

# Every Annotation Counts: Multi-label Deep Supervision for Medical Image Segmentation

Simon Reiß<sup>1,2</sup> Constantin Seibold<sup>1</sup> Alexander Freytag<sup>2</sup> Erik Rodner<sup>2,3</sup> Rainer Stiefelhagen<sup>1</sup>

<sup>1</sup>Karlsruhe Institute of Technology <sup>2</sup>Carl Zeiss AG <sup>3</sup>University of Applied Sciences Berlin

{simon.reiss, constantin.seibold, rainer.stiefelhagen}@kit.edu

## Abstract

*Pixel-wise segmentation is one of the most data and annotation hungry tasks in our field. Providing representative and accurate annotations is often mission-critical especially for challenging medical applications. In this paper, we propose a semi-weakly supervised segmentation algorithm to overcome this barrier. Our approach is based on a new formulation of deep supervision and student-teacher model and allows for easy integration of different supervision signals. In contrast to previous work, we show that care has to be taken how deep supervision is integrated in lower layers and we present multi-label deep supervision as the most important secret ingredient for success. With our novel training regime for segmentation that flexibly makes use of images that are either fully labeled, marked with bounding boxes, just global labels, or not at all, we are able to cut the requirement for expensive labels by 94.22% – narrowing the gap to the best fully supervised baseline to only 5% mean IoU. Our approach is validated by extensive experiments on retinal fluid segmentation and we provide an in-depth analysis of the anticipated effect each annotation type can have in boosting segmentation performance.*

## 1. Introduction

Medical imaging tools have become a central part in modern health care. In direct consequence, hundreds of millions of 2D- and 3D-images are recorded per year [41]. With this ever-growing number, medical personnel has become increasingly entangled in their evaluation. Deep learning with the use of artificial neural networks has found use to alleviate the necessary effort needed for interpretation of such images. However, training these models requires enormous amounts of annotated data, especially in case of semantic segmentation. The annotation process for this task is already extremely difficult in real-world scenarios such as in urban street-scenes, where the pixel-wise annotation process of a single image could span up to 90 minutes [15].

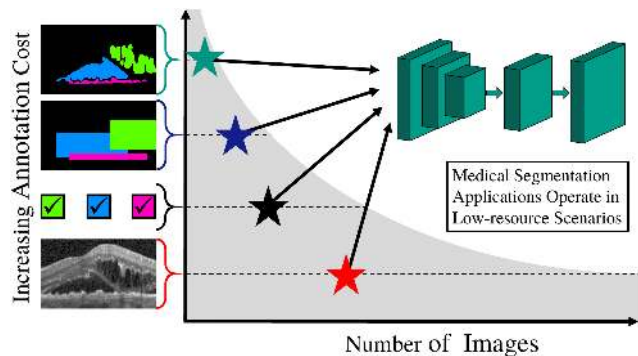


Figure 1. Annotations for segmentation are costly, especially when experts need to provide them. We show how our semi-weakly semantic segmentation method can use different annotation types and how the recognition performance benefits from them.

This problem is amplified in the medical domain as we further require domain experts to annotate data who are severely time-restricted due to their clinical work. Thus, we are faced with the problem of minimizing the needed annotation effort while maximizing model accuracy.

The majority of existing approaches follow two orthogonal paths: incorporating non-annotated data [24] which is fast to obtain or using cheap weak labels with different levels of quality ranging from image-level [10, 23, 32, 47] over single point [7] to bounding box annotations [29, 55]. Both these so called semi- and weakly-supervised approaches led to great insights and convincing results. However, they ignore that practical applications are often faced with several types of supervision simultaneously (Fig. 1).

In very common scenarios, we are provided with a small pixel-wise annotated data set, with automatically parsed image-level labels from medical reports, and with large amounts of additional unlabeled data from the same distribution. Currently, it is largely unanswered how such diverse supervision types can be unified to train a semantic segmentation system. To this end, we go beyond standard weakly- or semi-supervised learning and investigate also the semi-weakly supervised setting in the low-resource scenario of

the medical domain.

For dealing with only few annotated examples, we propose a novel pathway to integrate training signals deep into segmentation network layers via a new take on deep supervision. We then amplify these signals by enriching weakly- or entirely un-labeled images via our novel approach to infer robust pseudo-labels using a mean-teacher segmentation model.

Furthermore, as iterative training processes in low-data scenarios are often unstable, we present experiments following a rigorous evaluation protocol and report test accuracies with standard deviations along numerous data-splits.

Our contributions amount to:

- (1) We present the first thorough investigation of varying numbers of training samples and a large diversity of supervision types for semantic segmentation.
- (2) We introduce a novel perspective on the deep supervision paradigm adapting it to segmentation in our *Multi-label Deep Supervision* technique. With this, we introduce a flexible semi-weakly supervised pathway to integrate either un- or weakly labeled images: our novel *Self-Taught Deep Supervision* approach.
- (3) Finally, our best performing method *Mean-Taught Deep Supervision* adds invariance towards perturbations and a robust pseudo-label generation, achieving results close to fully supervised baselines while using only a fraction of 5.78% strong labels.

## 2. Related Work

**Mask supervision.** The most prominent research direction in semantic segmentation considers the task in a fully-supervised setting, assuming the availability of a sufficient amount of expensive pixel-wise labeled masks [5, 12, 13, 35, 39, 50, 67]. Based on this assumption, innovation in parameter-heavy models induces performance gains on large data sets [15, 18, 36], while in smaller-scale scenarios, *e.g.* domains where label acquisition is even more expensive, leveraging this progress is only partially possible.

**Bounding box supervision.** Coarse bounding boxes drawn around semantic regions offer strong location cues at manageable costs, but inherently include pixels not semantically associated to the box-label. On object-related data sets, leveraging this kind of supervision has led to impressive results [16, 29, 46, 55]. Yet, methods tend to leverage hand-crafted rules or tools like GrabCut [51], Multiscale Combinatorial Grouping [4], or Selective Search [60] to kick-start training. Therefore, such approaches are often not trivially transferred to data different from natural images. In medical imaging, bounding box supervision for segmentation has been studied for positron emission tomography scans [2] as well as magnetic resonance images [28, 49]. In this work,

we leverage box-supervision for retinal fluid segmentation in optical coherence tomography (OCT) scans.

**Image-level supervision.** Leveraging global labels with information about present classes in an image to train a segmentation model has recently seen growing interest in our community [10, 23, 30, 32, 47, 61]. While early approaches followed a Multiple Instance Learning (MIL) formulation [47], recent work builds upon feature attribution methods [53, 54, 68] for providing initial location cues and refine them using prior assumptions [10, 23, 30, 32, 61]. Such assumptions are often specific to the data set they were designed for, consequently having limited applicability for many medical domains. In consequence, only few approaches are known which investigate benefits of image-level labels in context of OCT scans [52, 63].

**Semi- and semi-weakly supervision.** Combining pixel-wise annotations with weaker annotations like global- or box-labels has been shown early in [46]. These scenarios are termed semi-weakly supervised segmentation [14]. With the emergence of generative adversarial networks, un- and pixel-wise labeled data were fused via adversarial objectives [24, 43, 56], exploring semi-supervised segmentation models. Making best use of bounding boxes in conjunction with pixel-wise labels has seen interest in [25, 29]. Recently, many attempts try to regularize the learner to form predictions invariant towards perturbations [34, 45]. Along this line of thought, the data augmentation strategy Cut-Mix [65] can serve as perturbation by mixing predictions and assembling new informative pseudo-labels [19, 44]. This approach is often accompanied by mean-teachers [58] serving as robust models for distilling invariance into the student. Closest to our work, medical image segmentation with a mean-teacher was done by changing the consistency term from [58] to better fit the segmentation task [48]. Our approach does not only enforce consistency between a teacher and student, we extend the idea with pseudo-targets. Yet, instead of adding complexity by generating them iteratively [6], we integrate pseudo-labels online and identify the missing link in coping with their noise: deep supervision.

**Deep supervision.** Early explorations of deep architectures for classification introduced companion objectives or dense pathways into lower layers [22, 31, 57]. The positive effects in low-data scenarios [64], on convergence, generalization and vanishing gradients [31] set off lots of work using deep supervision in segmentation models [11, 38, 40]. Especially the medical imaging community with their notoriously scarce-data applications showed strong interest [1, 17, 20, 33, 62, 64, 66, 69, 70]. The move from classification to segmentation is mostly done by up-scaling low resolution feature maps. We propose the counter intuitive reverse direction: down-scaling the pixel-wise annotations. By doing so in a novel semantic-preserving fashion, we uncover surprising properties when joining it with noisy pseudo-labels.

### 3. Proposed Approach

In this chapter, we introduce notation and the task of *semi-weakly supervised semantic segmentation*. Thereafter, we present our novel perspective on deep supervision: *Multi-label Deep Supervision* which we consider for traditional semantic segmentation and motivate its strong fit for semi- and semi-weakly supervised segmentation in our *Self-Taught* and *Mean-Taught Deep Supervision* frameworks.

#### 3.1. Preliminaries

##### 3.1.1 Supervision modalities and notation

For semi-weakly supervised semantic segmentation, we consider an image data set:

$$\mathcal{D} = \{x_1, \dots, x_n | x_i \in \mathbb{R}^{3 \times H \times W}\} . \quad (1)$$

An image  $x_i$  can have different available annotations, for example a pixel-wise annotated mask  $m_i$ , bounding box  $b_i$ , global image-level label  $g_i$  from:

$$\mathcal{M} = \{m_1, \dots, m_n | m_i \in [0, 1]^{(C+1) \times H \times W}\} , \quad (2)$$

$$\mathcal{B} = \{b_1, \dots, b_n | b_i \in [0, 1]^{(C+1) \times H \times W}\} , \quad (3)$$

$$\mathcal{G} = \{g_1, \dots, g_n | g_i \in [0, 1]^{(C+1)}\} , \quad (4)$$

or no annotation at all. Note, for  $m \in \mathcal{M}$ , we assume that at each spatial position  $(x, y) \in H \times W$  only a single class is set to one. We replace the common representation of bounding boxes as two points  $(x_1, y_1), (x_2, y_2) \in H \times W$  by a mask-like notation, where foreground classes  $c$  are set to one for spatial positions falling inside their bounding box. The background class  $C + 1$  is one at all left empty regions.

In semi-weakly supervised segmentation, we do not have full access on  $\mathcal{M}$  for all images  $x \in \mathcal{D}$  but each image is annotated with a mask, bounding box, image-level label or not at all. Along with its special case of having masks and raw images, i.e. semi-supervised segmentation, the target is to correctly infer the semantic mask for any given image.

##### 3.1.2 Supervision integration

Given an input feature map  $f \in \mathbb{R}^{d \times H \times W}$ , we refer to  $H$  and  $W$  as the spatial dimensions. To form a prediction  $p \in \mathbb{R}^{C \times H \times W}$  of the same spatial extent with  $C$  output dimensions, in semantic segmentation, most commonly  $f$  is transformed by  $C \ 1 \times 1$  convolutions. In our case, we transform  $f$  by a sequence of  $1 \times 1$  convolution, batch normalization [26] and ReLU non-linearity followed by a final  $1 \times 1$  convolution. We refer to this as an output-head of the model. In case we have multiple objective functions, we simply augment the model with additional output-heads on top of  $f$ , which is common practise in literature [8, 59].

Segmentation networks often employ encoder-decoder architectural designs, within which we later integrate auxiliary outputs (deep supervision). More specifically, these additional outputs operate on feature maps  $f_0, \dots, f_h$  produced in the *decoding process*. Here, we refer to  $f_0$  as the innermost feature map in the decoder, while  $f_h$  corresponds to the outermost feature map, on which standard segmentation models add an output-head for predictions. Given the spatial dimensions  $H_i \times W_i$  for the feature map  $f_i$ , we can generally assume that:

$$\forall_{i \in \{0, \dots, h\}; i < j} : H_i \leq H_j \ \& \ W_i \leq W_j . \quad (5)$$

In our experiments, we employ encoder-decoder models for which  $H_0 \ll H_h$  and  $W_0 \ll W_h$  hold. Computing predictions based on a certain feature map  $f_i$  will be referred to as  $\kappa_i(f_i) \in \mathbb{R}^{C \times H_i \times W_i}$ , with  $\kappa_i(\cdot)$  denoting an output-head.

##### 3.1.3 Supervision signals

Different modes of supervision often require different loss functions. The most common objective for training semantic segmentation models when pixel-wise masks are available is minimizing the *cross-entropy loss*:

$$\mathcal{L}_{CE}(f, m) = -\frac{1}{\Omega_1} \sum_{i,j,c=1}^{H,W,C} m^{c,i,j} \cdot \log(\alpha(\kappa(f))^{c,i,j}) , \quad (6)$$

where  $\alpha(\cdot)$  is the softmax along the first dimension and  $\Omega_1 = H \cdot W$ . In case of multi-label settings, the *binary cross-entropy loss* is commonly chosen:

$$BCE(o, t) = t \cdot \log(\sigma(o)) + (1 - t) \log(1 - \sigma(o)) \quad (7)$$

$$\mathcal{L}_{BCE}(f, m) = -\frac{1}{\Omega_2} \sum_{i,j,c=1}^{H,W,C} BCE(\kappa(f)^{c,i,j}, m^{c,i,j}) \quad (8)$$

With  $\Omega_2 = H \cdot W \cdot C$ , this loss most commonly uses a sigmoid normalization  $\sigma(\cdot)$ .

### 3.2. Multi-label Deeply Supervised Nets

**Parameter-efficient multi-label deep supervision.** In its inception, deep supervision was introduced for image classification [31], as such, having to identify *one* individual class corresponding to the input. Yet, semantic segmentation in itself has a different target: find the class for *each pixel* in the image. The problem we identify is the way deep supervision is commonly integrated into segmentation networks. Specifically, the challenge arises due to the mismatch in spatial dimensions between the full-scale ground-truth mask and the smaller spatial resolution within the network's feature maps (Equation 5).

