

ADAPTIVE EARLY-LEARNING CORRECTION FOR SEGMENTATION FROM NOISY ANNOTATIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

Deep learning in the presence of noisy annotations has been studied extensively in classification, but much less in segmentation tasks. In this work, we study the learning dynamics of deep segmentation networks trained on inaccurately annotated data. We observe a phenomenon that has been previously reported in the context of classification: the networks tend to first fit the clean pixel-level labels during an “early-learning” phase, before eventually memorizing the false annotations. However, in contrast to classification, memorization in segmentation does not arise simultaneously for all semantic categories. Inspired by these findings, we propose a new method for segmentation from noisy annotations with two key elements. First, we detect the beginning of the memorization phase separately for each category during training. This allows us to adaptively correct the noisy annotations in order to exploit early learning. Second, we incorporate a regularization term that enforces consistency across scales to boost robustness against annotation noise. Our method outperforms standard approaches on a medical-imaging segmentation task where noises are synthesized to mimic human annotation errors. It also provides robustness to realistic noisy annotations present in weakly-supervised semantic segmentation, achieving state-of-the-art results on PASCAL VOC 2012.

1 INTRODUCTION

Semantic segmentation is a fundamental problem in computer vision. The goal is to assign a label to each pixel in an image, indicating its semantic category. Deep learning models based on convolutional neural networks (CNNs) achieve state-of-the-art performance (Chen et al., 2017; Zhao et al., 2017; Sandler et al., 2018; Wang et al., 2018). These models are typically trained in a supervised fashion, which requires pixel-level annotations. Unfortunately, gathering pixel-level annotations is very costly, and may require significant domain expertise in some applications (Trembl et al., 2016; Havaei et al., 2017; Schlemper et al., 2018; Liu et al., 2021). Furthermore, annotation noise is inevitable in some applications. For example, in medical imaging, segmentation annotation may suffer from inter-reader annotation variations (Kats et al., 2019; Zhang et al., 2020c). Learning to perform semantic segmentation from noisy annotations is thus an important topic in practice.

Prior works on learning from noisy labels focus on classification tasks (Liu et al., 2020; Tanaka et al., 2018; Xia et al., 2021). There are comparatively fewer works on segmentation, where existing works focus on designing noise-robust network architecture (Wang et al., 2020a) or incorporating domain specific prior knowledge (Shu et al., 2019). We instead focus on improving the performance in a more general perspective by studying the learning dynamics. We observe that the networks tend to first fit the clean annotations during an “early-learning” phase, before eventually memorizing the false annotations, thus jeopardizing generalization performance. This phenomenon has been reported in the context of classification (Liu et al., 2020). However, this phenomenon in semantic segmentation differs significantly from its counterpart in classification in the following ways:

- The noise in segmentation labels is often spatially dependent. Therefore, it is beneficial to leverage spatial information during training.
- In semantic segmentation, early learning and memorization do not occur simultaneously for all semantic categories due to pixel-wise imbalanced labels. Previous methods (Li et al., 2020; Liu et al., 2020) in noisy label classification often assume class balanced data and thus either detecting or handling wrong labels for different classes at the same time.

- The annotation noise in semantic segmentation can be ubiquitous (all examples have some errors) while the state-of-the-art methods in classification (Li et al., 2020; Zhou et al., 2020; Liu et al., 2020) assume that some samples are completely clean.

Inspired by these observations, we propose a new method, ADELE (ADaptive Early-Learning corrEction), that is designed for segmentation from noisy annotations. Our method detects the beginning of the memorization phase by monitoring the Intersection over Union (IoU) curve for each category during training. This allows it to adaptively correct the noisy annotations in order to exploit early-learning for individual classes. We also incorporate a regularization term to promote spatial consistency, which further improves the robustness of segmentation networks to annotation noise.

To verify the effectiveness of our method, we consider a setting where noisy annotations are synthesized and controllable. We also consider a practical setting – Weakly-Supervised Semantic Segmentation (WSSS), which aims to perform segmentation based on weak supervision signals, such as image-level labels (Kolesnikov & Lampert, 2016; Wei et al., 2016), bounding box (Dai et al., 2015; Song et al., 2019), or scribbles (Lin et al., 2016). We focus on a popular pipeline in WSSS. This pipeline consists of two steps (See Figure 1). First, a classification model is used to generate pixel-level annotations. This is often achieved by applying variations of Class Activation Maps (CAM) (Zhou et al., 2016) combined with post-processing techniques (Krähenbühl & Koltun, 2011; Ahn & Kwak, 2018). Second, these pixel-level annotations are used to train a segmentation model (such as deeplabv1 Chen et al. (2014)). Generated by a classification model, the pixel-wise annotations supplied to the segmentation model are inevitably noisy, thus the second step is indeed a noisy segmentation problem. We therefore apply ADELE to the second step. In summary, our main contributions are:




Figure 1: A prevailing pipeline for training WSSS. We aim to improve the segmentation model from noisy annotations.

Generated by a classification model, the pixel-wise annotations supplied to the segmentation model are inevitably noisy, thus the second step is indeed a noisy segmentation problem. We therefore apply ADELE to the second step. In summary, our main contributions are:

- We analyze the behavior of segmentation networks when trained with noisy pixel-level annotations. We show that the training dynamics can be separated into an early-learning and a memorization stage in segmentation with annotation noise. Crucially, we discover that these dynamics differ across each semantic category.
- We propose a novel approach (ADELE) to perform semantic segmentation with noisy pixel-level annotations, which exploits early learning by adaptively correcting the annotations using the model output.
- We evaluate ADELE on the thoracic organ segmentation task where annotations are corrupted to resemble human errors. ADELE is able to avoid memorization, outperforming standard baselines. We also perform extensive experiments to study ADELE on various types and levels of noises.
- ADELE achieves the state of the art on PASCAL VOC 2012 for WSSS. We show that ADELE can be combined with several different existing methods for extracting pixel-level annotations (Ahn & Kwak, 2018; Wang et al., 2020b; Fan et al., 2020) in WSSS, consistently improving the segmentation performance by a substantial margin.

2 METHODOLOGY

2.1 EARLY LEARNING AND MEMORIZATION IN SEGMENTATION FROM NOISY ANNOTATIONS

In a typical classification setting with label noise, a subset of the images are incorrectly labeled. It has been observed in prior works that deep neural networks tend to first fit the training data with clean labels during an *early-learning* phase, before eventually *memorizing* the examples with incorrect labels (Arpit et al., 2017; Liu et al., 2020). Here, we show that this phenomenon also occurs in segmentation when the available pixel-wise annotations are noisy (i.e. some of the pixels are incorrect). We consider two different problems. First, segmentation in medical imaging, where annotation noise is mainly due to human error. Second, the annotation noise in weakly-supervised




Figure 2: We visualize the effect of early learning (IoU_{el} , green curves) and memorization (IoU_m , red curves) on segmentation models trained with (solid lines) and without (dashed lines) ADELE for each foreground category of a medical dataset SegThor Lambert et al. (2020). The model is a UNet trained with noisy annotations that mimic human errors. IoU_{el} is the IOU between the model output and the ground truth computed over the incorrectly-labeled pixels. IoU_m is the IOU between the model output and the incorrect annotations. For all classes, IoU_m increases substantially as training proceeds because the model gradually memorizes the incorrect annotations. This occurs at **different speeds** for different categories. In contrast, IoU_{el} first increases during an early-learning stage where the model learns to correctly segment the incorrectly-labeled pixels, but eventually decreases as memorization occurs. Like memorization, early-learning also happens at varying speeds for the different semantic categories. See Figure 8 in Appendix for the plot on PASCAL VOC.

semantic segmentation due to the bias of classification models, as they mostly focus on discriminative regions, and the post-processing errors may result in systematic over or under segmentation.

Given noisy annotations for which **we know the ground truth**, we can quantify the early-learning and memorization phenomena by analyzing the model output on the pixels that are incorrectly labeled:

- **early learning IoU_{el}** : We quantify early learning using the overlap (measured in terms of the Intersection over Union (IoU) metric) between the outputs and the corresponding ground truth label on the pixels that are incorrectly labeled, denoted by IoU_{el} .
- **memorization IoU_m** : We quantify memorization using the overlap (measured in IoU) between the CNN outputs and the incorrect labels, denoted by IoU_m .

Figure 2 demonstrates the phenomena of early-learning and memorization on a randomly corrupted CT-scan segmentation dataset (SegTHOR Lambert et al. (2020)). We analyze the learning curve on the incorrectly-annotated pixels during the training process. The plots show the IoU_m (dashed red line) and IoU_{el} (dashed green line) at different training epochs. For all classes, the IoU between the output and the incorrect labels (IoU_m) increases substantially as training proceeds because the model gradually memorizes the incorrect annotations. This memorization process occurs at **varying speeds** for different semantic categories (compare *heart* and *Aorta*s with *Trachea*s or *Esophagus* in the SegThor dataset). The IoU_{el} follows a completely different trajectory: it first increases during an early-learning stage where the model learns to correctly segment the incorrectly-labeled pixels, but eventually decreases as memorization occurs (for the WSSS dataset, we observe a very similar phenomenon shown in Figure 9 in the Appendix). Like memorization, early-learning also happens at varying speeds for the different semantic categories.

Figure 3 illustrates the effect of early learning and memorization on the model output. In the medical-imaging application, the noisy annotations (third column) are synthesized to resemble human annotation errors which either miss or encompass the ground truth regions (compare to second column). Right after early learning, these regions are identified by the segmentation model (fourth column), but after memorization the model overfits to the incorrect annotations and *forgets* how to segment these regions correctly (fifth column). Similar effects are observed in WSSS, in which the noisy annotations generated by the classification model are missing some object regions, perhaps because they are not particularly discriminative (e.g. the body of the dog, cat and people in the first, second, and fourth row respectively, or the upper half of the bus in the third row). The segmentation model first identify these regions but eventually overfits to the incorrect annotations.

