

Mixup-CAM: Weakly-supervised Semantic Segmentation via Uncertainty Regularization

Yu-Ting Chang¹
ychang39@ucmerced.edu

Qiaosong Wang²
qiaowang@ebay.com

Wei-Chih Hung²
whung8@ucmerced.edu

Robinson Piramuthu²
robinpir@amazon.com

Yi-Hsuan Tsai³
wasidennis@gmail.com

Ming-Hsuan Yang¹⁴
mhyang@ucmerced.edu

¹ UC Merced

² eBay Inc.

³ NEC Labs America

⁴ Google Research

Abstract

Obtaining object response maps is one important step to achieve weakly-supervised semantic segmentation using image-level labels. However, existing methods rely on the classification task, which could result in a response map only attending on discriminative object regions as the network does not need to see the entire object for optimizing the classification loss. To tackle this issue, we propose a principled and end-to-end trainable framework to allow the network paying attention to other parts of the object, while producing a more complete and uniform response map. Specifically, we introduce the mixup data augmentation scheme into the classification network and design two uncertainty regularization terms to better interact with the mixup strategy. In experiments, we conduct extensive analysis to demonstrate the proposed method and show favorable performance against state-of-the-art approaches.

1 Introduction

Semantic segmentation is one of the fundamental tasks in computer vision, with a wide range of applications such as image editing and scene understanding. In order to obtain reliable models and achieve promising performance, recently deep neural network (DNN) based methods [6, 19, 47] are learned from fully-supervised data that requires pixel-wise semantic annotations. However, acquiring such pixel-wise annotations is usually time-consuming and labor-intensive, which limits the application potentials in the real world. As a result, numerous approaches tackle this issue via training models only on weakly-annotated data, *e.g.*, image-level [10, 11, 26, 53, 55], bounding box [9, 24, 52], point-level [2], scribble-based

[61, 42], or video-level [10, 40, 52] labels. In this paper, we focus on utilizing the image-level label, which is the most efficient scheme for weak annotations but also a challenging scenario.

Existing weakly-supervised semantic segmentation (WSSS) algorithms mainly operate with three steps: 1) localizing objects via a categorical response map, 2) refining the response map to generate pseudo annotations, and 3) training the semantic segmentation network using pseudo ground truths. Recent methods [11, 21, 43, 45] have achieved significant progress for WSSS, but most of them focus on improving the latter two steps. Since the success of these sequential steps hinges on the quality of the initial response map generated in the first step, in this paper we present an effective solution to localize objects.

One common practice to produce the initial response map is using class activation map (CAM) [53]. However, since CAM is typically supervised by a classification loss that could be sufficiently optimized through seeing only a small portion of objects, the generated response map usually only attends on partial regions (see Figure 1(a)). To tackle this issue, recent methods [23, 28] make efforts to improve the response map via using the dropout strategy that increases the model uncertainty or aggregating maps produced at different stages to see more object parts. However, there remains a challenge whether there are better loss function designed to explicitly facilitate the model training and produce better response maps, which are not addressed in prior works.

In this paper, we propose a principled and end-to-end trainable network with loss functions designed to systematically control the generation of the response map. First, inspired by the mixup data augmentation in [50], we observe that including mixup could effectively calibrate the model uncertainty on overconfident predictions [49] and in return enables the model to attend to more object regions. However, it is challenging to control the mixup augmentation process and the model uncertainty, due to non-uniform response distributions (see Figure 1(b)), which may affect subsequent response refinement steps. Therefore, we introduce another two loss terms to the mixup process by regularizing the class-wise uncertainty and the spatial response distribution. We refer to our model as *Mixup-CAM* and show that the produced response map is more complete and balanced across object regions (see Figure 1(c)), which facilitates the latter response refinement and segmentation model training steps.

We conduct quantitative and qualitative experiments to demonstrate the effectiveness of the proposed Mixup-CAM method on the PASCAL VOC 2012 dataset [100]. To the best of our knowledge, our algorithm is the first to demonstrate that mixup could improve the WSSS task on complicated multi-labeled images, along with other designed loss functions to produce better response maps. In addition, we present the ablation study and more analysis to validate the importance of each designed loss. Finally, we show that our method achieves state-of-the-art semantic segmentation performance against existing approaches.

2 Related Work

Initial Prediction for WSSS. Initial cues are essential for segmentation tasks since they provide reliable priors to generate segmentation maps. The class activation map (CAM) [53] is a common practice for localizing objects. It highlights class-specific regions that are served as the initial cues. However, since the CAM model is trained by a classification network, it tends to attend to small discriminative parts of the object, leading to incomplete initial masks. Several methods have been developed to alleviate this problem. Approaches like [29, 37, 44, 51] deliberately hide or erase the regions of an object, forcing the model

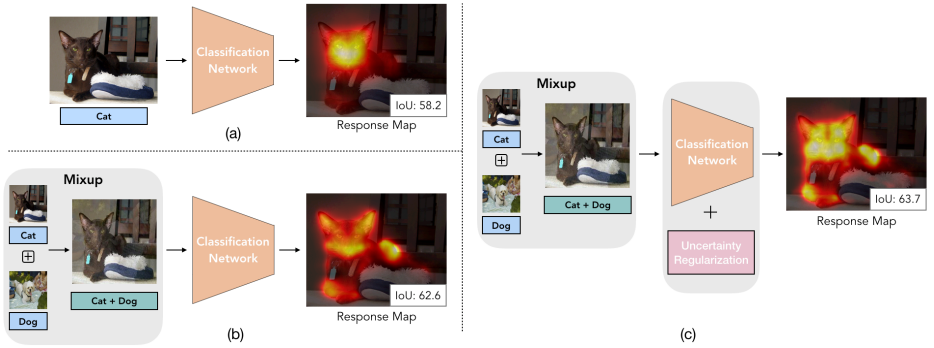


Figure 1: Comparisons of (a) the original CAM method; (b) CAM + the mixup data augmentation; and (c) the proposed Mixup-CAM framework that integrates the mixup scheme and the uncertainty regularization. Compared to (a) and (b), our final response map (c) attends to other object parts with more uniformly distributed response.

to look for more diverse parts. However, such strategies require iterative model training and response aggregation steps. After gradual expansion of the attention regions, non-object regions are prone to be activated, which leads to inaccurate attention maps. Other algorithms [20, 49] use both object and background cues to prevent the attention map from including more background regions, yet pixel-level saliency labels are used.

Instead of using the erasing scheme, recently the FickleNet approach [28] introduces the stochastic feature selection to obtain a diverse combination of locations on feature maps. Moreover, the OAA method [23] adopts an online attention accumulation strategy to collect various object parts discovered at different training stages. By aggregating the attention maps, it could obtain an initial cue that contains a larger region of the object. Unlike methods that mitigate the problem by discovering complementary regions via iterative erasing steps or consolidating attention maps, our proposed approach aims at harnessing the uncertainty in end-to-end classification learning. In addition, by regularizing both class-wise uncertainties and spatial response distributions, our approach averts the attention from focusing on small parts of the semantic objects, hence producing much improved response maps.

Response Refinement for WSSS. Various approaches [11, 12, 13, 21, 26, 43, 45] are proposed to refine the initial cue by expanding the region of attention maps. Other methods [11, 12, 13] are developed using affinity learning. The recent SSDD scheme [56] proposes a difference detection module to estimate and reduce the gap between the initial mask and final segmentation results in a self-supervised manner. However, the performance of these methods is limited as initial seeds are still obtained from CAM-like methods. If these seeds only come from the discriminative parts of the object, it is difficult to expand regions into non-discriminative parts. Moreover, the initial prediction may produce wrong attention regions, which would lead to even more inaccurate regions in subsequent refinement steps.

Label-preserving vs. Non-preserving Augmentations. Data augmentation is a common regularization technique in both supervised and unsupervised learning [8, 15, 30]. Conventional data augmentation techniques such as scaling, rotation, cropping, color augmentation, and Gaussian noise can change the pixel values of an image without altering its labels. These

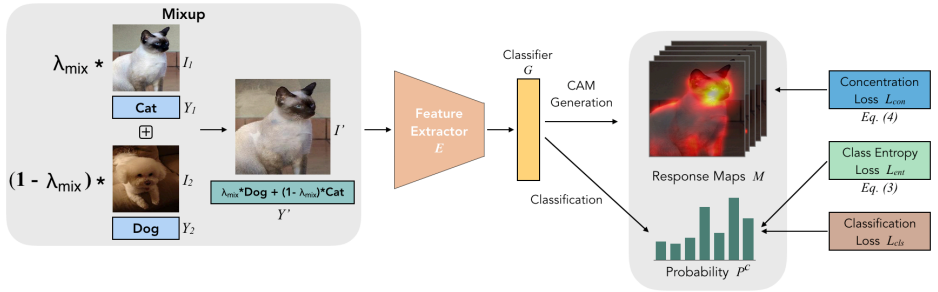


Figure 2: Overview of Mixup-CAM. We perform mixup data augmentation on input images with their corresponding labels via (2) and pass the mixed image through the feature extractor E and the classifier G to obtain the probability score P^c for each category c . For loss functions, in addition to the classification loss \mathcal{L}_{cls} on mixup samples, we design two terms to regularize class-wise entropy (\mathcal{L}_{ent} via (3)) and spatial distribution on CAM (\mathcal{L}_{con} via (4)).

label-preserving transformations are commonly applied in training deep neural networks to improve the model generalization capabilities.

Recent work has demonstrated that even non-label-preserving data augmentation can be surprisingly effective. Explicit label smoothing has been adopted successfully to improve the performance of deep neural models. The Mixup method [50] is proposed to train a neural network on a convex combination space of image pairs and their corresponding labels. It has been proven effective for the classification task and increases the robustness of neural networks. Numerous Mixup variants [8, 15, 38, 39, 40, 48] have been proposed to extend mixup for better prediction of uncertainty and calibration of the DNNs. These methods exhibit shared similarity of producing better-generalized models.

Entropy Regularization. Aside from mixup schemes for predictive uncertainty, another common uncertainty measure is entropy, which could act as a strong regularizer in both supervised and semi-supervised learning [14, 64]. In particular, [34] discourages the neural network from being over-confident by penalizing low-entropy distributions, while [14] utilizes entropy minimization in a semi-supervised setting as a training signal on unlabeled data. In this paper, we also adopt the entropy-based loss to regularize the uncertainty, coupled with the mixup data augmentation for producing better response maps on objects.

3 Weakly-supervised Semantic Segmentation

We first describe the overall algorithm and introduce details of the proposed Mixup-CAM framework with loss functions designed to improve the initial response map. We then detail how to generate the final semantic segmentation results.

3.1 Algorithm Overview

One typical way to generate response maps for annotated object categories is to use CAM [63]. However, these response maps tend to focus on discriminative object parts, which are less effective for the WSSS task. One reason is that CAM relies on the classification loss, which only requires partial object regions to be activated during training. As a result, when

the objective is already optimized with high confidence, the model may not attempt to learn other object parts.

In this paper, we propose to integrate the idea of mixup data augmentation [50], thereby calibrating the uncertainty in prediction [49] as well as allowing the model to attend to other regions of the image. Although we find that adding mixup could improve the response map, sometimes the response could diverge too much, resulting in more false-positive object regions. To further regularize this uncertainty, we introduce two additional loss terms: the spatial loss and the class-wise loss. We illustrate the overall model and loss designs in Figure 2 and provide more details in the next subsection.

After receiving the initial response map, we utilize the method in [10] to expand and refine the response. Finally, we generate pseudo ground truths from the refined response and train a semantic segmentation network to obtain the final segmentation output. Note that while we focus on the first step of the initial response map in this paper, the succeeding two steps could be replaced with alternative modules or models.

3.2 Mixup-CAM

CAM Generation. We first describe the CAM method for producing the initial response map as our baseline (see Figure 1(a)). The base network begins with a feature extractor E , followed by a global average pooling (GAP) layer and a fully-connected layer G as the output classifier. Next, given an input image I with its image-level labels Y , the network is trained with a multi-label classification loss $\mathcal{L}_{cls}(Y, G(E(I)))$ following [53]. After training this classification network, the activation map M^c for each category c is obtained by applying the c -channel classifier weight θ_G^c on the feature map $f = E(I)$:

$$M^c = \theta_G^{c\top} f. \quad (1)$$

Finally, the response is normalized by the maximum value of M^c .

Mixup Data Augmentation. Since the original classification network could easily obtain high confidence, in which the generated CAM only attends to small discriminative object parts, we utilize mixup data augmentation to calibrate the uncertainty in prediction [49]. Given an image pair $\{I_1, I_2\}$ and its label $\{Y_1, Y_2\}$ randomly sampled from the training set, we augment an image I' and its label Y' via:

$$\begin{aligned} I' &= \lambda_{mix} I_1 + (1 - \lambda_{mix}) I_2, \\ Y' &= \lambda_{mix} Y_1 + (1 - \lambda_{mix}) Y_2, \end{aligned} \quad (2)$$

where λ_{mix} is sampled from the $Beta(\alpha, \alpha)$ distribution following [50]. Using this augmented data, we feed it into the classification network to minimize the loss $\mathcal{L}_{cls}(Y', G(E(I')))$ and follow the same procedure in (1) to produce the response map (see Figure 1(b)).

Compared to the original CAM generation, our network no longer receives a pure image but a mixed image that could have multiple objects with their weights based on λ_{mix} as in (2). Therefore, the predictive uncertainty could be enhanced, leading to smoother output distributions and enforcing the model to pay attention to other regions in the image in order to satisfy the classification loss $\mathcal{L}_{cls}(Y', G(E(I')))$.

Uncertainty Regularization. Although mixup could improve the response map by looking at other parts in the image, sometimes the response could become too divergent and thus attend to pixels non-uniformly, *e.g.*, Figure 1(b). This is attributed to the difficulty in controlling the quality of mixed images, especially when the model faces more complicated images such as PASCAL VOC, *e.g.*, an object could appear at various locations of the image with noisy background clutters.

To further facilitate the mixup process, we propose to self-regularize the uncertainty via class-wise loss and spatial loss terms. The first term is to directly minimize the entropy in output prediction from the classifier to reduce uncertainty:

$$\mathcal{L}_{ent}(G(E(I'))) = -\frac{1}{HW} \sum_{h,w} \sum_{c \in C} P^c(h,w) \log P^c(h,w), \quad (3)$$

where C is the category number and $P^c \in \mathbb{R}^{H \times W}$ is the output probability for category c . Since our classifier G outputs multi-label probability, we concatenate the probabilities and normalize them by the maximum value, then calculate the final P with *softmax*.

Although the first term has the ability to minimize the uncertainty, it does not explicitly operate on the response map. To better balance the distribution on the response, we utilize a concentration loss similar to [22] and apply it directly on CAM for each category (*i.e.*, M^c), which encourages activated pixels to be spatially close to the response center:

$$\mathcal{L}_{con}(M) = \sum_{c \in \tilde{C}} \sum_{h,w} \|\langle h,w \rangle - \langle \mu_h^c, \mu_w^c \rangle\|^2 \cdot \hat{M}^c(h,w), \quad (4)$$

where $\mu_h^c = \sum_{h,w} h \cdot \hat{M}^c(h,w)$ is the center in height for category c (similarly for μ_w^c), \hat{M}^c is the normalized response of M^c to represent a spatially distributed probability map. Note that, here we only calculate the concentration loss on presented categories \tilde{C} as provided in the image-label Y to avoid confusing the model with invalid categories.

Overall Objective. We have described our proposed Mixup-CAM framework, including mixup data augmentation in (2) and two regularization terms, *i.e.*, (3) and (4). To train the entire model in an end-to-end fashion, we perform the online mixup procedure and jointly optimize the following loss functions:

$$\mathcal{L}_{all} = \mathcal{L}_{cls}(I', Y') + \lambda_{ent} \mathcal{L}_{ent}(I') + \lambda_{con} \mathcal{L}_{con}(M). \quad (5)$$

For simplicity, we omit the detailed notation inside each loss term. We also note that M is produced online via computing (1) on valid categories in each forward iteration.

3.3 Implementation Details

Classification Network. Similar to [10], we use the ResNet-38 architecture [46] as our classification network, which consists of 38 convolution layers with wide channels, followed by a 3×3 convolution layer with 512 channels for better adaptation to the classification task, a global average pooling layer, and two fully-connected layers for classification. In training, we adopt the pre-trained model on ImageNet [10] and finetune it on the PASCAL VOC 2012 dataset. Typical label-preserving data augmentations, *i.e.*, horizontal flip, random cropping, random scaling, and color jittering, are utilized on the training set.

Table 1: IoU results of CAM and its refinement on the PASCAL VOC training set.

Method	CAM	CAM + Refinement
AffinityNet [10]	48.0	58.1
Mixup \mathcal{L}_{cls}	49.3	60.5
Mixup $\mathcal{L}_{cls} + \mathcal{L}_{ent}$	49.5	61.6
Mixup $\mathcal{L}_{cls} + \mathcal{L}_{con}$	49.9	61.7
Mixup $\mathcal{L}_{cls} + \mathcal{L}_{ent} + \mathcal{L}_{con}$	50.1	61.9

We implement the proposed Mixup-CAM framework using PyTorch with a Titan X GPU with 12 GB memory. For training the classification network, we use the Adam optimizer [25] with initial learning rate of $1e-3$ and the weight decay of $5e-4$. For mixup, we use $\alpha = 0.2$ in the $Beta(\alpha, \alpha)$ distribution. For uncertainty regularization, we set λ_{ent} as 0.02 and λ_{con} as $2e-4$. Unless specified otherwise, we use the same parameters in all the experiments. In the experimental section, we show studies for the sensitivity of different parameters.

Semantic Segmentation Generation. Based on the response map generated by our Mixup-CAM, we adopt the random walk approach via affinity [10] to refine the response and produce pixel-wise pseudo ground truths for semantic segmentation. In addition, similar to existing methods, we adopt dense conditional random fields (CRF) [27] to further refine the response and obtain better object boundaries. Finally, we utilize the Deeplab-v2 framework [6] with the ResNet-101 architecture [27] and train the segmentation network.

4 Experimental Results

In this section, we present our main results of the proposed Mixup-CAM method for the WSSS task. First, we show that our approach achieves better initial response maps and further improves the subsequent refinement step. Second, we demonstrate the importance of each designed component. Finally, we provide evaluations on final semantic segmentation outputs in the PASCAL VOC dataset [10] against the state-of-the-art approaches. More results can be found in the supplementary material.

4.1 Evaluated Dataset and Metric

We conduct experiments on the PASCAL VOC 2012 semantic segmentation benchmark [10] with 21 categories, including one background class. Following existing WSSS methods, we use augmented 10,528 training images [10] to train our network. For evaluation of response maps of the training set, we use the set without augmentation with 1,464 examples, following the setting in [10]. For final semantic segmentation results, we use 1,449 images in the validation set to compare our results with other methods¹. In all experiments, the mean Intersection-over-Union (mIoU) ratio is used as the evaluation metric.

¹Although there is a test set that can be evaluated on the official PASCAL VOC website, by the submission deadline the website is still out of service for returning the evaluated performance.



Figure 3: Sample results of initial responses. Our approach often produces the response map that covers more complete region of the object (i.e., attention on the body of the animal), while the initial cue obtained by CAM [53] is prone to focus on small discriminative parts.




Figure 4: Enhancement on refinement. Our regularization enforces a more uniform response on objects, which can facilitate the refinement step. The examples illustrate that the IoU difference of the resultant refined map is significantly larger than the one of initial response.

4.2 Ablation Study and Analysis

Improvement on Response Map. We first present results of the initial and refined response maps. In Table 1, we show the performance for the original CAM used by the baseline AffinityNet [4], our CAM using the mixup data augmentation (Mixup \mathcal{L}_{cls}), and our final Mixup-CAM with mixup and uncertainty regularization (Mixup $\mathcal{L}_{cls} + \mathcal{L}_{ent} + \mathcal{L}_{con}$). In both results of CAM and its refinement, our IoU improvements are consistent after gradually adding the mixup augmentation and regularization. In addition, Figure 3 shows example results of the initial response, which illustrates our Mixup-CAM is able to make the network attend to more object parts and produce more uniform response distributions on objects.

Effectiveness of Regularization. One interesting aspect we find is that adding regularization could enhance the effectiveness of the refinement step. In Table 1, compared to Mixup \mathcal{L}_{cls} , adding either \mathcal{L}_{ent} (3rd row) or \mathcal{L}_{con} (4th row) improves the CAM IoU by 0.2% and 0.6% respectively. Nevertheless, with the refinement, the corresponding improvements in IoU is 1.1% and 1.2%, which are larger than the ones before refining the response. This is because our regularization enforces a more uniform response on objects, which greatly facilitates the refinement step (e.g., via region expanding). In addition, we illustrate one example in Figure 4, where the IoU difference of initial response is relatively small, but the resultant refined map could differ significantly.




Figure 5: Sensitivity analysis for parameters. (a) α for mixup augmentation; (b) λ_{ent} and (c) λ_{con} for uncertainty regularization.

Parameter Sensitivity. In this paper, we mainly study three parameters in our Mixup-CAM framework, i.e., α for mixup regularization and $\{\lambda_{ent}, \lambda_{con}\}$ in uncertainty regularization. In Figure 5(a), when increasing the α value, the *Beta* distribution would become more uniform, which encourages a more uniform λ_{mix} in (2) and results in mixed images that are more challenging to optimize. Nevertheless, the IoUs of Mixup \mathcal{L}_{cls} under various α are consistently better than the CAM baseline. For regularization terms, we fix $\lambda_{con} = 2e - 4$ and adjust λ_{ent} in Figure 5(b), while fixing $\lambda_{ent} = 0.2$ and change λ_{con} in Figure 5(c). Both figures show that these two parameters are robust to the performance under a wide range.

4.3 Semantic Segmentation Performance

After generating the pseudo ground truths using the refined response map, we use them to train the semantic segmentation network. First, we compare our method with state-of-the-art algorithms using the ResNet-101 architecture or other similarly powerful ones in Table 2. Note that, while our method focuses on improving the initial responses on the object, most methods aim to improve the refinement step or segmentation network training. In Table 3, we further present detailed performance for each category. We show two groups of results without (top rows) or with (bottom rows) applying CRF [27] to refine final segmentation outputs. Compared to the recent FickleNet [28] approach that also tries to improve the initial response map, our proposed Mixup-CAM shows favorable performance in final semantic segmentation results.

5 Conclusions

In this paper, we propose the Mixup-CAM framework to improve the localization of object response maps, as an initial step towards weakly-supervised semantic segmentation task using image-level labels. To this end, we propose to integrate the mixup data augmentation strategy for calibrating the uncertainty in network prediction. Furthermore, we introduce another two regularization terms as the interplay with the mixup scheme, thereby producing more complete and uniform response maps. In experimental results, we provide comprehensive analysis of each component in the proposed method and show that our approach achieves state-of-the-art performance against existing algorithms.

Table 2: Comparison of WSSS methods using image-level labels on the PASCAL VOC 2012 validation set. ✓ indicates the methods that focus on improving the initial response. The result of † on AffinityNet is re-produced by training the same ResNet-101 as our pipeline.

Method	Backbone	Init. Resp.	IoU on Val
MCOF <small>CVPR'18</small> [13]	ResNet-101		60.3
DCSP <small>BMVC'17</small> [9]	ResNet-101		60.8
DSRG <small>CVPR'18</small> [14]	ResNet-101		61.4
AffinityNet <small>CVPR'18</small> [10]	Wide ResNet-38		61.7
AffinityNet [†] <small>CVPR'18</small> [10]	ResNet-101		61.9
SeeNet <small>NIPS'18</small> [15]	ResNet-101	✓	63.1
Zeng <i>et al</i> <small>ICCV'19</small> [16]	DenseNet-169		63.3
BDSSW <small>ECCV'18</small> [17]	ResNet-101		63.6
OAA <small>ICCV'19</small> [18]	ResNet-101	✓	63.9
CIAN <small>CVPR'19</small> [19]	ResNet-101		64.1
FickleNet <small>CVPR'19</small> [20]	ResNet-101	✓	64.9
SSDD <small>ICCV'19</small> [6d]	Wide ResNet-38		64.9
Ours	ResNet101	✓	65.6

Table 3: Semantic segmentation performance on the PASCAL VOC 2012 validation set. The bottom group contains results with CRF refinement, while the top group is without CRF. The best three results are in red, green and blue, respectively.

Method	bkg	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	motor	person	plant	sheep	sofa	train	tv	mIoU
AffinityNet [10]	88.2	68.2	30.6	81.1	49.6	61.0	77.8	66.1	75.1	29.0	66.0	40.2	80.4	62.0	70.4	73.7	42.5	70.7	42.6	68.1	51.6	61.7
Ours (w/o CRF)	87.6	54.5	30.7	73.0	46.5	72.0	86.5	74.8	87.6	31.3	80.8	50.3	82.6	74.5	67.2	68.7	39.6	79.2	44.8	64.9	51.1	64.2
MCOF [13]	87.0	78.4	29.4	68.0	44.0	67.3	80.3	74.1	82.2	21.1	70.7	28.2	73.2	71.5	67.2	53.0	47.7	74.5	32.4	71.0	45.8	60.3
Zeng <i>et al.</i> [16]	90.0	77.4	37.5	80.7	61.6	67.9	81.8	69.0	83.7	13.6	79.4	23.3	78.0	75.3	71.4	68.1	35.2	78.2	32.5	75.5	48.0	63.3
FickleNet [20]	89.5	76.6	32.6	74.6	51.5	71.1	83.4	74.4	83.6	24.1	73.4	47.4	78.2	74.0	68.8	73.2	47.8	79.9	37.0	57.3	64.6	64.9
SSDD [6d]	89.0	62.5	28.9	83.7	52.9	59.5	77.6	73.7	87.0	34.0	83.7	47.6	84.1	77.0	73.9	69.6	29.8	84.0	43.2	68.0	53.4	64.9
Ours (w/ CRF)	88.4	57.0	31.2	75.2	47.8	72.4	87.2	76.0	89.2	32.7	83.1	51.1	85.4	77.3	68.4	70.2	40.0	81.5	46.2	65.4	51.8	65.6

Acknowledgments. This work is supported in part by the NSF CAREER Grant #1149783, and gifts from eBay and Google.

References

- [1] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *CVPR*, 2018. 1, 2, 3, 5, 6, 7, 8, 10
- [2] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *ECCV*, 2016. 1
- [3] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *NeurIPS*, 2019. 4
- [4] Yu-Ting Chang, Qiaosong Wang, Wei-Chih Hung, Robinson Piramuthu, Yi-Hsuan

- Tsai, and Ming-Hsuan Yang. Weakly-supervised semantic segmentation via sub-category exploration. In *CVPR*, 2020. 1
- [5] Arslan Chaudhry, Puneet K Dokania, and Philip HS Torr. Discovering class-specific pixels for weakly-supervised semantic segmentation. In *BMVC*, 2017. 10
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *CoRR*, abs/1606.00915, 2016. 1, 7
- [7] Yi-Wen Chen, Yi-Hsuan Tsai, Yen-Yu Lin, and Ming-Hsuan Yang. Vostr: Video object segmentation via transferable representations. *International Journal of Computer Vision*, 02 2020. 2
- [8] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *CVPR*, 2019. 3
- [9] Jifeng Dai, Kaiming He, and Jian Sun. Boxsup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *ICCV*, 2015. 1
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 6
- [11] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2010 (VOC2010) Results. "<http://www.pascal-network.org/challenges/VOC/voc2010/workshop/index.html>". 2, 7
- [12] Junsong Fan, Zhaoxiang Zhang, and Tieniu Tan. Cian: Cross-image affinity net for weakly supervised semantic segmentation. In *CVPR*, 2019. 3, 10
- [13] Ruochen Fan, Qibin Hou, Ming-Ming Cheng, Gang Yu, Ralph R Martin, and Shi-Min Hu. Associating inter-image salient instances for weakly supervised semantic segmentation. In *ECCV*, 2018. 3, 10
- [14] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *NIPS*, 2005. 4
- [15] Hongyu Guo, Yongyi Mao, and Richong Zhang. Mixup as locally linear out-of-manifold regularization. In *AAAI*, 2019. 4
- [16] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *ICCV*, 2011. 7
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 7
- [18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 3
- [19] Jonathan helhamer, Evan an Long and Trevor Darrell. Fully convolutional networks for semantic segmentation. *TPAMI*, 2016. 1
- [20] Qibin Hou, PengTao Jiang, Yunchao Wei, and Ming-Ming Cheng. Self-erasing network for integral object attention. In *NIPS*, 2018. 3, 10

- [21] Zilong Huang, Xinggang Wang, Jiasi Wang, Wenyu Liu, and Jingdong Wang. Weakly-supervised semantic segmentation network with deep seeded region growing. In *CVPR*, 2018. [2](#), [3](#), [10](#)
- [22] Wei-Chih Hung, Varun Jampani, Sifei Liu, Pavlo Molchanov, Ming-Hsuan Yang, and Jan Kautz. Scops: Self-supervised co-part segmentation. In *CVPR*, 2019. [6](#)
- [23] Peng-Tao Jiang, Qibin Hou, Yang Cao, Ming-Ming Cheng, Yunchao Wei, and Hong-Kai Xiong. Integral object mining via online attention accumulation. In *CVPR*, 2019. [2](#), [3](#), [10](#)
- [24] A. Khoreva, R. Benenson, J. Hosang, M. Hein, and B. Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *CVPR*, 2017. [1](#)
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. [7](#)
- [26] Alexander Kolesnikov and Christoph H Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *ECCV*, 2016. [1](#), [3](#)
- [27] Philipp Krahenbuhl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, 2011. [7](#), [9](#)
- [28] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *CVPR*, 2019. [2](#), [3](#), [9](#), [10](#)
- [29] Kunpeng Li, Ziyang Wu, Kuan-Chuan Peng, Jan Ernst, and Yun Fu. Tell me where to look: Guided attention inference network. In *CVPR*, 2018. [2](#)
- [30] Sungbin Lim, Ildoo Kim, Taesup Kim, Chiheon Kim, and Sungwoong Kim. Fast autoaugmentation. In *NeurIPS*, 2019. [3](#)
- [31] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *CVPR*, 2016. [2](#)
- [32] George Papandreou, Liang-Chieh Chen, Kevin Murphy, and Alan L Yuille. Weakly-and semi-supervised learning of a dcnn for semantic image segmentation. In *ICCV*, 2015. [1](#)
- [33] Deepak Pathak, Philipp Krahenbuhl, and Trevor Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *ICCV*, 2015. [1](#)
- [34] Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548*, 2017. [4](#)
- [35] Pedro O Pinheiro and Ronan Collobert. From image-level to pixel-level labeling with convolutional networks. In *CVPR*, 2015. [1](#)
- [36] Wataru Shimoda and Keiji Yanai. Self-supervised difference detection for weakly-supervised semantic segmentation. In *ICCV*, 2019. [3](#), [10](#)

- [37] Krishna Kumar Singh and Yong Jae Lee. Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In *ICCV*, 2017. 2
- [38] Cecilia Summers and Michael J Dinneen. Improved mixed-example data augmentation. In *WACV*, 2019. 4
- [39] Sunil Thulasidasan, Gopinath Chennupati, Jeff Bilmes, Tanmoy Bhattacharya, and Sarah Michalak. On mixup training: Improved calibration and predictive uncertainty for deep neural networks. In *NeurIPS*, 2019. 2, 4, 5
- [40] Yi-Hsuan Tsai, Guangyu Zhong, and Ming-Hsuan Yang. Semantic co-segmentation in videos. In *ECCV*, 2016. 2
- [41] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, Aaron Courville, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. *arXiv preprint arXiv:1806.05236*, 2018. 4
- [42] Paul Vernaza and Manmohan Chandraker. Learning random-walk label propagation for weakly-supervised semantic segmentation. In *CVPR*, 2017. 2
- [43] Xiang Wang, Shaodi You, Xi Li, and Huimin Ma. Weakly-supervised semantic segmentation by iteratively mining common object features. In *CVPR*, 2018. 2, 3, 10
- [44] Yunchao Wei, Jiashi Feng, Xiaodan Liang, Ming-Ming Cheng, Yao Zhao, and Shuicheng Yan. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *CVPR*, 2017. 2
- [45] Yunchao Wei, Huaxin Xiao, Honghui Shi, Zequn Jie, Jiashi Feng, and Thomas S Huang. Revisiting dilated convolution: A simple approach for weakly-and semi-supervised semantic segmentation. In *CVPR*, 2018. 2, 3
- [46] Zifeng Wu, Chunhua Shen, and Anton Van Den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition*, 90:119–133, 2019. 6
- [47] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *ICLR*, 2016. 1
- [48] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. 4
- [49] Yu Zeng, Yunzhi Zhuge, Huchuan Lu, and Lihe Zhang. Joint learning of saliency detection and weakly supervised semantic segmentation. In *CVPR*, 2019. 3, 10
- [50] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. 2, 4, 5
- [51] Xiaolin Zhang, Yunchao Wei, Jiashi Feng, Yi Yang, and Thomas S Huang. Adversarial complementary learning for weakly supervised object localization. In *CVPR*, 2018. 2
- [52] Guangyu Zhong, Yi-Hsuan Tsai, and Ming-Hsuan Yang. Weakly-supervised video scene co-parsing. In *ACCV*, 2016. 2
- [53] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016. 2, 4, 5, 8