how to reduce conflicts in our loss to further control the network behavior around boundaries efficiently. In addition, we plan to explore how to design boundary-aware loss to improve the boundary details in the task of depth prediction.

Acknowledgments

We thank the reviewers for their constructive comments. Weiwei Xu is partially supported by NSFC (No. 61732016).

References

- Berman, M.; Triki, A. R.; and Blaschko, M. B. 2018. The Lovász-Softmax Loss: A Tractable Surrogate for the Optimization of the Intersection-Over-Union Measure in Neural Networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 4413–4421.
- Bertasius, G.; Shi, J.; and Torresani, L. 2016. Semantic segmentation with boundary neural fields. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 3602–3610.
- Bertasius, G.; Torresani, L.; Yu, S. X.; and Shi, J. 2017. Convolutional random walk networks for semantic image segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 858–866.
- Borse, S.; Wang, Y.; Zhang, Y.; and Porikli, F. 2021. Inverse-Form: A Loss Function for Structured Boundary-Aware Segmentation. In *CVPR*, 5901–5911.
- Boykov, Y.; Kolmogorov, V.; Cremers, D.; and DeLong, A. 2006. An integral solution to surface evolution PDEs via geo-cuts. In *Eur. Conf. Comput. Vis.*, 409–422. Springer.
- Caelles, S.; Maninis, K.-K.; Pont-Tuset, J.; Leal-Taixé, L.; Cremers, D.; and Van Gool, L. 2017. One-shot video object segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 221–230.
- Chen, L.; Papandreou, G.; Schroff, F.; and Adam, H. 2017a. Rethinking Atrous Convolution for Semantic Image Segmentation. arXiv:1706.05587.
- Chen, L.-C.; Barron, J. T.; Papandreou, G.; Murphy, K.; and Yuille, A. L. 2016. Semantic image segmentation with task-specific edge detection using cnns and a discriminatively trained domain transform. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 4545–4554.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2017b. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4): 834–848.
- Chen, X.; Lian, Y.; Jiao, L.; Wang, H.; Gao, Y.; and Shi, L. 2020. Supervised Edge Attention Network for Accurate Image Instance Segmentation. In Vedaldi, A.; Bischof, H.; Brox, T.; and Frahm, J., eds., *Eur. Conf. Comput. Vis.*, volume 12372 of *Lecture Notes in Computer Science*, 617–631. Springer.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The cityscapes dataset for semantic urban scene understanding. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 3213–3223.
- Ding, H.; Jiang, X.; Liu, A. Q.; Thalmann, N. M.; and Wang, G. 2019. Boundary-aware feature propagation for scene segmentation. In *Int. Conf. Comput. Vis.*, 6819–6829.
- Ding, H.; Jiang, X.; Shuai, B.; Qun Liu, A.; and Wang, G. 2018. Context contrasted feature and gated multi-scale aggregation for scene segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2393–2402.
- Duta, I. C.; Liu, L.; Zhu, F.; and Shao, L. 2020. Pyramidal Convolution: Rethinking Convolutional Neural Networks for Visual Recognition. arXiv:2006.11538.
- Fu, J.; Liu, J.; Tian, H.; Li, Y.; Bao, Y.; Fang, Z.; and Lu, H. 2019a. Dual attention network for scene segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 3146–3154.
- Fu, J.; Liu, J.; Wang, Y.; Li, Y.; Bao, Y.; Tang, J.; and Lu, H. 2019b. Adaptive context network for scene parsing. In *Int. Conf. Comput. Vis.*, 6748–6757.
- Gong, K.; Liang, X.; Li, Y.; Chen, Y.; Yang, M.; and Lin, L. 2018. Instance-level human parsing via part grouping network. In *Eur. Conf. Comput. Vis.*, 770–785.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 770–778.
- Hu, Y.-T.; Huang, J.-B.; and Schwing, A. 2017. Maskrcnn: Instance level video object segmentation. In *Adv. Neural Inform. Process. Syst.*, 325–334.
- Kass, M.; and Witkin, A. 1988. Snakes: Active contour models. *Int. J. Comput. Vis.*, 1: 321–331.
- Ke, T.-W.; Hwang, J.-J.; Liu, Z.; and Yu, S. X. 2018. Adaptive affinity fields for semantic segmentation. In *Eur. Conf. Comput. Vis.*, 587–602.
- Kervadec, H.; Bouchtiba, J.; Desrosiers, C.; Granger, E.; Dolz, J.; and Ayed, I. B. 2019. Boundary loss for highly unbalanced segmentation. In *International conference on medical imaging with deep learning*, 285–296.
- Krähenbühl, P.; and Koltun, V. 2011. Efficient Inference in Fully Connected CRFs with Gaussian Edge Potentials. In *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain*, 109–117.
- Lafferty, J. D.; McCallum, A.; and Pereira, F. C. 2001. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In *ICML*, 282–289.
- Li, Y.; Qi, H.; Dai, J.; Ji, X.; and Wei, Y. 2017. Fully convolutional instance-aware semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2359–2367.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 10012–10022.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 3431–3440.
- Minaee, S.; Boykov, Y. Y.; Porikli, F.; Plaza, A. J.; Kehtarnavaz, N.; and Terzopoulos, D. 2021. Image segmentation using deep learning: A survey. *PAMI*.
- Oh, S. W.; Lee, J.-Y.; Xu, N.; and Kim, S. J. 2019. Video object segmentation using space-time memory networks. In *Int. Conf. Comput. Vis.*, 9226–9235.

- Perazzi, F.; Pont-Tuset, J.; McWilliams, B.; Van Gool, L.; Gross, M.; and Sorkine-Hornung, A. 2016a. A benchmark dataset and evaluation methodology for video object segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 724–732.
- Perazzi, F.; Pont-Tuset, J.; McWilliams, B.; Van Gool, L.; Gross, M.; and Sorkine-Hornung, A. 2016b. A Benchmark Dataset and Evaluation Methodology for Video Object Segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241. Springer.
- Ruan, T.; Liu, T.; Huang, Z.; Wei, Y.; Wei, S.; and Zhao, Y. 2019. Devil in the details: Towards accurate single and multiple human parsing. In *AAAI*, volume 33, 4814–4821.
- Shin Yoon, J.; Rameau, F.; Kim, J.; Lee, S.; Shin, S.; and So Kweon, I. 2017. Pixel-level matching for video object segmentation using convolutional neural networks. In *Int. Conf. Comput. Vis.*, 2167–2176.
- Su, J.; Li, J.; Zhang, Y.; Xia, C.; and Tian, Y. 2019. Selectivity or invariance: Boundary-aware salient object detection. In *Int. Conf. Comput. Vis.*, 3799–3808.
- Sudre, C. H.; Li, W.; Vercauteren, T.; Ourselin, S.; and Cardoso, M. J. 2017. Generalised Dice Overlap as a Deep Learning Loss Function for Highly Unbalanced Segmentations. In *Deep learning in medical image analysis and multimodal learning for clinical decision support.*, volume 10553, 240–248. Springer, Cham.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2818–2826.
- Takikawa, T.; Acuna, D.; Jampani, V.; and Fidler, S. 2019. Gated-scnn: Gated shape cnns for semantic segmentation. In *Int. Conf. Comput. Vis.*, 5229–5238.
- Voigtlaender, P.; Chai, Y.; Schroff, F.; Adam, H.; Leibe, B.; and Chen, L.-C. 2019. Feelvos: Fast end-to-end embedding learning for video object segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 9481–9490.
- Wang, J.; Sun, K.; Cheng, T.; Jiang, B.; Deng, C.; Zhao, Y.; Liu, D.; Mu, Y.; Tan, M.; Wang, X.; Liu, W.; and Xiao, B. 2019. Deep High-Resolution Representation Learning for Visual Recognition. *TPAMI*.
- Xiao, T.; Liu, Y.; Zhou, B.; Jiang, Y.; and Sun, J. 2018. Unified Perceptual Parsing for Scene Understanding. In *Eur. Conf. Comput. Vis.* Springer.
- Xu, D.; Ouyang, W.; Wang, X.; and Sebe, N. 2018a. Padnet: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 675–684.
- Xu, N.; Yang, L.; Fan, Y.; Yang, J.; Yue, D.; Liang, Y.; Price, B.; Cohen, S.; and Huang, T. 2018b. Youtube-vos: Sequence-to-sequence video object segmentation. In *Eur. Conf. Comput. Vis.*, 585–601.
- Yin, M.; Yao, Z.; Cao, Y.; Li, X.; Zhang, Z.; Lin, S.; and Hu, H. 2020. Disentangled Non-local Neural Networks. In Vedaldi, A.; Bischof, H.; Brox, T.; and Frahm, J., eds., *Eur. Conf. Comput. Vis.*, volume 12360 of *Lecture Notes in Computer Science*, 191–207. Springer.
- Yu, C.; Wang, J.; Gao, C.; Yu, G.; Shen, C.; and Sang, N. 2020. Context Prior for Scene Segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 12416–12425.
- Yuan, Y.; Chen, X.; and Wang, J. 2020a. Object-Contextual Representations for Semantic Segmentation. In *ECCV*, volume 12351 of *Lecture Notes in Computer Science*, 173–190. Springer.
- Yuan, Y.; Xie, J.; Chen, X.; and Wang, J. 2020b. Segfix: Model-agnostic boundary refinement for segmentation. In *Eur. Conf. Comput. Vis.*, 489–506. Springer.
- Zhang, H.; Dana, K.; Shi, J.; Zhang, Z.; Wang, X.; Tyagi, A.; and Agrawal, A. 2018. Context encoding for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 7151–7160.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 633–641.
- Zhu, F.; Zhu, Y.; Zhang, L.; Wu, C.; Fu, Y.; and Li, M. 2021. A Unified Efficient Pyramid Transformer for Semantic Segmentation. In *Int. Conf. Comput. Vis.*, 2667–2677.
- Zhu, Y.; Sapra, K.; Reda, F. A.; Shih, K. J.; Newsam, S.; Tao, A.; and Catanzaro, B. 2019. Improving semantic segmentation via video propagation and label relaxation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 8856–8865.