Our goal in this work is to modify the training of segmentation models on noisy annotations in order to prevent memorization. This is achieved by combining two strategies described in the next two

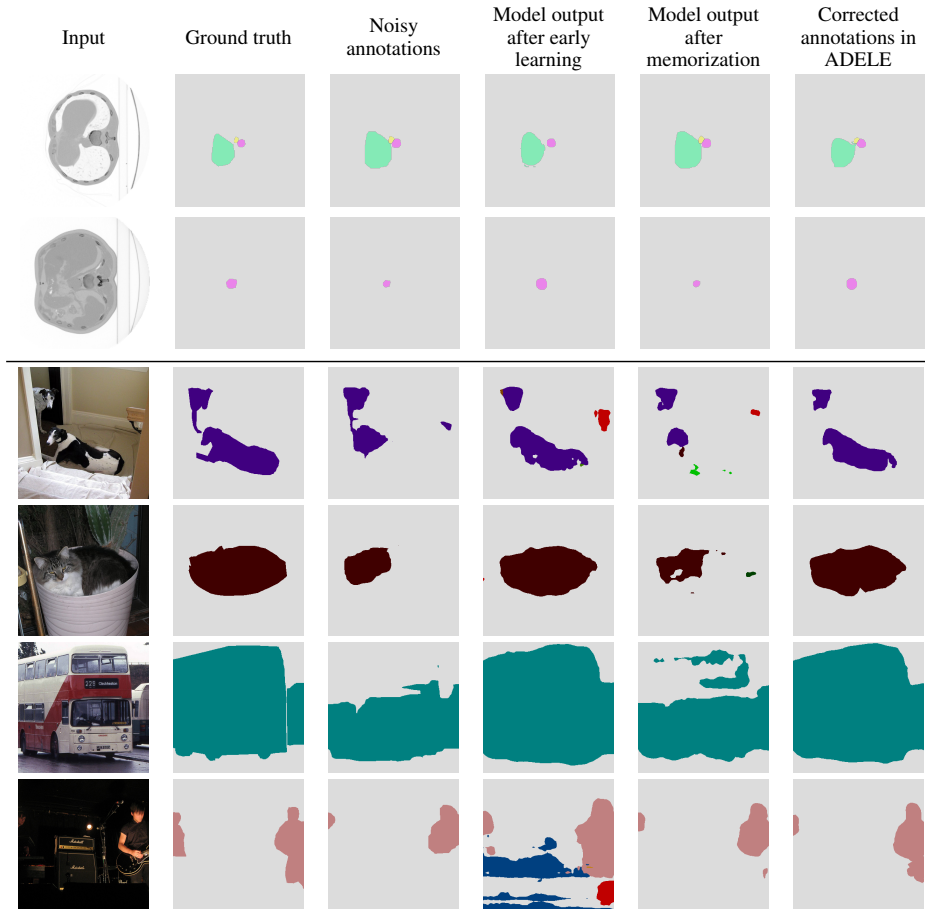


Figure 3: Visual examples illustrating the early-learning and memorization phenomena. For several images in a medical dataset Segthor (Lambert et al., 2020) (top two rows) and the WSSS dataset VOC 2012 (Everingham et al., 2015) (bottom four rows), we show the ground-truth annotations (second column), noisy annotations (third column) obtained by a synthetic corruption process for the medical data and by the classification-based SEAM (Wang et al., 2020b) model for WSSS, the output of a model segmentation model trained on the noisy annotations after early learning (fourth column), and the output of the same model after memorization (fifth column). The model for the medical dataset is a UNet. The WSSS model is a standard DeepLab-v1 network trained with the SEAM annotations. As suggested by the graphs in Figure 2 after early learning the model corrects some of the annotation errors, but these appear again after memorization. ADELE is able to correct the labels leveraging the early learning output, thereby avoiding memorization (sixth column). We set the background color to light gray for ease of visualization.

sections. Figure 2 and Figure 3 shows that the resulting method substantially mitigates memorization (solid red lines) and promotes continued learning beyond the early-learning stage (solid green lines).

2.2 ADAPTIVE LABEL CORRECTION BASED ON EARLY-LEARNING

The early-learning phenomenon described in the previous section suggests a strategy to enhance segmentation models: correcting the annotations using the model output. Similar ideas have inspired works in classification with noisy labels (Tanaka et al., 2018; Yi & Wu, 2019; Reed et al., 2015; Liu et al., 2020). However, different from the classification task where the noise is mainly sample-wise, **the annotation noise is ubiquitous across examples and distributed in a pixel-wise manner.** There is a key consideration for this approach to succeed: the annotations cannot be corrected too soon, because this degrades their quality. Determining when to correct the pixel-level annotations using the model output is challenging for two reasons:




Figure 4: Illustration of the proposed curve fitting method to decide when to begin label correction in ADELE (Results on SegThor). On the left, we plot the IoU between the model predictions and the initial noisy annotations for the same model used in Figures 2 and 3 and the corresponding fit with the parametric model in Equation 1. The label correction beginning iteration is based on the relative slope change of the fitted curve. The center image shows the label correction times for different semantic categories, showing that they are quite different. On the right graph, the green line shows the IoU_{el} for a given category *Heart*. The IoU_{el} equals the IoU between the model output and the ground truth computed over the incorrectly-labeled pixels, and therefore quantifies early-learning. The label correction begins close to the end of the early-learning phase, as desired. More result in section A.1 in Appendix shows that this also occurs for VOC 2012.

- Correcting all classes at the same time can be sub-optimal.
- During training, we do not have access to the performance of the model on ground-truth annotations (otherwise we would just use them to train the model in the first place!).

To overcome these challenges we propose to update the annotations corresponding to different categories at different times by detecting when early learning has occurred and memorization is about to begin using the training performance of the model.

In our experiments, we observe that the segmentation performance on the training set (measured by the IoU between the model output and the noisy annotations) improves rapidly during early learning, and then much more slowly during memorization (see the rightmost graph in Figure 4). We propose to use this deceleration to decide when to update the noisy annotations. To estimate the deceleration we first fit the following exponential parametric model to the training IoU using least squares:

$$f(t) = a \left(1 - e^{-b \cdot t^c} \right), \quad (1)$$

where t represents training time and $0 < a \leq 1$, $b \geq 0$, and $c \geq 0$ are fitting parameters. Then we compute the derivative $f'(t)$ of the parametric model with respect to t at $t = 1$ and at the current iteration.¹ For each semantic category, the annotations are corrected when the relative change in derivative is above a certain threshold r , i.e. when

$$\frac{|f'(1) - f'(t)|}{|f'(1)|} > r, \quad (2)$$

which we set to 0.9, and at every subsequent epoch. We only correct annotations for which the model output has confidence above a certain threshold τ , which we set to 0.8. As shown in Table 2, adaptive label correction based on early learning improves segmentation models in the medical-imaging applications and WSSS, both on its own and in combination with multiscale-consistency regularization. Figure 3 shows some examples of annotation corrections (rightmost column).

2.3 MULTISCALE CONSISTENCY

As we previously mentioned, model outputs after early-learning are used to correct noisy annotations. Therefore, the quality of model outputs is crucial for the effectiveness of the proposed method. Following a common procedure that has shown to result in more accurate segmentation from the outputs (Lin et al., 2018; Yang et al., 2018), we average model outputs corresponding to multiple rescaled copies of inputs to form the final segmentation, and use them to correct labels. Furthermore, we incorporate a regularization that imposes consistency of the outputs across multi-scales and is able to make averaged outputs more accurate (See the right graph of Figure 5). This idea is inspired

¹The derivative is given by $f'(t) = abce^{-bt^c} t^{c-1}$.

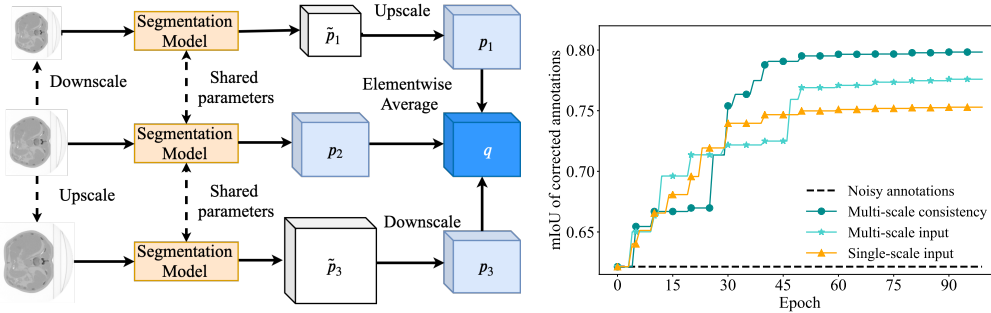


Figure 5: **Left:** In the proposed multiscale-consistency regularization, rescaled copies of the same input (here upscaled $\times 1.5$ and downscaled $\times 0.7$) are fed into the segmentation model. The outputs (\tilde{p}_1 , p_2 and \tilde{p}_3) are rescaled to have the same dimensionality (p_1 , p_2 and p_3). Regularization promotes consistency between these rescaled outputs and their elementwise average q . **Right:** Multi-scale consistency regularization leads to more accurate corrected annotations (results on SegThor, results for VOC 2012 can be found in Figure 12).

by consistency regularizations, a popular concept in the semi-supervised learning literature (Laine & Aila, 2018; Tarvainen & Valpola, 2017; Miyato et al., 2018; Berthelot et al., 2019; Sohn et al., 2020; French et al., 2019; Kim et al., 2020) that encourages the model to produce predictions that are robust to arbitrary semantic-preserving spatial perturbations.

To be more specific, let s be the number of scaling operations. In our experiments we set $s = 3$ (downscaling $\times 0.7$, no scaling, and upscaling $\times 1.5$). We denote by $p_k(x)$, $1 \leq k \leq s$, the model predictions for an input x rescaled according to these operations (see Figure 5). We propose to use a regularization term $\mathcal{L}_{\text{Multiscale}}$ to promote consistency between $p_k(x)$, $1 \leq k \leq s$, and the average $q(x) = \frac{1}{s} \sum_{k=1}^s p_k(x)$:

$$\mathcal{L}_{\text{Multiscale}}(x) = -\frac{1}{s} \sum_{k=1}^s \text{KL}(p_k(x) \parallel q(x)), \quad (3)$$

where KL denotes the Kullback-Leibler divergence. The term is only applied to the input x where the maximum entry of $q(x)$ is above a threshold ρ (equal to 0.8 for all experiments). The regularization is weighted by a parameter λ (set to one in all experiments) and then combined with a cross-entropy loss based on the available annotations. As shown in Tables 2, with multiscale consistency regularization, adaptive label correction further improves segmentation performance in both medical-imaging applications and the WSSS.

3 RELATED WORK

Classification from noisy labels. Early learning and memorization were first discovered in image classification from noisy labels (Liu et al., 2020). Several methods exploit early learning to improve classification models by correcting the labels (Tanaka et al., 2018; Yi & Wu, 2019; Reed et al., 2015; Liu et al., 2020; Xia et al., 2021). Here we show that segmentation from noisy labels also exhibits early learning and memorization. However, these dynamics are different for different semantic categories. ADELE exploits this to perform correction in a class-adaptive fashion.

Segmentation from noisy annotations. Segmentation from noisy annotations is an important problem, especially in the medical domain (Asman & Landman, 2012). Some recent works address this problem by explicitly taking into account systematic human labeling errors (Zhang et al., 2020c), and by modifying the segmentation loss to increase robustness (Shu et al., 2019; Wang et al., 2020a). Min et al. (2019) propose to discover noisy gradient by collecting information from two networks connected with mutual attention. Luo et al. (2021) shows that the network learns high-level spatial structures for fluorescence microscopy images. These structures are then leveraged as supervision signals to alleviate influence from wrong annotations. These methods mainly focus on improving the robustness by exploiting some setting-specific information (*e.g.* network architecture, dataset, requiring some samples with completely clean annotation). In contrast, we propose to study the

learning dynamics of noisy segmentation and propose ADELE, which performs label correction by exploiting early learning.

Weakly supervised semantic segmentation (WSSS). Recent methods for WSSS (Ahn & Kwak, 2018; Zhang et al., 2020a; Fan et al., 2020) are mostly based on the approach introduced by Ref. (Kolesnikov & Lampert, 2016; Wei et al., 2016), where a classification model is first used to produce pixel-level annotations (Zhou et al., 2016), which are then used to train a segmentation model. These techniques mostly focus on improving the initial pixel-level annotations, by modifying the classification model itself (Wei et al., 2017; 2018; Li et al., 2018; Wang et al., 2020b), or by post-processing these annotations (Vernaza & Chandraker, 2017; Ahn & Kwak, 2018; Ahn et al., 2019). However, the resulting annotations are still noisy (Zhang et al., 2020b) (see Figure 3). Our goal is to improve the segmentation model by accounting for this noise. The work that is most similar to our label-correction strategy is (Huang et al., 2018), inspired by traditional *seeded region-growing* techniques (Adams & Bischof, 1994). This method estimates the foreground using an additional model (Jiang et al., 2013), and initializes the foreground segmentation estimate with classification-based annotations. This estimate is used to train a segmentation model, which is then used to iteratively update the estimate. ADELE seeks to *correct* the initial annotations, as opposed to *growing them*, and does not need to identify the foreground estimate or an initial subset of highly-accurate annotations.

4 SEGMENTATION ON MEDICAL IMAGES WITH ANNOTATION NOISE

Segmentation from noisy annotations is a fundamental challenge in the medical domain, where available annotations are often hampered by human error (Zhang et al., 2020c). Here, we evaluate ADELE on a segmentation task where the goal is to identify organs from computed tomography images.

Settings. The dataset consists of 3D CT scans from the SegTHOR dataset (Lambert et al., 2020). Each pixel is assigned to the esophagus, heart, trachea, aorta, or background. We treat each 2D slice of the 3D scan as an example, resizing to 256×256 pixels. We randomly split the slices into a training set of 3638 slices, a validation set of 570 slices, and a test set of 580 slices. Each patient only appears in one of these subsets. We generate annotation noise by applying random degrees of dilation and erosion to the ground-truth segmentation labels, mimicking common human errors (Zhang et al., 2020c) (see Figure 3). In the main experiment, the noisy annotation is with a mIoU of 0.6 w.r.t the ground truth annotation. We further control the degree of dilation and erosion to simulate noisy annotation sets with different noise levels for testing the model robustness. We corrupt all annotations in the training set, but not in the validation and test sets. Our **evaluation metric** is Mean Intersection over Union (mIoU).

	Baseline	ADELE w/o class adaptive	ADELE
Best val	62.6±2.3	40.7±2.5	71.1±0.7
Max test	63.3±2.0	40.7±2.4	71.2±0.6
Last Epoch	59.1±1.3	40.5±2.3	70.8±0.7

Table 1: The mIoU (%) comparison of the baseline and ADELE with or without class-adaptively correcting labels, on the test set of SegTHOR (Lambert et al., 2020). We report the test mIoU of the model that performs best on the validation set (**Best Val**), the test mIoU at the last epoch (**Last Epoch**), and the highest test performance during training (**Max Test**). We report the mean and standard deviation after training the model with five realizations of the noisy annotations.




Figure 6: The performance comparison of the baseline and ADELE on the test set of SegTHOR (Lambert et al., 2020). The model is trained on noisy annotations with various levels of corruption (measured in mIoU with the clean ground truth annotations). ADELE is able to improve the model performance across a wide range of corruption levels.

Results. For a fair comparison, we choose a UNet trained with multi-scale inputs as our **baseline**. We report the mIoU of the baseline and ADELE on the test set of SegTHOR dataset in Table 1. ADELE outperforms the baseline method at all three evaluation epochs. Moreover, correcting labels at the same time for all classes will have a detrimental effect on the performance.

Impacts of noise levels. Figure 6 provides empirical evidence that ADELE is robust to a wide range of noises. The mIoU of noisy annotations (x -axis) indicates the correctness of the noisy annotations. Thus the smaller the mIoU shows the higher noise levels. The improvements achieved by ADELE are substantial when the noise levels are moderate.

Ablation study for each part of ADELE. We perform an ablation study to understand how different parts of ADELE contribute to the final performance. From Table 2, we observe that the model trained with multiple rescaled versions of the input (illustrated in left graph of Figure 5) performs better than the model trained only with the original scale of the input. The proposed spatial consistency regularization further improves the performance. Most importantly, combining any of these methods with label correction would substantially improve the performance. ADELE, which combines label correction with the proposed regularization, achieves the best performance. We also include ablation studies for the hyperparameters r , τ and ρ in Appendix C. Additional segmentation results are provided in Appendix A.1.

Label correction	SegTHOR			PASCAL VOC 2012		
	Single scale	Multiscale input augmentation	Multiscale consistency regularization	Single scale	Multiscale input augmentation	Multiscale consistency regularization
\times	58.8	60.7	62.5	64.5	65.5	66.7
\checkmark	65.2	69.8	72.2	65.6	67.3	69.3

Table 2: Ablation study for ADELE on SegTHOR (Lambert et al., 2020) and PASCAL VOC 2012 (Everingham et al., 2015). We report the mIoU achieved at the last epoch on the validation set for both dataset. Class-adaptive label correction mechanism achieves the best performance when combined with multi-scale consistency regularization.

5 NOISY ANNOTATIONS IN WEAKLY-SUPERVISED SEMANTIC SEGMENTATION

We adopt a prevailing pipeline for training WSSS (described in detail in Section 1), in which some pixel-wise annotations are generated using image level labels to supervise a segmentation network. These pixel-wise annotations are noisy. Therefore, we apply ADELE to this WSSS pipeline.

We evaluate ADELE on a standard WSSS dataset – PASCAL VOC 2012 (Everingham et al., 2015), which has 21 annotation classes (including background), and contains 1464, 1449 and 1456 images in the *training*, *validation (val)* and *test* sets respectively. Following (Shimoda & Yanai, 2019; Zhang et al., 2020a; Wang et al., 2020b; Zhang et al., 2020b; Sun et al., 2020; Yao et al., 2021), we use an augmented training set with 10582 images with annotations from (Hariharan et al., 2011).

Baseline Models. To demonstrate the broad applicability of our approach, we apply ADELE using pixel-level annotations generated by three popular WSSS models: AffinityNet (Ahn & Kwak, 2018), SEAM (Wang et al., 2020b) and ICD (Fan et al., 2020), which do not rely on external datasets or external saliency maps. The annotations are produced by a classification model combined with the post-processing specified in (Ahn & Kwak, 2018; Wang et al., 2020b; Fan et al., 2020). We provide details on the training procedure in Section B in the Appendix. We use the same inference pipeline as SEAM (Wang et al., 2020b), which includes multi-scale inference (Ahn & Kwak, 2018; Wang et al., 2020b; Fan et al., 2020; Zhang et al., 2020d) and CRF (Krähenbühl & Koltun, 2011).

Comparison with the state-of-the-art. Table 3 compares the performance of the proposed method ADELE to state-of-the-art WSSS methods on PASCAL VOC 2012. ADELE improves the performance of AffinityNet (Ahn & Kwak, 2018), SEAM (Wang et al., 2020b) and ICD (Fan et al., 2020) substantially on the validation and test sets. Moreover, ADELE combined with SEAM (Wang et al., 2020b) and ICD (Fan et al., 2020) achieves state-of-the-art performance on both sets. Although it uses only image-level labels, ADELE outperforms state-of-the-art methods (Jiang et al., 2019; Zhang et al., 2020d; Sun et al., 2020; Yao et al., 2021) that rely on external saliency models (Jiang et al., 2013) (see Appendix A.2).

	Previous methods						ADELE +			
	DSRG (Huang et al., 2018)	ICD (Fan et al., 2020)	SCE (Chang et al., 2020)	AffinityNet (Ahn & Kwak, 2018)	SDD (Shimoda & Yanai, 2019)	SEAM (Wang et al., 2020b)	CONTA (Zhang et al., 2020b)	AffinityNet (Ahn & Kwak, 2018)	SEAM (Wang et al., 2020b)	ICD (Fan et al., 2020)
	ResNet-101			ResNet-38						
Val	61.4	64.1	66.1	61.7	64.9	64.5	66.1	64.8	69.3	68.6
Test	63.2	64.3	65.9	63.7	65.5	65.7	66.7	65.5	68.8	68.9

Table 3: Comparison with state-of-the-art methods on the Pascal VOC 2012 dataset using mIoU (%). The **best** and **the best previous method** performance under each set are highlighted in red and blue respectively. The version of CONTA (Zhang et al., 2020b) reported here is deployed combined with SEAM (Wang et al., 2020b). The results clearly show that ADELE-M outperforms other approaches.

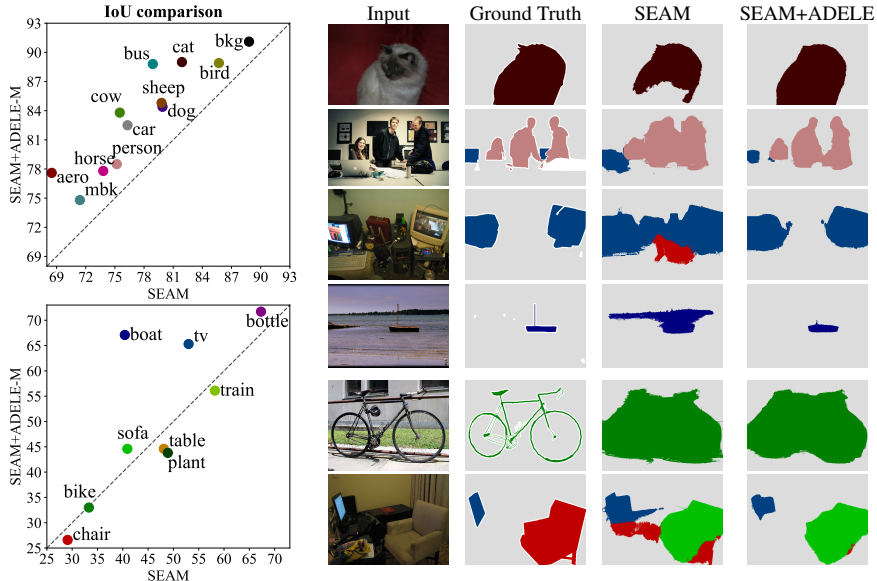


Figure 7: **Left:** Category-wise comparison of the IoU (%) of SEAM (Wang et al., 2020b) and SEAM combined with the proposed method ADELE on the validation set of PASCAL VOC 2012. **Right:** Visualization of the segmentation results of both methods for several examples. ADELE successfully improves segmentation for the first four examples, but not for the last two. We set the background color to gray for ease of visualization. More examples can be found in Appendix A.1

Figure 7 compares the performance of SEAM and the performance of ADELE combined with SEAM on the validation set separately for each semantic category. ADELE improves performance for most categories, with the exception of a few categories where the baseline model does not perform well (e.g. chair, bike). On the right of Figure 7, we show some qualitative segmentation results from the validation set. The first four rows show examples where ADELE successfully improves the SEAM segmentation. The last two rows show examples where it does not. In both the output of SEAM has highly structured segmentation errors: the prediction encompasses the bike but completely fails to capture its inner structure, and the chair is misclassified as a sofa. This supports the conclusion that ADELE provides less improvement when the baseline method performs poorly.

6 CONCLUSION

In this work, we introduce a novel method to improve the robustness of segmentation models trained on noisy annotations. Inspired from the early-learning phenomenon, we proposed ADELE to boost the performance on the segmentation of thoracic organ, where noise is incorporated to resemble human annotation errors. Moreover, standard segmentation networks, equipped with ADELE, achieve the state-of-the-art results for WSSS on PASCAL VOC 2012. We hope that this work will trigger interest in the design of new forms of segmentation methods that provide robustness to annotation noise, as this is a crucial challenge in applications such as medicine. We also hope that the work will motivate further study of the early-learning and memorization phenomena in settings beyond classification.

REFERENCES

- Rolf Adams and Leanne Bischof. Seeded region growing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(6):641–647, 1994.
- Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- Devansh Arpit, Stanisław Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, and Simon Lacoste-Julien. A closer look at memorization in deep networks. In *International Conference on Machine Learning*, 2017.
- Andrew J Asman and Bennett A Landman. Formulating spatially varying performance in the statistical fusion framework. *IEEE Transactions on Medical Imaging*, 31(6):1326–1336, 2012.
- David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*, 2019.
- 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 Conference on Computer Vision and Pattern Recognition*, 2020.
- Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014.
- 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.
- 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*, 2015.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Ieee, 2009.
- 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.
- Junsong Fan, Zhaoxiang Zhang, Chunfeng Song, and Tieniu Tan. Learning integral objects with intra-class discriminator for weakly-supervised semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- 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.
- Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *Proceedings of the IEEE International Conference on Computer Vision*, 2011.
- Mohammad Havaei, Axel Davy, David Warde-Farley, Antoine Biard, Aaron Courville, Yoshua Bengio, Chris Pal, Pierre-Marc Jodoin, and Hugo Larochelle. Brain tumor segmentation with deep neural networks. *Medical Image Analysis*, 35:18–31, 2017.

- 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*, 2018.
- Huaizu Jiang, Jingdong Wang, Zejian Yuan, Yang Wu, Nanning Zheng, and Shipeng Li. Salient object detection: A discriminative regional feature integration approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013.
- Peng-Tao Jiang, Qibin Hou, Yang Cao, Ming-Ming Cheng, Yunchao Wei, and Hong-Kai Xiong. Integral object mining via online attention accumulation. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019.
- Eytan Kats, Jacob Goldberger, and Hayit Greenspan. A soft staple algorithm combined with anatomical knowledge. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 510–517. Springer, 2019.
- Jongmok Kim, Jooyoung Jang, and Hyunwoo Park. Structured consistency loss for semi-supervised semantic segmentation. *arXiv preprint arXiv:2001.04647*, 2020.
- Alexander Kolesnikov and Christoph H Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *Proceedings of the European Conference on Computer Vision*. Springer, 2016.
- Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. *Advances in Neural Information Processing Systems*, 24, 2011.
- Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *International Conference on Learning Representations*, 2018.
- Zoé Lambert, Caroline Petitjean, Bernard Dubray, and Su Kuan. Segthor: Segmentation of thoracic organs at risk in ct images. In *2020 Tenth International Conference on Image Processing Theory, Tools and Applications (IPTA)*. IEEE, 2020.
- Junnan Li, Richard Socher, and Steven C.H. Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *International Conference on Learning Representations*, 2020.
- Qizhu Li, Anurag Arnab, and Philip HS Torr. Weakly-and semi-supervised panoptic segmentation. In *Proceedings of the European Conference on Computer Vision*, 2018.
- Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- Di Lin, Yuanfeng Ji, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Multi-scale context intertwining for semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, 2018.
- Kangning Liu, Yiqiu Shen, Nan Wu, Jakub Chledowski, Carlos Fernandez-Granda, and Krzysztof J Geras. Weakly-supervised high-resolution segmentation of mammography images for breast cancer diagnosis. In *Medical Imaging with Deep Learning*, 2021.
- Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. *Advances in Neural Information Processing Systems*, 2020.
- Yaoru Luo, Guole Liu, Wenjing Li, Yuanhao Guo, and Ge Yang. Deep neural networks learn meta-structures to segment fluorescence microscopy images. *arXiv preprint arXiv:2103.11594*, 2021.
- Shaobo Min, Xuejin Chen, Zheng-Jun Zha, Feng Wu, and Yongdong Zhang. A two-stream mutual attention network for semi-supervised biomedical segmentation with noisy labels. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 4578–4585, 2019.

- Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8):1979–1993, 2018.
- Scott E. Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. Training deep neural networks on noisy labels with bootstrapping. *CoRR*, abs/1412.6596, 2015.
- 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*. Springer, 2015.
- Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- Jo Schlemper, Ozan Oktay, Liang Chen, Jacqueline Matthew, Caroline Knight, Bernhard Kainz, Ben Glocker, and Daniel Rueckert. Attention-gated networks for improving ultrasound scan plane detection. *arXiv preprint arXiv:1804.05338*, 2018.
- Wataru Shimoda and Keiji Yanai. Self-supervised difference detection for weakly-supervised semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019.
- Yucheng Shu, Xiao Wu, and Weisheng Li. Lvc-net: Medical image segmentation with noisy label based on local visual cues. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019.
- Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685*, 2020.
- 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*, 2019.
- Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Van Gool. Mining cross-image semantics for weakly supervised semantic segmentation. In *Proceedings of the European Conference on Computer Vision*. Springer, 2020.
- Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- 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*, 2017.
- Michael Trembl, José Arjona-Medina, Thomas Unterthiner, Rupesh Durgesh, Felix Friedmann, Peter Schuberth, Andreas Mayr, Martin Heusel, Markus Hofmarcher, Michael Widrich, et al. Speeding up semantic segmentation for autonomous driving. In *MLITS, NIPS Workshop*, 2016.
- Paul Vernaza and Manmohan Chandraker. Learning random-walk label propagation for weakly-supervised semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- Guotai Wang, Xinglong Liu, Chaoping Li, Zhiyong Xu, Jiugen Ruan, Haifeng Zhu, Tao Meng, Kang Li, Ning Huang, and Shaoting Zhang. A noise-robust framework for automatic segmentation of covid-19 pneumonia lesions from ct images. *IEEE Transactions on Medical Imaging*, 39(8): 2653–2663, 2020a.
- Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

- 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 Conference on Computer Vision and Pattern Recognition*, 2020b.
- Yunchao Wei, Xiaodan Liang, Yunpeng Chen, Xiaohui Shen, Ming-Ming Cheng, Jiashi Feng, Yao Zhao, and Shuicheng Yan. Stc: A simple to complex framework for weakly-supervised semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(11):2314–2320, 2016.
- Yunchao Wei, Jiashi Feng, Xiaodan Liang, Ming-Ming Cheng, Yao Zhao, and Shuicheng Yan. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, July 2017.
- Yunchao Wei, Huaxin Xiao, Honghui Shi, Zequn Jie, Jiashi Feng, and Thomas S Huang. Revisiting dilated convolution: A simple approach for weakly-and semi-supervised semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- Zifeng Wu, Chunhua Shen, and Anton Van Den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition*, 90:119–133, 2019.
- Xiaobo Xia, Tongliang Liu, Bo Han, Chen Gong, Nannan Wang, Zongyuan Ge, and Yi Chang. Robust early-learning: Hindering the memorization of noisy labels. In *International Conference on Learning Representations*, 2021.
- Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. Denseaspp for semantic segmentation in street scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- Yazhou Yao, Tao Chen, Guosen Xie, Chuanyi Zhang, Fumin Shen, Qi Wu, Zhenmin Tang, and Jian Zhang. Non-salient region object mining for weakly supervised semantic segmentation. *arXiv preprint arXiv:2103.14581*, 2021.
- Kun Yi and Jianxin Wu. Probabilistic end-to-end noise correction for learning with noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- Bingfeng Zhang, Jimin Xiao, Yunchao Wei, Mingjie Sun, and Kaizhu Huang. Reliability does matter: An end-to-end weakly supervised semantic segmentation approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020a.
- Dong Zhang, Hanwang Zhang, Jinhui Tang, Xiansheng Hua, and Qianru Sun. Causal intervention for weakly-supervised semantic segmentation. *arXiv preprint arXiv:2009.12547*, 2020b.
- Le Zhang, Ryutaro Tanno, