Bottom-Up Top-Down Cues for Weakly-Supervised Semantic Segmentation

Qibin Hou[1*]   Puneet Kumar Dokania[2*]   Daniela Massiceti[2]   Yunchao Wei[3]
Ming-Ming Cheng[1]   Philip H. S. Torr[2]

[1]CCCE, Nankai University   [2] University of Oxford   [3]NUS

Abstract

We consider the task of learning a classifier for semantic segmentation using weak supervision in the form of image labels which specify the object classes present in the image. Our method uses deep convolutional neural networks (CNNs) and adopts an Expectation-Maximization (EM) based approach. We focus on the following three aspects of EM: (i) initialization; (ii) latent posterior estimation (E-step) and (iii) the parameter update (M-step). We show that saliency and attention maps, our bottom-up and top-down cues respectively, of simple images provide very good cues to learn an initialization for the EM-based algorithm. Intuitively, we show that before trying to learn to segment complex images, it is much easier and highly effective to first learn to segment a set of simple images and then move towards the complex ones. Next, in order to update the parameters, we propose minimizing the combination of the standard softmax loss and the KL divergence between the true latent posterior and the likelihood given by the CNN. We argue that this combination is more robust to wrong predictions made by the expectation step of the EM method. We support this argument with empirical and visual results. Extensive experiments and discussions show that: (i) our method is very simple and intuitive; (ii) requires only image-level labels; and (iii) consistently outperforms other weakly-supervised state-of-the-art methods with a very high margin on the PASCAL VOC 2012 dataset.

1. Introduction

The semantic segmentation task has rapidly advanced with the use of Convolutional Neural Networks (CNNs) [5, 6, 24, 37]. The performance of CNNs, however, is largely dependent on the availability of a large corpus of annotated training data, which is both cost- and time-intensive to acquire. The pixel-level annotation of an image in PASCAL VOC takes on average 4 minutes [4], which is likely a conservative estimate given that it is based on the COCO dataset [23] in which ground-truths are obtained by annotating polygon corners rather than pixels directly. In response, recent works focus on weakly-supervised semantic segmentation [4, 19, 27, 28, 30, 31, 35]. These differ from fully-supervised cases in that rather than having pixel-level ground-truth segmentations, the supervision available is to some lesser degree (image-level labels [19, 27, 28, 30], bounding boxes [27], or points and scribbles [4, 22, 34]).

In this work, we address the semantic segmentation task using only image labels, which specify the object categories present in the image. Our motivation for this is two-fold: (i) the annotation of an image with 20 object classes in PASCAL VOC is estimated to take 20 seconds, which is at least 12 times faster than a pixel-level annotation and is also scalable, and (ii) images with their image labels or tags can easily be downloaded from the Internet, providing a rich and virtually infinite source of training data. The method we adopt, similarly to the weakly-supervised semantic segmentation of [27], takes the form of Expectation-Maximization (EM) [10, 25]. An EM-based approach has three key steps: (i) initialization; (ii) latent posterior estimation (E step); and (ii) parameter update (M step). We focus on all of these aspects. In what follows we briefly talk about each of them.

We provide an informed initialization to the EM algorithm by training an initial model for the semantic segmentation task using an approximate ground-truth obtained using the combination of class-agnostic *saliency* map [15] and class-specific *attention* maps [36] on a set of simple images with one object category (*ImageNet* [11] classification dataset). This intuition comes from the obvious fact that it is easier to learn to segment simple images and then progress to more complex ones, similar to the work of [35]. Given an image, our saliency map *finds* the object (Figure 1b) - this is a class-agnotic 'bottom-up' cue. Added to this, once provided with the class present in the image (from the image label - 'boat' in this case), our attention map (Figure 1c) gives the 'top-down' class-specific regions in the image. Since both saliency and attention maps are tasked to find the same object, their combination is more powerful than if either one is used in isolation, as shown

---
[*]Joint first author

1

| (a) Input Image | (b) Saliency Map | (c) Attention Map | (d) Combined Map |

Figure 1: An example of combining bottom-up and top-down cues for simple images. (a) Input image with 'boat'. (b) Saliency map [15] ('bottom-up' cue). (c) Attention map [36] ('top-down' cue). (d) Combined map - notice that the saliency map puts high probability mass on regions missed by the attention map, and vice versa.

in Figure 1d. The combined probability map is then used as the ground-truth for training an initial model for the semantic segmentation task. The trained initial model provides the initialization parameters for the follow-up EM algorithm. Notice that this initialization is in contrast to [27] where the initial model is trained for the *image classification task* on the same *ImageNet* dataset. Experimentally we have found that this simple way of combining bottom-up with top-down cues on the *ImageNet* dataset (with *no* images from PASCAL VOC 2012) allows us to obtain an initial model capable of outperforming all current state-of-the-art algorithms for the weakly-supervised semantic segmentation task on the PASCAL VOC 2012 dataset. *This is surprising since these algorithms are significantly more complex and they rely on higher degrees of supervision such as bounding boxes, points/squiggles and superpixels.* This clearly indicates the importance of learning from simple images before delving into more complex ones. With the trained initial model, we then incorporate PASCAL VOC images (with multiple objects) for the E and M Steps of our EM-based algorithm.

In the E-step, we obtain the latent posterior probability distribution by constraining (or regularizing) the CNN likelihood using image labels based prior. This reduces many false positives by redistributing the probability masses (which were initially over the 20 object categories) among only the labels present in the image and the background. In the M-step, the parameter update step, we then minimize a combination of the standard *softmax* loss (where the ground-truth is assumed to be a Dirac delta distribution) and the KL divergence [21] between the latent posterior distribution (obtained using the E-step) and the likelihood given by the CNN. In the weakly-supervised setting this makes the approach more robust than using the *softmax* loss alone since in the case of confusing classes, the latent posterior (from the E-step) can sometimes be completely wrong. In addition to this, to obtain better CNN parameters, we add a probabilistic approximation of the Intersection-over-Union (IoU) [1, 9, 26] to the above loss function.

With this intuitive approach we obtain state-of-the-art re-

sults in the weakly-supervised semantic segmentation task on the PASCAL VOC 2012 benchmark [12]. We beat the best method which uses image label supervision by 10%.

2. Related Work

Work in weakly-supervised semantic segmentation has explored varying levels of supervision including combinations of image labels [19, 27, 28, 35], annotated points [4], squiggles [22, 34], and bounding boxes [27]. Papandreou *et al*. [27] employ an EM-based approach with supervision from image labels and bounding boxes. Their method iterates between inferring a latent segmentation (E-step) and optimizing the parameters of a segmentation network (M-step) by treating the inferred latents as the ground-truth segmentation. Similarly, [35] train an initial network using saliency maps, following which a more powerful network is trained using the output of the initial network. The MIL frameworks of [30] and [29] use fully convolutional networks to learn pixel-level semantic segmentations from only image labels. The image labels, however, provide no information about the position of the objects in an image. To address this, localization cues can be used [30, 31], obtained through indirect methods like bottom-up proposal generation (for example, MCG [3]), and saliency- [35] and attention-based [36] mechanisms. Localization cues can also be obtained directly through point- and squiggle-level annotations [4, 22, 34].

Our method is most similar to the EM-based approach of [27]. We use *saliency* and *attention* maps to learn a network for a simplified semantic segmentation task which provides better initialization of the EM algorithm. This is in contrast to [27] where a network trained for a classification task is used as initialization. Also different from [27] where the latent posterior is approximated by a Dirac delta function (which we argue is too harsh of a constraint in a weakly-supervised setting), we instead propose to use the combination of the true posterior distribution and the Dirac delta function to learn the parameters.

3. The Semantic Segmentation Task

Consider an image I consisting of a set of pixels $\{y_1, \cdots, y_n\}$. Each pixel can be thought of as a random variable taking on a value from a discrete semantic label set $\mathcal{L} = \{l_0, l_1, \cdots, l_c\}$, where c is the number of classes (l_0 for the background). Under this setting, a semantic segmentation is defined as the assignment of all pixels to their corresponding semantic labels, denoted as $\mathbf{y}$.

CNNs are extensively used to model the class-conditional likelihood for this task. Specifically, assuming each random variable (or pixel) to be independent, a CNN models the likelihood function of the form $P(\mathbf{y}|I;\theta) = \prod_{m=1}^{n} p(y_m|I;\theta)$, where $p(y_m = l|I;\theta)$ is the *softmax* probability (or the marginal) of assigning label l to the m-th pixel, obtained by applying the *softmax*[1] function to the CNN outputs $f(y_m|I;\theta)$ such that $p(y_m = l|I;\theta) \propto \exp(f(y_m = l|I;\theta))$. Given a training dataset $\mathcal{S} = \{I_i, \mathbf{y}_i\}_{i=1}^{N}$, where I_i and $\mathbf{y}_i$ represent the i-th image and its corresponding ground-truth semantic segmentation, respectively, the log-likelihood is maximized by minimizing the cross-entropy loss function using the back-propagation algorithm to obtain the optimal θ. At test time, for a given image, the learned θ is used to obtain the *softmax* probabilities for each pixel. These probabilities are either post-processed or used directly to assign semantic labels to each pixel.

4. Weakly-Supervised Semantic Segmentation

As mentioned in Sec. 3, to find the optimal θ for the semantic segmentation task, we need a dataset with ground-truth pixel-level semantic labels, obtaining which is highly time-consuming and expensive: for a given image, annotating it with 20 object classes takes nearly 20 seconds, while pixel-wise segmentation takes nearly 239.7 seconds [4]. This is highly non-scalable to higher numbers of images and classes. Motivated by this, we use an Expectation-Maximization (EM) [10, 25] based approach for weakly-supervised semantic segmentation using only image labels. Let us denote $Z = \mathcal{L} \setminus l_0$ as the set of object labels we are interested in, and a weak dataset as $\mathcal{D} = \{I_i, \mathbf{z}_i\}_{i=1}^{N}$ where I_i and $\mathbf{z}_i \subseteq Z$ are the i-th image and the labels present in the image. The task is to learn an optimal θ using $\mathcal{D}$.

4.1. The EM Algorithm

Similar to [27], we treat the unknown semantic segmentation $\mathbf{y}$ as the latent variable. Our probabilistic graphical model is of the following form (Fig. 2):

$$P(I, \mathbf{y}, \mathbf{z}; \theta) = P(I)P(\mathbf{y}|I, \mathbf{z}; \theta)P(\mathbf{z}), \qquad (1)$$

where we assume that $P(\mathbf{y}|I, \mathbf{z}; \theta) = P(\mathbf{y}|I; \theta)P(\mathbf{y}|\mathbf{z})$. Briefly, to learn θ while maximizing the above joint proba-

---

[1]The *softmax* function is defined as $\sigma(f_k) = \frac{e^{f_k}}{\sum_{j=0}^{c} e^{f_j}}$.


Figure 2: The graphical model. I is the image and $\mathbf{z}$ is the set of labels present in the image. $\mathbf{y}$ is the latent variable (semantic segmentation) and θ is the set of parameters.

bility distribution, the three major steps of an EM algorithm are: (i) initialize the parameter θ_t; (ii) E-step: compute the expected complete-data log-likelihood $F(\theta; \theta_t)$; and (iii) M-step: update θ by maximizing $F(\theta; \theta_t)$. In what follows, we first talk about how to obtain a good initialization θ_t in order to avoid poor local maxima and then talk about optimizing parameters (E and M steps) for a given θ_t.

4.2. Initialization: Skipping Poor Local Maxima Using Bottom-Up Top-Down Cues

It is well known that in case that the log-likelihood has several maxima or saddle points, an EM-based approach is highly susceptible to mediocre local maxima and a good initialization is crucial [16]. We argue that instead of initializing the algorithm with parameters learned for the classification task using the *ImageNet* classification dataset [11], as is done by most state-of-the-art methods irrespective of their nature, it is much more effective and intuitive to initialize with parameters learned for solving the task at hand - that is *semantic segmentation* - using the same dataset. This would allow for the full power of the dataset to be harnessed. For weakly-supervised semantic segmentation, however, the challenge is that only image-level labels are accessible during the training process. In the following we address this issue to obtain a good initialization θ_t by training an initial CNN model over simple *ImageNet* images for the weakly-supervised semantic segmentation task.

Let us denote $\mathcal{D}(I)$ as a subset of images from the *ImageNet* dataset containing objects of the categories in which we are interested (for details of this dataset, see Section 5). Dataset $\mathcal{D}(I)$ mostly contains images with centered and clutter-free single objects, unlike the challenging PASCAL VOC 2012 [12]. Given $\mathcal{D}(I)$, in order to train the initial model to obtain θ_t, we need pixel-level semantic labels which are not available in the weakly-supervised setting. To circumvent this, we use the combination of a class-agnostic *saliency* map [8, 15] (bottom-up cue) and a class-specific *attention* map [36] (top-down cue) to obtain an approximate ground-truth probability distribution of labels for each pixel. Intuitively, a saliency map gives the probability of each pixel belonging to any foreground class, and an attention map provides the probability of it belonging to the given object class. Combining these two maps allow us to

---

Algorithm 1 Approximate Ground Truth Distribution

---

input Image I with one object category; Image label z
1: $M = zeros(n)$, n is the number of pixels.
2: $\mathbf{s} \leftarrow SaliencyMap(I)$ [15]
3: $\mathbf{a} \leftarrow AttentionMap(I, z)$ [36]
4: **for** each pixel $m \in I$ **do**
5:      $M(m) = h(\mathbf{s}(m), \mathbf{a}(m))$
6: **end for**
output M

---

obtain a very accurate ground-truth probability distribution of the labels for each pixel in the image (see Figure 1).

Precisely, as shown in Algorithm 1, for a given simple image $I \in \mathcal{D}(I)$ and its corresponding class label $z \in Z$, we combine the attention and the saliency values per pixel to obtain $M(m) \in [0, 1]$ for all the pixels in the image. The value of $M(m)$ for the m-th pixel denotes the probability of it being the z-th object category. Similarly, $1 - M(m)$ is the probability of it being the background. The combination function $h(.,.)$ in Algorithm 1 is a user-defined function that combines the saliency and the attention maps. In this work we employ the max function which takes the union of the attention and saliency maps (Figure 1), thus complement each other. Let us define the approximate ground-truth label distribution for the m-th pixel as δ_m^I. Thus, $\delta_m^I \in [0, 1]^{|\mathcal{L}|}$, where $\delta_m^I(z) = M(m)$ at the z-th index for the object category z, $\delta_m^I(0) = 1 - M(m)$ at the 0-th index for the background, and zero otherwise. Given δ_m^I for each pixel, we find θ_t by using a CNN and optimizing the per-pixel cross-entropy loss $\sum_{k \in \mathcal{L}} \delta_m^I(k) \log p(k|I; \theta)$ between δ_m^I and $p(y_m|I, \theta)$, where $p(y_m|I, \theta)$ is the CNN likelihood. Notice that using this approach to obtain the approximate ground-truth label distribution requires only image-level labels. No human annotator involvement is required. By using the probability value $M(m)$ directly instead of using a standard Dirac delta distribution makes the approach more robust to noisy attention and saliency maps. This approach can be seen as a way of mining class-specific noise-free pixels, and is motivated by the work of Bearman et al. [4] where humans annotate points and squiggles in complex images. Their work showed that instead of training a network with a fully-supervised dataset, the learning process can be sufficiently guided using only a few supervised pixels which are easy to obtain. Their approach still requires human intervention in order to obtain points and squiggles, whereas, our approach requires only image-level labels which makes it highly scalable.

4.3. Optimizing Parameters

Let us now talk about how to define and optimize the expected complete-data log-likelihood $F(\theta; \theta_t)$. By definition, $F(\theta; \theta_t) = \sum_{\mathbf{y}} P(\mathbf{y}|I, \mathbf{z}; \theta_t) \log P(I, \mathbf{y}, \mathbf{z}; \theta)$,

where the expectation is taken over the posterior over the latent variables at a given set of parameters θ_t, denoted as $P(\mathbf{y}|I, \mathbf{z}; \theta_t)$. In the case of semantic segmentation, the latent space is exponentially large $|\mathcal{L}|^n$, therefore, computing $F(\theta; \theta_t)$ is infeasible. However, as will be shown, the independence assumption over the random variables, namely $P(\mathbf{y}|I; \theta) = \prod_{m=1}^{n} p(y_m|I; \theta)$, allows us to maximize $F(\theta; \theta_t)$ efficiently by decomposition. By using Eqn. (1), the independence assumption, the identity $\sum_{\mathbf{y}} P(\mathbf{y}|I, \mathbf{z}; \theta_t) = 1$, and ignoring the terms independent of θ, $F(\theta; \theta_t)$ can be written in a simplified form as:

$$\bar{F}(\theta; \theta_t) = \sum_{m=1}^{n} \sum_{\mathbf{y}} P(\mathbf{y}|I, \mathbf{z}; \theta_t) \log p(y_m|I; \theta) \quad (2)$$

Without loss of generality, we can write $P(\mathbf{y}|I, \mathbf{z}; \theta_t) = P(\mathbf{y} \setminus y_m|I, \mathbf{z}, y_m; \theta_t) p(y_m|I, \mathbf{z}; \theta_t)$, and using the identity $\sum_{\mathbf{y} \setminus y_m} P(\mathbf{y} \setminus y_m|I, \mathbf{z}, y_m; \theta_t) = 1$, we obtain:

$$\bar{F}(\theta; \theta_t) = \sum_{m=1}^{n} \sum_{y_m \in \mathcal{L}} p(y_m|I, \mathbf{z}; \theta_t) \log p(y_m|I; \theta) \quad (3)$$

Therefore, the M-step parameter update, which is maximizing $\bar{F}(\theta; \theta_t)$ w.r.t. θ, can be written as:

$$\theta_{t+1} = \underset{\theta}{\operatorname{argmax}} \sum_{m=1}^{n} \sum_{y_m \in \mathcal{L}} p(y_m|I, \mathbf{z}; \theta_t) \log p(y_m|I; \theta) \quad (4)$$

As mentioned in Sec. 4.1, the latent posterior distribution is $P(\mathbf{y}|I, \mathbf{z}; \theta_t) = P(\mathbf{y}|I; \theta_t) P(\mathbf{y}|\mathbf{z})$, where $P(\mathbf{y}|I; \theta_t)$ is the likelihood obtained using the CNN at a given θ_t. The distribution $P(\mathbf{y}|\mathbf{z})$ can be used to *regularize* the CNN likelihood $P(\mathbf{y}|I; \theta_t)$ by imposing constraints based on the image label information given in the dataset. Note that, $P(\mathbf{y}|\mathbf{z})$ is independent of θ and is a task-specific user-defined distribution that depends on the image labels or object categories. For example, if we know that there are only two classes in a given training image such as 'cat' and 'person' out of many other possible classes, then we would like to push the latent posterior probability $P(\mathbf{y}|I, \mathbf{z}; \theta_t)$ of absent classes to zero and increase the probability of the present classes. To impose this constraint, let us assume that similar to the likelihood, $P(\mathbf{y}|\mathbf{z})$ also decomposes over the pixels and belongs to the exponential family distribution $p(y_m|\mathbf{z}) \propto \exp(g(y_m, \mathbf{z}))$, where $g(.,.)$ is the user-define function designed to impose the desired constraints. Thus, the posterior can be written as $p(y_m|I, \mathbf{z}; \theta_t) \propto \exp(f(y_m|I; \theta_t) + g(y_m, \mathbf{z}))$. In order to impose the above mentioned constraints, we use the following form of $g(.,.)$:

$$g(y_m, \mathbf{z}) = \begin{cases} -\infty, & if \ y_m \notin \mathbf{z} \cup l_0, \\ 0, & otherwise. \end{cases} \quad (5)$$

4

Practically speaking, imposing the above constraint is equivalent to obtaining *softmax* probabilities for only those classes (including background) present in the image and assigning a probability of zero to all other classes. In other words, the above definition of $g(.,.)$ inherently defines $p(y_m|\mathbf{z})$ such that it is uniform for the classes present in the image including the background (l_0) and zero for the remaining ones. Other forms of $g(.,.)$ can also be used to impose different task-specific label-dependent constraints.

Cross entropy functions Consider the parameter update (M-step) as defined in Eq. 4. Solving this is equivalent to minimizing the cross entropy or the KL divergence between the latent posterior distribution and the CNN likelihood. Notice that, as opposed to [27], which uses a Dirac delta approximation $\hat{p}$ of the posterior distribution, where $\hat{p}(\hat{l}_m) = 1$ at $\hat{l}_m = \arg\max_{l \in \mathcal{L}} p(y_m = l|I, \mathbf{z}; \theta_t)$ and otherwise zero, we use the true posterior distribution (or the regularized likelihood) itself. We argue that using a Dirac delta distribution imposes a *hard* constraint that is suitable only when we are very confident about the true label assignment (for example, in the fully-supervised setting where the label is provided by a human annotator). In the weakly-supervised setting where the latent posterior, which decides the label, can be noisy (mostly seen in the case of confusing classes), it is more suitable to use the true posterior distribution, obtained using the combination of the CNN likelihood $p(y_m|I, \theta_t)$ and the class label-based prior $p(y_m|\mathbf{z})$. We propose to optimize θ by combining this true posterior distribution and its Dirac delta approximation:

$$J_m(I, \mathbf{z}, \theta_t; \theta) = \sum_{y_m \in \mathcal{L}} \bar{p}(y_m|I, \mathbf{z}; \theta_t) \log p(y_m|I; \theta) \quad (6)$$

where, $\bar{p}(y_m|I, \mathbf{z}; \theta_t) = (1 - \epsilon)p(y_m|I, \mathbf{z}; \theta_t) + \epsilon\hat{p}(y_m)$. To investigate the interplay between the two terms, we define a criterion which we call the *Relative Heuristic* to compute the value of ϵ given the pixel-wise latent posterior $p(y_m = l|I, \mathbf{z}; \theta)$ and a user-defined hyper-parameter $\eta \in [0, 1]$:

$$\epsilon = \begin{cases} 1, & if \ r \geq \eta, \\ r, & otherwise. \end{cases} \quad (7)$$

where $r = (p_1 - p_2)/p_1$, and p_1 and p_2 are the highest and the second highest probability values in the latent posterior distribution. Intuitively, $\eta = 0.05$ implies that the most probable score should be at least 5% better compared to the second most probable score.

The IoU gain function Along with minimizing the cross entropy losses as shown in the Eq. 6, in order to obtain better parameter estimate, we also maximize the probabilistic approximation of the intersection-over-union (IoU) between

Algorithm 2 Our Final Algorithm

---

input Datasets $\mathcal{D}(P)$ and $\mathcal{D}(I)$; θ_0; η; K
1: Use $\mathcal{D}(I)$ and θ_0 to obtain initialization parameter θ_t using method explained in Section 4.2.
2: **for** $k = 1 : K$ **do**
3:     $\theta \leftarrow \theta_t$
4:     **for** each pixel m in $\mathcal{D}(P) \cup \mathcal{D}(I)$ **do**
5:        Obtain latent posterior: $p_m(y_m|I, \mathbf{z}; \theta) \propto \exp(f_m(y_m|I; \theta) + g(y_m, \mathbf{z}))$
6:     **end for**
7:     Optimize Eq. 9 using CNN to update θ_t.
8: **end for**

---

the posterior distribution and the likelihood [1, 9, 26]:

$$\mathcal{J}_{IOU}(P(\mathbf{y}|I, \mathbf{z}; \theta_t), P(\mathbf{y}|I; \theta)) \approx$$
$$\frac{1}{|\mathcal{L}|} \sum_{l \in \mathcal{L}} \frac{\sum_{m=1}^{n} p_m^t(l) p_m^\theta(l)}{\sum_{m=1}^{n}\{p_m^t(l) + p_m^\theta(l) - p_m^t(l)p_m^\theta(l)\}} \quad (8)$$

where, $p_m^t(l) = p(y_m = l|I, \mathbf{z}; \theta_t)$ and $p_m^\theta(l) = p(y_m = l|I; \theta)$. Refer to [9] for further details about Eq. 8.

Overall objective function and the algorithm Combining the cross entropy loss function (Eq. 6) and the IoU gain function (Eq. 8), the M-step parameter update problem is:

$$\theta_{t+1} = \arg\max_\theta \sum_{m=1}^{n} J_m(I, \mathbf{z}, \theta_t; \theta) + \mathcal{J}_{IOU} \quad (9)$$

We use a CNN model along with the back-propagation algorithm to optimize the above objective function. Recall that our evaluation is based on the PASCAL VOC 2012 dataset, therefore, during the M-step of the algorithm we use both the *ImageNet* $\mathcal{D}(I)$ and the PASCAL trainval $\mathcal{D}(P)$ datasets (see Section 5 for details). Our overall approach is summarized in Algorithm 2.

5. Experimental Results and Comparisons

We show the efficacy of our method on the challenging PASCAL VOC 2012 benchmark and outperform all existing state-of-the-art methods by a large margin. Specifically, we improve on the current state-of-the-art weakly-supervised method using only image labels[31] by 10%.

5.1. Setup

Dataset $\mathcal{D}(I)$. To train our initial model (Section 4.2), we download 80,000 images from the *ImageNet* dataset which contain objects in the 20 foreground object categories of the PASCAL VOC 2012 segmentation task. We filter this dataset using simple heuristics. First, we discard images with width or height less than 200 or greater than 500 pixels. Using the attention model of [36] (which is trained with

Method	Dataset	Dependencies	Supervision	CRF [20]	mIoU (Val)	mIoU (Test)
EM Adapt [27]	$\mathcal{D}(I), \mathcal{D}(P)$	No	Image labels	✗	–	–
				✓	38.2%	39.6%
CCNN [28]	$\mathcal{D}(I), \mathcal{D}(P)$	No	Image labels	✗	33.3%	35.6%
				✓	35.3%	–
		Class size		✗	*40.5%*	*43.3%*
				✓	42.4%	45.1%
SEC [19]	$\mathcal{D}(I), \mathcal{D}(P)$	Saliency [32] & Localization [38]	Image labels	✗	*44.3%*	–
				✓	50.7%	51.7%
MIL [30]	$\mathcal{D}(I)$	Superpixels [13]	Image labels	✗	36.6%	35.8%
		BBox BING [7]			*37.8%*	*37.0%*
		MCG [3]			42.0%	40.6%
WTP [4]	$\mathcal{D}(I), \mathcal{D}(P)$	Objectness [2]	Image labels	–	*32.2%*	–
			Image labels + 1 Point/Class		42.7%	–
			Image labels + 1 Squiggle/Class		49.1%	–
STC [35]	$\mathcal{D}(I), \mathcal{D}(P), \mathcal{D}(F)$	Saliency [18]	Image labels	✓	49.8%	51.2%
AugFeed [31]	$\mathcal{D}(I), \mathcal{D}(P)$	SS [33]	Image labels	✗	*46.98%*	*47.8%*
				✓	52.62%	52.7%
		MCG [3]		✗	50.41%	50.6%
				✓	54.34%	55.5%
Ours — Initial Model	$\mathcal{D}(I)$	Saliency [15] & Attention [36]	Image labels	✗	53.53%	54.34%
				✓	55.19%	56.24%
Final	$\mathcal{D}(I), \mathcal{D}(P)$			✗	**56.91%**	**57.74%**
				✓	58.71%	59.58%

Table 1: **Comparison table**. All dependencies, datasets, and degrees of supervision used by most of the existing methods for weakly-supervised semantic segmentation. Table 3 complements this table by providing the degrees of supervision used by the dependencies themselves. $\mathcal{D}(I)$ is the simple filtered *ImageNet* dataset, $\mathcal{D}(P)$ is the complex filtered PASCAL VOC 2012 dataset (see Section 5) and $\mathcal{D}(F)$ contains 41K images from *Flickr* [35]. Note that the cross validation of CRF hyper-parameters and the training of MCG are performed using a fully-supervised pixel-level semantic segmentation dataset. To be fair, our method should only be compared with the numbers shown in italic with underline.

only image-level labels), we generate an attention map for each image and record the most probable class label with its corresponding probability. We discard the image if (i) its most probable label does not match the image label or (ii) its most probable label matches the image label but its corresponding probability is less than 0.2. We then generate saliency maps using the saliency model of [15] (which is trained with only class-agnostic saliency masks). For the remaining images, we combine attention and saliency by finding the pixel-wise intersection between the saliency and the attention binary masks. The saliency mask is obtained by setting the pixel's value to 1 if its corresponding saliency probability is greater than 0.5. The same is done to obtain the attention mask. For each object category, the images are sorted by this intersection area (i.e. the number of overlapping pixels between the two masks) with the intuition that larger intersections correspond to higher quality saliency and attention maps. The top 1500 images are then selected for each category, with the exception of the 'person' category where the top 2500 images are kept, and

any category with fewer than 1500 images, in which case all images are kept. *This complete filtering process leaves us with* 24,000 *simple images*, which contain uncluttered and mainly-centered single objects. We denote this dataset as $\mathcal{D}(I)$ and will make this dataset publicly available. *We highlight that* $\mathcal{D}(I)$ *does not contain any additional images relative to those used by other weakly supervised works* (see Dataset column in Table 1).

Datasets $\mathcal{D}(P)$ **and** $\mathcal{D}(I)$ **for M-step.** For the M-step, we use a filtered subset of PASCAL VOC 2012 images, denoted $\mathcal{D}(P)$, and a subset of $\mathcal{D}(I)$. To obtain $\mathcal{D}(P)$, we take complex PASCAL VOC 2012 images (10,582 in total, made up of 1,464 training images [12] and the extra images provided by [14]), and use the trained initial model (i.e. θ_t) to generate a (hard) ground-truth segmentation for each. The hard segmentations are obtained by assigning each pixel with the class label that has the highest probability. The ratio of the foreground area to the whole image area (where area is the sum of the number of pixels) is computed,

and if the ratio is below 0.05, the image is discarded. This leaves $10,000$ images. We also further filter $\mathcal{D}(I)$: using the trained initial model, we generate (hard) segmentations for all simple *ImageNet* images in $\mathcal{D}(I)$. We compute the intersection area (as above) between the attention mask and the predicted segmentation (rather than the saliency mask as before) and select the top $10,000$ of $24,000$ images based on this metric. Together $\mathcal{D}(P)$ and this subset of $\mathcal{D}(I)$ consist of $20,000$ images used for the M-step.

CNN architecture and parameter settings Similar to [19, 27, 35], both our initial model and our EM model are based on the largeFOV DeepLab architecture [6]. We use simple bilinear interpolation to map the downsampled feature maps to the original image size as suggested in [24]. We use the publicly available Caffe toolbox [17] for our implementation. We use weight decay (0.0005), momentum (0.9), and iteration size (10) for gradient accumulation. The learning rate is 0.001 at the beginning and is divided by 10 every 10 epochs. We use a batch size of 1 and randomly crop the input image to 321×321. Images with width or height less than 321 are padded with the mean pixel values and the corresponding places in the ground-truth are padded with ignore labels to nullify the effect of padding. We flip the images horizontally, resulting in an augmented set twice the original one. We train our networks for 30K iterations by optimizing Eq. 9 as per Algorithm 2 with $\eta = 0.05$ and $K = 2$. Performance gains beyond two EM iterations were not significant compared to the computational cost.

5.2. Results, Comparisons, and Analysis

We provide Table 1 and 3 for an extenstive comparison between our and current methods, their dependencies, and degrees of supervision. Regarding the dependencies of our method, our *saliency* network [15] is trained using salient region masks. These masks are class-agnostic, therefore, once trained the network can be used for any semantic object category, so there is no issue with scalability and no need to train the *saliency* network again for new object categories. Our second dependency, the *attention* network [36] is trained using solely image labels.

State-of-the-art. Our method outperforms all existing state-of-the-art methods by a very high margin. The most directly comparable method in terms of supervision and dependencies is *AugFeed* [31] which uses super-pixels. Our method obtains almost 10% better mIoU than *AugFeed* on both the *val* and *test* sets. Even if we disregard 'equivalent' supervision and dependencies, our method is still almost 2.6% and 2.2% better than the best method in the *val* and *test* sets, respectively. Table 2 shows class-wise performance of our method.

Simplicity vs sophistication. The initial model is essential to the success of our method. We train this model in a very simple and intuitive way by employing a filtered subset of simple *ImageNet* images and training for the semantic segmentation task. Importantly, this process uses only image labels and is fully automatic, requiring no human intervention. The learned θ_t provides a *very* good initialization for the EM algorithm, enabling it to avoid poor local maxima. This is shown visually in Figure 3: the initial model (third column) is already a good prediction, with the first and second EM iterations (fourth and fifth columns), improving the semantic segmentation even further. We highlight that with this simple approach, surprisingly, our initial model beats *all* current state-of-the-art methods, which are more complex and often use higher degrees of supervision. By implementing this intuitive modification, we believe that many methods can easily boost their performance.

To CRF or not to CRF? In our work, we specifically choose not to employ a CRF [20] as a post-processing step nor inside our models, during training and testing, for the following reasons: (1) CRF hyper-parameters are normally cross validated over a fully-supervised pixel-wise segmentation dataset which contradicts a fully "weak" supervision. This is likewise the case for MCG [3] which is trained on a pixel-level semantic segmentation dataset. (2) The CRF hyper-parameters are incredibly sensitive, and if we wish to extend our framework to incorporate new object categories, this would require a pixel-level annotated dataset of the new categories along with the old ones for the cross-validation of the CRF hyper-parameters. This is highly non-scalable. For completeness, however, we include our method with a CRF applied (Table 1 & 2) which boosts our accuracy by 1.8%. We note that even without a CRF, our best approach exceeds the state-of-the-art (which uses a CRF *and* a higher degree of supervision) by 2.2% on the *test* set.

6. Conclusions and Future Work

We have addressed weakly-supervised semantic segmentation using only image labels. We proposed an EM-based approach and focus on the three key components of the algorithm: (i) initialization, (ii) E-step and (iii) M-step. Using only the image labels of a filtered subset of *ImageNet*, we learn a set of parameters for the semantic segmentation task which provides an informed initialization of our EM algorithm. Following this, with each EM iteration, we empirically and qualitatively verify that our method improves the segmentation accuracy on the challenging PASCAL VOC 2012 benchmark. Furthermore, we show that our method outperforms all state-of-the-art methods by a large margin.

Future directions include making our method more robust to noisy labels, for example, when images downloaded from the Internet have incorrect labels, as well as

Data	Method	CRF	bkg	plane	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	motor	person	plant	sheep	sofa	train	tv	mIoU
Val	Initial	✗	87.1	74.7	29.0	69.8	55.8	55.6	73.3	65.2	63.4	15.8	61.5	15.9	60.0	56.4	57.5	53.7	32.9	65.6	23.9	64.6	42.2	53.53
		✓	87.7	79.7	32.6	73.4	58.4	57.8	74.3	64.8	66.0	15.9	63.1	15.0	62.3	59.6	57.7	54.9	33.8	69.1	23.7	65.0	44.3	55.19
	Final	✗	87.8	72.4	28.7	67.9	58.8	55.8	78.0	69.7	70.2	17.8	63.3	23.2	65.7	60.5	63.1	58.7	40.0	68.2	28.9	70.9	45.5	**56.91**
		✓	88.6	76.1	30.0	71.1	62.0	58.0	79.2	70.5	72.6	18.1	65.8	22.3	68.4	63.5	64.7	60.0	41.5	71.7	29.0	72.5	47.3	58.71
Test	Initial	✗	87.9	69.2	29.2	74.9	41.7	53.4	70.6	69.6	59.9	18.3	66.1	24.9	62.5	63.3	68.8	55.4	33.7	63.8	18.6	64.3	44.9	54.35
		✓	88.5	72.6	32.6	80.3	44.6	55.4	70.9	69.6	62.2	18.9	68.4	24.6	65.2	66.8	71.2	57.2	37.2	66.7	17.4	64.8	45.9	56.24
	Final	✗	88.2	69.5	29.7	72.2	45.1	57.3	73.2	72.7	69.3	20.5	65.4	33.5	67.8	64.0	72.3	58.9	45.5	69.8	26.8	63.8	46.8	**57.74**
		✓	88.9	72.7	31.0	76.3	47.7	59.2	74.3	73.2	71.7	19.9	67.1	34.0	70.3	66.6	74.4	60.2	48.1	73.1	27.8	66.9	47.9	59.58

Table 2: Per-class results (in %) on PASCAL VOC 2012 *val* and *test* sets. Initial shows results for initial model trained using $\mathcal{D}(I)$ for EM algorithm initialization. Final shows results at final EM iteration ($K = 2$). Results show with and without application of post-process CRF [20] (although we do not consider using a CRF as true weak supervision, see Section 5).

Dependency	Supervision
Class size	Image labels + Bboxes
Saliency	[32] Image labels
	[18] Bboxes
	[15] Saliency masks
Attention [36]	Image labels
Objectness [2]	Image labels + Bboxes
Localization [38]	Image labels
Superpixels [13]	None
Bbox BING [7]	Bboxes
MCG [3]	Pixel labels
SS [33]	Bboxes
CRF [20]	Pixel labels (parameter cross-val)

Table 3: **Dependency table**. The degrees of supervision required by the dependencies of several weakly-supervised methods for the semantic segmentation task. 'Bboxes' represents bounding boxes.

better handling images with multiple classes of objects.

References

[1] F. Ahmed, D. Tarlow, and D. Batra. Optimizing expected intersection-over-union with candidate-constrained crfs. In *ICCV*, 2015. 2, 5

[2] B. Alexe, T. Deselares, and V. Ferrari. Measuring the objectness of image windows. In *PAMI*, 2012. 6, 8

[3] P. Arbelaez, J. Pont-Tuset, J. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping. In *CVPR*, 2014. 2, 6, 7, 8

[4] A. Bearman, O. Russakovsky, V. Ferrari, and F.-F. Li. What's the point: Semantic segmentation with point supervision. In *ECCV*, 2016. 1, 2, 3, 4, 6

[5] S. Chandra and I. Kokkinos. Fast, exact and multi-scale inference for semantic image segmentation with deep gaussian crfs. In *ECCV*, 2016. 1

[6] L.-G. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected. In *ICLR*, 2015. 1, 7

[7] M. Cheng, Z. Zhang, W. Lin, and P. H. S. Torr. Bing: Binarized normed gradients for objectness estimation at 300fps. In *CVPR*, 2014. 6, 8

[8] M.-M. Cheng, N. J. Mitra, X. Huang, P. H. S. Torr, and S.-M. Hu. Global contrast based salient region detection. *IEEE TPAMI*, 37(3):569–582, 2015. 3

[9] M. Cogswell, X. Lin, S. Purushwalkam, and D. Batra. Combining the best of graphical models and convnets for semantic segmentation. In *arXiv:1412.4313*, 2014. 2, 5

[10] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. In *Journal of the Royal Statistical Society*, 1977. 1, 3

[11] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009. 1, 3

[12] M. Everingham, S. M. A. Eslami, L. V. Gool, C. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge a retrospective. In *IJCV*, 2014. 2, 3, 6

[13] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph based image segmentation. In *IJCV*, 2004. 6, 8

[14] B. Hariharan, P. Arbelaez, L. Bourdev, S. Maji, and J. Malik. Semantic contours from inverse detectors. In *ICCV*, 2011. 6

[15] Q. Hou, M.-M. Cheng, X.-W. Hu, A. Borji, Z. Tu, and P. Torr. Deeply supervised salient object detection with short connections. In *IEEE CVPR*, 2017. 1, 2, 3, 4, 6, 7, 8

[16] C. F. Jeff Wu. On the convergence properties of the em algorithm. In *The Annals of Statistics*, 1983. 3

[17] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM International Conference on Multimedia*, 2014. 7

[18] H. Jiang, J. Wang, Z. Yuan, Y. Wu, Z. N., and S. Li. Salient object detection: A discriminative regional feature integration approach. In *CVPR*, 2013. 6, 8

[19] A. Kolesnikov and C. H. Lampert. Seed, Expand and Constrain: Three Principles for Weakly-Supervised Image Segmentation. In *ECCV*, 2016. 1, 2, 6, 7

[20] P. Krahenbuhl and V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, 2011. 6, 7, 8

[21] S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 1951. 2

[22] D. Lin, J. Dai, J. Jia, K. He, and J. Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *CVPR*, 2016. 1, 2

[23] T. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 1

[24] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 1, 7

[25] G. J. McLachlan and T. Krishnan. *The EM algorithm and extensions*. Wiley, 1997. 1, 3

[26] S. Nowozin. Optimal decisions from probabilistic models: the intersection-over-union case. In *CVPR*, 2014. 2, 5

[27] G. Papandreou, L.-C. Chen, K. P. Murphy, and A. L. Yuille. Weakly- and semi-supervised learning of a DCNN for semantic image segmentation. In *ICCV*, 2015. 1, 2, 3, 5, 6, 7

[28] D. Pathak, P. Krahenbuhl, and T. Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *ICCV*, 2015. 1, 2, 6

[29] D. Pathak, E. Shelhamer, J. Long, and T. Darrell. Fully convolutional multi-class multiple instance learning. In *ICLR*, 2014. 2

[30] P. O. Pinheiro and R. Collobert. From image-level to pixel-level labeling with convolutional networks. In *CVPR*, 2015. 1, 2, 6

[31] X. Qi, Z. Liu, J. Shi, H. Zhao, and J. Jia. Augmented feedback in semantic segmentation under image level supervision. In *ECCV*, 2016. 1, 2, 5, 6, 7

[32] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *ICLR*, 2014. 6, 8

[33] J. R. R. Uijlings, K. E. A. van de Sande, T. Gevers, and A. W. M. Smeulders. Selective search for object recognition. In *IJCV*, 2013. 6, 8

[34] J. Xu, A. Schwing, and R. Urtasun. Learning to segment under various forms of weak supervision. In *CVPR*, 2015. 1, 2

[35] W. Yunchao, L. Xiaodan, C. Yunpeng, S. Xiaohui, M.-M. Cheng, Z. Yao, and Y. Shuicheng. STC: A simple to complex framework for weakly-supervised semantic segmentation. In *arXiv:1509.03150*, 2015. 1, 2, 6, 7

[36] J. Zhang, Z. Lin, J. Brandt, X. Shen, and S. Sclaroff. Top-down neural attention by excitation backprop. In *ECCV*, 2016. 1, 2, 3, 4, 5, 6, 7, 8

[37] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. Torr. Conditional random fields as recurrent neural networks. In *ICCV*, 2015. 1

[38] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016. 6, 8

| (1) Source | (2) Ground truth | (3) Initial model | (4) EM 1st iteration | (5) EM 2nd iteration |

Figure 3: Qualitative results for weakly-supervised semantic segmentation using our proposed EM-based method. It is clear that as we move from the initial model (3rd column) to the second iteration of the EM algorithm (5th column), the segmentation quality improves incrementally. The two bottom rows show failure cases.