Apart from [37] which use a lossy nearest interpolation, most work forces the network to learn an up-scaling to address the mismatch between low-resolution spatial features

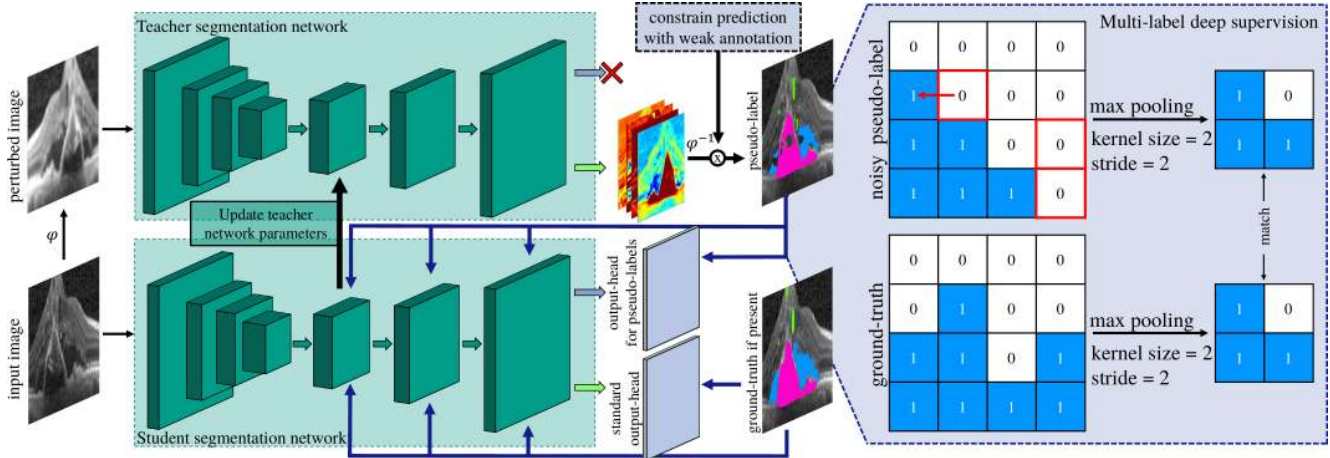


Figure 2. Our proposed approach combines pseudo-labeling via a mean-teacher with a novel perspective on deep supervision. By perturbing the input to the teacher and reversing geometric transformations in output-space, we streamline mean-teachers for segmentation. The new take on deep supervision, *i.e.* *Multi-label Deep Supervision* introduces a much needed smoothing effect for noisy pseudo-labels: At smaller scales, erroneous predictions (red) like small shifts or few missed pixel-classifications get smoothed out and match the unavailable mask.

and the ground-truth mask. Then, after up-scaling, a standard output-head transforms the features into a prediction with the same size as the mask. Consider the toy example with the feature map  $f_{small} \in \mathbb{R}^{d \times 10 \times 10}$  and a corresponding ground-truth mask  $m_{big} \in \mathbb{R}^{c \times 100 \times 100}$ . To enrich the segmentation model via deep supervision, previous approaches would for the  $d$  dimensional feature  $f_{small}^{:,x,y}$  at the spatial location  $(x, y) \in H_{small} \times W_{small}$  up-scale and infer a patch of size  $10 \times 10$  in the ground-truth. Not only is this a tremendously hard task, it is precisely the task the entirety of the decoder is trying to solve.

To summarize, we identify two shortcomings, (1) the network has to learn an up-scaling, at the cost of additional parameters and (2) intermediate features are burdened to model complex classification information and spatial relations in *output space* that we question to be useful in the decoding process, but presumably only serve as skip-connections for more stable gradients.

Instead, we propose to model each feature vector  $f_{:,x,y}$  at each location  $(x, y)$  in feature maps as patch-descriptors for their receptive field in the input image. Thus, we aim at ingraining the semantic information of all pixels contained in the receptive field of the patch-descriptors into the model. We argue that this can be achieved by enforcing a multi-label loss with a label containing *all semantic classes* present in the receptive field. It follows, that instead of up-scaling feature maps, we can simply down-scale the ground-truth mask to match the size of the feature maps at no cost of parameters and contain labels of *all classes confined* in the receptive fields to preserve semantic information.

By formulating ground-truth masks as binary tensors (2), this down-scaling process can be efficiently implemented, by applying max-pooling with a fitting kernel-size and

stride to match the feature map's spatial extent. The new down-scaled target  $m_i^* \in \mathbb{R}^{C \times H_i \times W_i}$ , therefore, contains the multi-label ground-truth  $m_i^{*,x,y}$  of the feature  $f_i^{:,x,y}$  (*i.e.* patch-descriptor) aggregated from all spatial positions within the feature's receptive field.

As shown in Fig. 3, we integrate such semantic-preserving, down-scaled multi-label ground-truths by applying output-heads on top of the feature maps and enforce:

$$\mathcal{L}(f_1, \dots, f_h, m_1^*, \dots, m_h^*) = \frac{1}{h} \sum_{k=1}^h \mathcal{L}_{BCE}(f_k, m_k^*). \quad (9)$$

We refer to this way of integrating supervision as *Multi-label Deep Supervision* since it respects the presence of multiple classes within the receptive field at each spatial location in the hierarchical, low-resolution feature maps. During training, we pair this loss for the outermost feature maps with a *standard output-head* using (6), and only use the latter for inference in a usual fashion.

**Self-taught deep supervision.** For semi-supervised segmentation, some authors consider generating pseudo-labels for unlabeled images to add noisy annotated samples. In a similar fashion, we propose to generate pseudo-labels but integrate them by providing *Multi-label Deep Supervision*. The motivation behind this is that pseudo-labels at full resolution will contain a significant amount of incorrect predictions. Yet, making use of *Multi-label Deep Supervision* we down-sample the pseudo-labels, naturally smoothing the noisy supervision signal. For earlier decoder outputs with smaller spatial resolution, the down-sampled pseudo-label will better match the unavailable actual ground-truth, as shown on the right-hand side of Fig. 2.

We introduce *Self-Taught Deep Supervision*, which generates a binary ground-truth tensor for an unlabeled sam-




Figure 3. Our proposed method of integrating deep supervision into the decoder of segmentation networks by down-sampling the pixel-mask and enforcing a multi-label classification loss.

ple in the following fashion: First the unlabeled image  $x_i$  is passed through the segmentation network to obtain a prediction  $p_i$ , which we transform into a pseudo-mask by attributing ones to classes with the highest probability at each location  $(x, y)$ . Using this pseudo-label, as described before, we enforce *Multi-label Deep Supervision*. We use separate output-heads on the outermost feature map for the ground-truth and the generated pseudo-labels, which turned out to be important. Thus, we integrate one output-head which is trained with standard cross-entropy loss and the clean labels, as well as a second output-head which is trained as part of Eq. (9) using pseudo-labels. For inferring pseudo-labels, we use the output-head trained with clean labels (see Fig. 2).

In case an image-level label  $g_i$  is available, we can further constrain the generated pseudo-label to the classes contained in  $g_i$ . In a similar fashion, an associated bounding box label  $b_i$  can constrain the pseudo-labels to lie within coarse regions. This results in a flexible integration of weakly-labeled images to improve pseudo-label quality.

**Mean-taught deep supervision.** In our final variant, we extend the *Self-Taught Deep Supervision* approach with a stronger pseudo-label generation. We generate more robust pseudo-labels by (1) enforcing consistent predictions with respect to perturbations and (2) using a teacher model, which is a combination of all models of previous iterations.

Mean-teachers (i.e. the exponential moving average over previous model parameters) were successfully introduced for semi-supervised classification [58] and saw some use in segmentation [19, 44, 48]. The idea is that better predictions can be obtained by maintaining a teacher model via continuously updating its parameters with the moving average of the student model and the previous teacher model. Hence, the teacher is not trained separately, but just updated by:

$$\theta_t^{\text{teacher}} = \alpha \cdot \theta_{t-1}^{\text{teacher}} + (1 - \alpha) \cdot \theta_t^{\text{student}}. \quad (10)$$

Here,  $\theta$  denotes either the parameters of the student- or teacher model, while the index denotes the training iteration

and  $\alpha$  the smoothing coefficient. The authors of [58] use the teacher to align its softmax predictions with the student’s given differently perturbed inputs using a mean squared error (MSE) loss. We extend this by leveraging the mean-teacher to provide pseudo-labels for our *Multi-label Deep Supervision* that the student network is trained with.

Previous adaptations to semantic segmentation either did not perturb the input image to the teacher [48] or used Cut-Mix variants [19, 44] as perturbations. Instead, we propose to perturb the input image to the teacher in a different way as the student input and still obtain matching outputs. To this intent, we transform the teacher’s input  $x_i$  by photo-metric (e.g. color jittering) and geometric perturbations (e.g. flipping)  $\varphi(x_i)$ . After the forward pass, we reverse only the geometric perturbations on the predictions via  $\varphi^{-1}(p_i)$ . This alteration is the crucial detail to make mean-teachers work for semantic segmentation using common, simple data augmentations as perturbations. As the mean-teacher is deemed to produce better predictions, we use its *standard output-head* for inference. See also Fig. 2 for visualizations.

## 4. Experiments

### 4.1. Data set

Our experiments are built on the publicly available RE-TOUCH data set [9] for retinal fluid segmentation. It is a collection of optical coherence tomography (OCT) volumes (stacks of b-scans) containing different retinal diseases. These volumes can be obtained by imaging tools of different vendors of which the data set features three: Spectralis, Cirrus, and Topcon. Generally, b-scans differ across manufacturers in appearance and will be considered separately in our experiments. Further, as part of our work lies in investigating the scarce-data scenario, our main experiments are carried out on the smallest of the three, Spectralis (49 b-scans per volume), while the performance on the remaining vendors (128 b-scans per volume) is evaluated to emphasize the generalization of our approach. The data set is fully annotated with pixel-wise labels for three types of retinal fluids: Intraretinal fluid, Subretinal fluid and Pigment Epithelial Detachments. For our experiments, we derive bounding boxes and image-level labels from the masks.

**Experimental setup.** As low-data scenarios entail sensitivity in the optimization process, we carry out each experimental configuration on 10 cross-validation splits. Each split is generated randomly and independent for each vendor: First, we randomly select 5 volumes for validation and 5 for testing. The remaining volumes (Spectralis: 14, Cirrus: 14, Topcon: 12) containing fluid annotations form the training set. In our experiments, we consider scenarios where 3, 6, 12, 24 labeled b-scans are present (corresponding to 1, 2, 4, 8 b-scans *per class*), therefore, we enumerate the scans in the training set and make sure that in an inter-

val of size 3 all diseases are present (as long as available). Thus, *e.g.* the first 6 images in a split contain each disease at least twice, which ensures that all models access the exact same annotations and the scenarios with more annotations subsume the smaller ones (*e.g.*  $\mathcal{D}_3 \subset \dots \subset \mathcal{D}_{24}$ ).

We consider two settings: (1) the models are provided with pixel-wise labeled masks ( $\mathcal{M}$ ) or (2) with coarse bounding boxes ( $\mathcal{B}$ ) as strongest annotation. In those scenarios, we train models having access to either one of these labels ( $\mathcal{M}$  or  $\mathcal{B}$ ), models having additionally access to the remaining unlabeled examples ( $\mathcal{U}$ , *i.e.* semi-supervised) or having access to the remaining b-scans with global image-level labels ( $\mathcal{G}$ , *i.e.* semi-weakly supervised).

**Evaluation metric.** Our evaluation follows the standard procedure for segmentation models. We calculate the prediction for all pixels in all testing images and then calculate the Intersection over Union (IoU) for all disease classes with the corresponding ground-truth pixels. Then, these disease-wise IoUs are averaged resulting in the mean IoU (mIoU). We carry out each experiment on all splits (10 trained models for one result), the final measures we report are the average mIoU and standard deviation over the splits.

## 4.2. Implementation details

**Pre-processing.** We employ the pre-processing of [3] on the individual b-scans which includes anisotropic filtering and warping the lower retinal edge onto a straight line. Finally, we resize all scans to the size  $200 \times 200$ .

**Network parameters.** For our encoder-decoder segmentation model, we use a standard UNet architecture [50], ubiquitous in medical image analysis. The UNet is augmented with batchnorm layers [26] and a total of four down-scaling and four up-scaling (bilinear interpolation) convolutional blocks. This way, the decoder has a total of four feature maps  $f_0, \dots, f_3$  increasing in spatial dimensions as well as a feature map  $f_4$  directly before the standard output  $1 \times 1$  convolution. As noted before, we exchange the output layers with output-heads as stated in Sec. 3.1.2, enabling each head to learn an individual non-linear transformation of the features. In case we use deep supervision, we add output-heads for all feature maps  $f_{\{0-4\}}$  and enforce the respective loss function. All experiments are carried out with the same basic training-hyperparameters: We use a batch size of 16 b-scans, which are transformed by a horizontal flip 50% of the time and by a random adjustment of brightness, contrast, hue and saturation by a factor of 0.0 – 0.1. Network parameters are initialized using Xavier initialization [21] and the optimization process spans over 100 epochs with SGD optimization and a momentum term of 0.9. The learningrate is adjusted once after 80 epochs from 0.01 to 0.001. Each model is evaluated every 10 epochs on the validation set, we perform early stopping and at the end of our experimentation phase evaluate the best model once on the testing set.

larger $\varphi$	inference	$\alpha$	MSE	validation (mIoU)
✓	teacher	0.5	–	$61.24 \pm 3.69$
✓	teacher	0.5	✓	<b><math>61.36 \pm 4.73</math></b>
✓	teacher	0.1	✓	$60.15 \pm 4.14$
✓	student	0.1	✓	$58.54 \pm 3.62$
–	student	0.1	✓	$58.26 \pm 4.27$
–	student	0.0	–	$57.80 \pm 4.68$

Table 1. Ablation for *Mean-Taught Deep Supervision* using 24 pixel-masks and the remaining image-level labels. Last line shows *Self-Taught Deep Supervision* performance for comparison.

## 4.3. Baselines and methods

Our first *Baseline* model is a UNet only using the strongest available supervision (mask or bounding box) as training signal, integrating it via the loss of (6) on  $f_4$ . To accommodate for the small amount of seen images (only  $\{3, 6, 12, 24\}$ ), we train the model for 10 times the epochs. All other models should outperform this model.

Next, we consider the semi-supervised scenario. Here, we leverage a UNet which uses (6) for available masks or bounding boxes and additionally the self-supervised invariant information clustering (IIC) loss of [27] for all labeled and unlabeled images on  $f_4$ . This *IIC Baseline* sets the bar for all models integrating unlabeled data.

For the semi-weakly supervised scenario with access to image-level labels  $g$ , we introduce a Multiple-Instance Learning (MIL) model. This lower bound leverages the weakly labeled images by pooling a feature map in the spatial dimension, classifying this pooled feature  $\bar{f}$  and enforcing  $\mathcal{L}_{BCE}(\bar{f}, g)$  from (8). Here, we use a pixel-wise cross-entropy loss for either masks or bounding boxes on  $f_4$  and for image-level labeled images we pool  $f_4$  via average pooling enforcing the MIL loss.

To show the effect of integrating deep supervision into segmentation models, we show performance of the previous two models but use all feature maps  $f_{\{0-4\}}$ . We term these models *Deeply Supervised IIC* and *Deeply Supervised MIL*. As we will see, the latter will provide an even stronger lower bound on semi-weakly supervised segmentation.

To grasp whether *Multi-Label Deep Supervision* is beneficial in training, we extend the *Baseline* UNet with our novel loss integration for all decoder feature maps  $f_{\{0-4\}}$ .

*Self-Taught Deep Supervision* can be applied to both the semi- and semi-weakly segmentation scenario. We integrate pseudo-labels as described in Sec. 3.2. When image-level labels are available, we refine the noisy pseudo-labels by zeroing out absent classes. As a natural extension of this approach, *Mean-Taught Deep Supervision* is setup similarly. Further, we benchmark the consistency-based approach of *Perone and Cohen-Adad* [48]. Re-implementing their approach as described led to diverging models. Thus, we modify it by using cross-entropy- instead of the DICE loss [42],

Method	$\mathcal{G}$	$\mathcal{U}$	3	6	12	24	Full Access
<b>Mask Supervision</b>							
Baseline [50]			14.80 ± 6.50	26.98 ± 7.83	35.39 ± 6.36	48.63 ± 5.17	62.09 ± 4.77
Multi-Label Deep Supervision (Ours)			<b>17.98 ± 8.20</b>	<b>32.92 ± 7.35</b>	<b>42.96 ± 6.71</b>	<b>52.68 ± 6.82</b>	<b>65.82 ± 4.64</b>
IIC Baseline <sup>8</sup> [27]	✓		<b>22.45 ± 9.36</b>	32.02 ± 7.23	41.48 ± 7.26	53.08 ± 6.13	65.16 ± 3.80
Deeply Supervised IIC <sup>8</sup>	✓		20.78 ± 8.83	31.39 ± 10.26	39.18 ± 6.94	50.10 ± 7.92	65.18 ± 3.85
<i>Perone and Cohen-Adad</i> <sup>10</sup> [48]	✓		16.17 ± 10.74	33.10 ± 10.24	45.80 ± 7.51	54.75 ± 5.96	65.49 ± 4.14
Self-Taught Deep Supervision (Ours)	✓		10.37 ± 8.29	28.62 ± 12.96	43.57 ± 9.97	56.11 ± 6.30	66.24 ± 4.67
Mean-Taught Deep Supervision <sup>10</sup> (Ours)	✓		16.31 ± 15.48	<b>35.17 ± 11.35</b>	<b>53.52 ± 8.72</b>	<b>58.84 ± 6.57</b>	<b>66.31 ± 4.66</b>
MIL Baseline	✓		15.44 ± 11.10	25.46 ± 8.57	41.34 ± 9.66	49.07 ± 8.20	61.50 ± 5.64
Deeply Supervised MIL	✓		20.02 ± 9.17	31.50 ± 8.88	44.29 ± 5.03	51.13 ± 3.93	62.04 ± 3.92
Self-Taught Deep Supervision (Ours)	✓		20.47 ± 8.62	36.40 ± 8.91	49.39 ± 9.95	59.29 ± 7.52	66.34 ± 3.81
Mean-Taught Deep Supervision <sup>10</sup> (Ours)	✓		<b>21.91 ± 13.49</b>	<b>42.14 ± 14.25</b>	<b>54.70 ± 9.26</b>	<b>60.45 ± 5.71</b>	<b>66.39 ± 4.29</b>
<b>Bounding Box Supervision</b>							
Baseline [50]			12.49 ± 4.28	18.32 ± 4.94	25.62 ± 3.08	29.55 ± 2.77	38.45 ± 4.44
Multi-Label Deep Supervision (Ours)			<b>14.59 ± 5.81</b>	<b>19.62 ± 6.21</b>	<b>27.89 ± 3.44</b>	<b>32.02 ± 4.78</b>	<b>38.66 ± 3.36</b>
IIC Baseline <sup>8</sup> [27]	✓		<b>15.40 ± 7.07</b>	18.15 ± 7.49	26.05 ± 6.00	30.07 ± 4.32	38.45 ± 4.65
Deeply Supervised IIC <sup>8</sup>	✓		12.77 ± 7.15	17.76 ± 6.26	<b>28.99 ± 4.60</b>	30.64 ± 3.05	38.81 ± 4.48
<i>Perone and Cohen-Adad</i> <sup>10</sup> [48]	✓		11.17 ± 7.41	<b>19.02 ± 8.46</b>	27.44 ± 5.81	31.72 ± 3.87	39.38 ± 3.56
Self-Taught Deep Supervision (Ours)	✓		5.14 ± 3.84	9.62 ± 7.35	24.47 ± 6.12	32.71 ± 3.56	<b>39.39 ± 3.63</b>
Mean-Taught Deep Supervision <sup>10</sup> (Ours)	✓		8.21 ± 3.96	14.28 ± 7.48	24.79 ± 5.79	<b>34.14 ± 3.10</b>	39.04 ± 4.15
MIL Baseline	✓		15.82 ± 6.55	16.95 ± 6.19	22.56 ± 4.56	26.48 ± 5.51	37.15 ± 4.06
Deeply Supervised MIL	✓		<b>17.14 ± 8.06</b>	20.18 ± 4.61	24.15 ± 4.95	29.12 ± 4.75	37.94 ± 3.35
Self-Taught Deep Supervision (Ours)	✓		16.04 ± 8.52	<b>22.15 ± 6.29</b>	28.63 ± 4.04	32.37 ± 3.75	<b>38.97 ± 3.59</b>
Mean-Taught Deep Supervision <sup>10</sup> (Ours)	✓		15.81 ± 8.59	21.97 ± 8.17	<b>29.83 ± 5.30</b>	<b>34.81 ± 3.62</b>	38.66 ± 4.73

Table 2. Results for a diverse variety of models using varying amounts (3,6,12,24,all) of annotations, i.e. ground-truth masks or coarse bounding boxes. Approaches with  $\mathcal{U}$  use additional unlabeled images, with  $\mathcal{G}$  global image labels (**best results in category bold**). Accuracy in average mIoU over 10 testing splits with standard deviation. Superscripts indicate smaller batch sizes due to memory constraints.

$\alpha = 0.5$ , no ramp up phase and balanced loss weighting.

#### 4.4. Ablation studies

For *Mean-Taught Deep Supervision*, we perform a sequence of ablation experiments on the validation set as shown in Table 1. First, we move from not using a teacher model (*Self-Taught Deep Supervision*) to using one. Further improvement is made when adding stronger perturbations  $\varphi(\cdot)$ . To this extent instead of using a factor 0.1, we use 0.4 for the photo-metric components described in Sec. 4.2. Inferring predictions with the teacher instead of the student for inference and tuning the hyperparameter  $\alpha$  led to the best configuration. To show that the mean squared error term taken over from [58] is not the vital component, we omit it and observe a minor decrease of 0.12% mIoU.

#### 4.5. Quantitative results

**Mask supervision.** In the upper half of Table 2, we study results of the diverse methods using pixel-wise annotations. The first observation we can make is that just one sample per class, i.e. the 3-scenario, is not sufficient as all methods struggle in producing meaningful segmentations. The step to 2 pixel-wise annotations per class sounds small, yet we can observe a giant leap of 20.23% to respectable 42.14%

mIoU for the *Mean-Taught Deep Supervision* model which integrates weak global labels. Throughout the scenarios 6, 12, and 24 annotated samples, our *Mean-Taught Deep Supervision* outperforms all other models consistently in both semi- and semi-weakly supervised segmentation. With just 24 masks (i.e. 5.78% of all annotations), our results come close to the full access *Baseline* that requires 415.5 masks (average over splits). Note the large margins between *Perone and Cohen-Adad* and our mean-teacher semi-supervised model, stretching up to 8.12%. Integrating our *Multi-label Deep Supervision* as compared to the standard Cross-Entropy loss (*Baseline*) increases the performance by 3.18%, 5.94%, 7.57%, 4.05% and even fully supervised by 3.73% showcasing the strength of our novel deep supervision mechanism.

**Bounding box supervision.** In the lower half of Table 2, we exchange pixel-wise masks with coarsely annotated bounding boxes. As expected, the upper bound of full access at 39.39% is limited by the nature of weak annotation. Surprisingly, our *Mean-Taught Deep Supervision* model having access to just 6 *pixel-wise masks* and additional global label (in the top half of Table 2) already outperforms all models with full bounding box access. Further, most of the best performing models build on *Multi-label Deep Supervision*.

Method	$\mathcal{U}$	6	12	24	Full Access
<b>Cirrus</b>					
Baseline [50]		$12.31 \pm 5.41$	$19.43 \pm 8.00$	$30.10 \pm 9.34$	$48.92 \pm 11.94$
Multi-label Deep Supervision (Ours)		<b><math>15.99 \pm 6.87</math></b>	<b><math>25.12 \pm 8.58</math></b>	$33.53 \pm 9.44$	$50.47 \pm 10.84$
<i>Perone and Cohen-Adad</i> <sup>10</sup> [48]	✓	$12.36 \pm 6.12$	$24.99 \pm 6.49$	$33.79 \pm 10.15$	$49.75 \pm 12.87$
Mean-Taught Deep Supervision <sup>10</sup> (Ours)	✓	$9.18 \pm 8.53$	$23.33 \pm 7.37$	<b><math>35.82 \pm 11.40</math></b>	<b><math>51.24 \pm 10.94</math></b>
<b>Topcon</b>					
Baseline [50]		$14.79 \pm 9.34$	$21.19 \pm 11.57$	$27.61 \pm 10.31$	$42.22 \pm 10.42$
Multi-label Deep Supervision (Ours)		<b><math>18.20 \pm 10.48</math></b>	$20.92 \pm 13.02$	$33.71 \pm 11.92$	<b><math>45.85 \pm 10.32</math></b>
<i>Perone and Cohen-Adad</i> <sup>10</sup> [48]	✓	$15.26 \pm 12.74$	$21.88 \pm 12.48$	$27.67 \pm 13.81$	$41.43 \pm 8.18$
Mean-Taught Deep Supervision <sup>10</sup> (Ours)	✓	$14.39 \pm 11.19$	<b><math>23.92 \pm 15.25</math></b>	<b><math>33.87 \pm 8.25</math></b>	$42.70 \pm 10.97$

Table 3. Segmentation accuracy of different methods on Cirrus and Topcon OCT vendors in average mIoU over 10 testing splits.

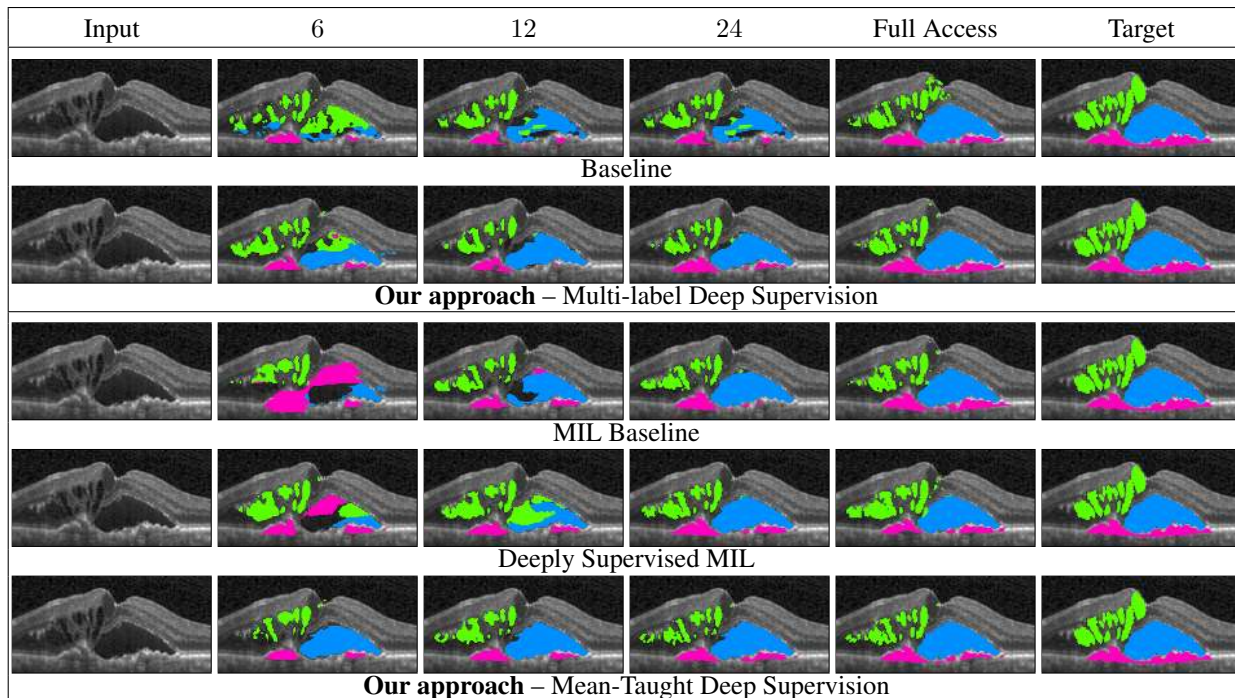


Figure 4. Segmentation progression when increasing the number of pixel-wise annotated masks from 6 to 12, 24, and full access. The upper two methods leverage only pixel-wise masks while the remaining three methods have access to weak annotations (image-level labels).

**Results for other vendors.** Due to the immense resources required to carry out the full evaluation as in Table 2, i.e. training a total amount of 1100 models, we show that our approaches generalize to other vendors (Cirrus and Topcon) on a smaller subset of experiments in Table 3.

#### 4.6. Qualitative results

In Fig. 4, we show segmentation results as more annotations are added for the *Baseline* and *Multi-label Deep Supervision* models as well as some semi-weakly supervised results. Comparing the first two rows, we observe the advantageous effect of our deep supervision technique, especially in the low-data setting. Similarly, when comparing our deep supervision in *Mean-Taught* models (bottom row) to the rest, we directly see in column two and three that

our models nicely segment the fluid regions, even with extremely few samples while other methods fail severely.

## 5. Conclusion

We introduced a new segmentation approach that allows for learning from different annotation types (pixel-wise labels, bounding boxes, global labels, unlabeled images). We showed that a key ingredient is the deep supervision principle, which we adapted to efficiently introduce supervision signals even in intermediate layers with a small spatial resolution. Furthermore, we introduced a student-teacher model to stabilize training and increase generalization. Our experiments for a challenging medical application demonstrated the flexibility of our approach and the power of the simple concept of deep supervision for semi-weakly learning.



## References

- [1] Nabila Abraham and Naimul Mefraz Khan. A novel focal tversky loss function with improved attention u-net for lesion segmentation. In *2019 IEEE 16th International Symposium on Biomedical Imaging*, pages 683–687. IEEE, 2019. 2
- [2] Saeedeh Afshari, Aïcha BenTaieb, Zahra Mirikharaji, and Ghassan Hamarneh. Weakly supervised fully convolutional network for pet lesion segmentation. In *Medical Imaging 2019: Image Processing*, volume 10949, page 109491K. International Society for Optics and Photonics, 2019. 2
- [3] Stefanos Apostolopoulos, Carlos Ciller, Sandro De Zanet, Sebastian Wolf, and Raphael Sznitman. Retinet: Automatic amd identification in oct volumetric data. *Investigative Ophthalmology & Visual Science*, 58(8):387–387, 2017. 6
- [4] Pablo Arbeláez, Jordi Pont-Tuset, Jonathan T Barron, Ferran Marques, and Jitendra Malik. Multiscale combinatorial grouping. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 328–335, 2014. 2
- [5] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017. 2
- [6] Wenjia Bai, Ozan Oktay, Matthew Sinclair, Hideaki Suzuki, Martin Rajchl, Giacomo Tarroni, Ben Glocker, Andrew King, Paul M Matthews, and Daniel Rueckert. Semi-supervised learning for network-based cardiac mr image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 253–260. Springer, 2017. 2
- [7] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *European conference on computer vision*, pages 549–565. Springer, 2016. 1
- [8] Benjamin Bischke, Patrick Helber, Joachim Folz, Damian Borth, and Andreas Dengel. Multi-task learning for segmentation of building footprints with deep neural networks. In *2019 IEEE International Conference on Image Processing*, pages 1480–1484. IEEE, 2019. 3
- [9] Hrvoje Bogunović, Freerk Venhuizen, Sophie Klimscha, Stefanos Apostolopoulos, Alireza Bab-Hadiashar, Ulas Bagci, Mirza Faisal Beg, Loza Bekalo, Qiang Chen, Carlos Ciller, Karthik Gopinath, Amirali K. Gostar, Kiwan Jeon, Zexuan Ji, Sung Ho Kang, Dara D. Koozekanani, Donghuan Lu, Dustin Morley, Keshab K. Parhi, Hyoung Suk Park, Abdolreza Rashno, Marinko Sarunic, Saad Shaikh, Jayanthi Sivaswamy, Ruwan Tennakoon, Shivin Yadav, Sandro De Zanet, Sebastian M. Waldstein, Bianca S. Gerendas, Caroline Klaver, Clara I. Sánchez, and Ursula Schmidt-Erfurth. RETOUCH - The Retinal OCT Fluid Detection and Segmentation Benchmark and Challenge. *IEEE Transactions on Medical Imaging*, 38(8):1858–1874, Aug 2019. 5
- [10] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8991–9000, 2020. 1, 2
- [11] Kaiqiang Chen, Michael Weinmann, Xian Sun, Menglong Yan, Stefan Hinz, Boris Jutzi, and Martin Weinmann. Semantic segmentation of aerial imagery via multi-scale shuffling convolutional neural networks with deep supervision. *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences*, 4(1), 2018. 2
- [12] 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. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 2
- [13] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision*, pages 801–818, 2018. 2
- [14] Junsuk Choe, Seong Joon Oh, Seungho Lee, Sanghyuk Chun, Zeynep Akata, and Hyunjung Shim. Evaluating weakly supervised object localization methods right. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3133–3142, 2020. 2
- [15] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 1, 2
- [16] Jifeng Dai, Kaiming He, and Jian Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1635–1643, 2015. 2
- [17] Qi Dou, Lequan Yu, Hao Chen, Yueming Jin, Xin Yang, Jing Qin, and Pheng-Ann Heng. 3d deeply supervised network for automated segmentation of volumetric medical images. *Medical image analysis*, 41:40–54, 2017. 2
- [18] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111(1):98–136, 2015. 2
- [19] Geoff French, Timo Aila, Samuli Laine, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, high-dimensional perturbations. *arXiv preprint arXiv:1906.01916*, 2019. 2, 5
- [20] Huazhu Fu, Jun Cheng, Yanwu Xu, Damon Wing Kee Wong, Jiang Liu, and Xiaochun Cao. Joint optic disc and cup segmentation based on multi-label deep network and polar transformation. *IEEE transactions on medical imaging*, 37(7):1597–1605, 2018. 2
- [21] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010. 6
- [22] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional net-

- works. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 2
- [23] Zilong Huang, Xinggang Wang, Jiasi Wang, Wenyu Liu, and Jingdong Wang. Weakly-supervised semantic segmentation network with deep seeded region growing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7014–7023, 2018. 1, 2
- [24] Wei-Chih Hung, Yi-Hsuan Tsai, Yan-Ting Liou, Yen-Yu Lin, and Ming-Hsuan Yang. Adversarial learning for semi-supervised semantic segmentation. *arXiv preprint arXiv:1802.07934*, 2018. 1, 2
- [25] Mostafa S Ibrahim, Arash Vahdat, Mani Ranjbar, and William G Macready. Semi-supervised semantic image segmentation with self-correcting networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12715–12725, 2020. 2
- [26] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 3, 6
- [27] Xu Ji, João F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9865–9874, 2019. 6, 7
- [28] Hoel Kervadec, Jose Dolz, Shanshan Wang, Eric Granger, and Ismail Ben Ayed. Bounding boxes for weakly supervised segmentation: Global constraints get close to full supervision. *arXiv preprint arXiv:2004.06816*, 2020. 2
- [29] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 876–885, 2017. 1, 2
- [30] Alexander Kolesnikov and Christoph H Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *European conference on computer vision*, pages 695–711. Springer, 2016. 2
- [31] Chen-Yu Lee, Saining Xie, Patrick Gallagher, Zhengyou Zhang, and Zhuowen Tu. Deeply-supervised nets. In *Artificial intelligence and statistics*, pages 562–570, 2015. 2, 3
- [32] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5267–5276, 2019. 1, 2
- [33] Bingbing Li, Chengdong Wu, Jianning Chi, Xiaosheng Yu, and Gang Wang. A deeply supervised convolutional neural network for brain tumor segmentation. In *2020 39th Chinese Control Conference (CCC)*, pages 6262–6267. IEEE, 2020. 2
- [34] Xiaomeng Li, Lequan Yu, Hao Chen, Chi-Wing Fu, and Pheng-Ann Heng. Semi-supervised skin lesion segmentation via transformation consistent self-ensembling model. *arXiv preprint arXiv:1808.03887*, 2018. 2
- [35] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1925–1934, 2017. 2
- [36] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2
- [37] Wen Liu, Yankui Sun, and Qingge Ji. Mdan-unet: Multi-scale and dual attention enhanced nested u-net architecture for segmentation of optical coherence tomography images. *Algorithms*, 13(3):60, 2020. 3
- [38] Yun Liu, Ming-Ming Cheng, Xinyu Zhang, Guang-Yu Nie, and Meng Wang. Dna: Deeply-supervised nonlinear aggregation for salient object detection. *arXiv preprint arXiv:1903.12476*, 2019. 2
- [39] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 2
- [40] Dimitrios Marmanis, Jan D Wegner, Silvano Galliani, Konrad Schindler, Mihai Datcu, and Uwe Stilla. Semantic segmentation of aerial images with an ensemble of cnns. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2016, 3:473–480, 2016. 2
- [41] Matej Mikulic. Number of examinations with magnetic resonance imaging (mri) in selected countries as of 2019. <https://www.statista.com/statistics/271470/mri-scanner-number-of-examinations-in-selected-countries/>. Accessed: 2020-11-16. 1
- [42] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision*, pages 565–571. IEEE, 2016. 6
- [43] Sudhanshu Mittal, Maxim Tatarchenko, and Thomas Brox. Semi-supervised semantic segmentation with high-and low-level consistency. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 2
- [44] Viktor Olsson, Wilhelm Tranheden, Juliano Pinto, and Lennart Svensson. Classmix: Segmentation-based data augmentation for semi-supervised learning. *arXiv preprint arXiv:2007.07936*, 2020. 2, 5
- [45] Yassine Ouali, Céline Hudelot, and Myriam Tami. Semi-supervised semantic segmentation with cross-consistency training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12674–12684, 2020. 2
- [46] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1742–1750, 2015. 2
- [47] Deepak Pathak, Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional multi-class multiple instance learning. *arXiv preprint arXiv:1412.7144*, 2014. 1, 2
- [48] Christian S Perone and Julien Cohen-Adad. Deep semi-supervised segmentation with weight-averaged consistency

- targets. In *Deep learning in medical image analysis and multimodal learning for clinical decision support*, pages 12–19. Springer, 2018. 2, 5, 6, 7, 8
- [49] Martin Rajchl, Matthew CH Lee, Ozan Oktay, Konstantinos Kamnitsas, Jonathan Passerat-Palmbach, Wenjia Bai, Mellisa Damodaram, Mary A Rutherford, Joseph V Hajnal, Bernhard Kainz, et al. Deepcut: Object segmentation from bounding box annotations using convolutional neural networks. *IEEE transactions on medical imaging*, 36(2):674–683, 2016. 2
- [50] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015. 2, 6, 7, 8
- [51] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. ” grabcut” interactive foreground extraction using iterated graph cuts. *ACM transactions on graphics (TOG)*, 23(3):309–314, 2004. 2
- [52] Thomas Schlegl, Sebastian M Waldstein, Wolf-Dieter Vogl, Ursula Schmidt-Erfurth, and Georg Langs. Predicting semantic descriptions from medical images with convolutional neural networks. In *International Conference on Information Processing in Medical Imaging*, pages 437–448. Springer, 2015. 2
- [53] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. 2
- [54] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. 2014. 2
- [55] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3136–3145, 2019. 1, 2
- [56] Nasim Souly, Concetto Spampinato, and Mubarak Shah. Semi supervised semantic segmentation using generative adversarial network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5688–5696, 2017. 2
- [57] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. 2
- [58] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*, pages 1195–1204, 2017. 2, 5, 7
- [59] Marvin Teichmann, Michael Weber, Marius Zoellner, Roberto Cipolla, and Raquel Urtasun. Multinet: Real-time joint semantic reasoning for autonomous driving. In *2018 IEEE Intelligent Vehicles Symposium*, pages 1013–1020. IEEE, 2018. 3
- [60] Jasper RR Uijlings, Koen EA Van De Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *International journal of computer vision*, 104(2):154–171, 2013. 2
- [61] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12275–12284, 2020. 2
- [62] Yan Xu, Yang Li, Mingyuan Liu, Yipei Wang, Maode Lai, I Eric, and Chao Chang. Gland instance segmentation by deep multichannel side supervision. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 496–504. Springer, 2016. 2
- [63] Hyun-Lim Yang, Jong Jin Kim, Jong Ho Kim, Yong Koo Kang, Dong Ho Park, Han Sang Park, Hong Kyun Kim, and Min-Soo Kim. Weakly supervised lesion localization for age-related macular degeneration detection using optical coherence tomography images. *PLoS one*, 14(4):e0215076, 2019. 2
- [64] Xin Yang, Lequan Yu, Shengli Li, Xu Wang, Na Wang, Jing Qin, Dong Ni, and Pheng-Ann Heng. Towards automatic semantic segmentation in volumetric ultrasound. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 711–719. Springer, 2017. 2
- [65] 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 *Proceedings of the IEEE International Conference on Computer Vision*, pages 6023–6032, 2019. 2
- [66] Yishuo Zhang and Albert CS Chung. Deep supervision with additional labels for retinal vessel segmentation task. In *International conference on medical image computing and computer-assisted intervention*, pages 83–91. Springer, 2018. 2
- [67] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. 2
- [68] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016. 2
- [69] Yuyin Zhou, Lingxi Xie, Elliot K Fishman, and Alan L Yuille. Deep supervision for pancreatic cyst segmentation in abdominal ct scans. In *International conference on medical image computing and computer-assisted intervention*, pages 222–230. Springer, 2017. 2
- [70] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pages 3–11. Springer, 2018. 2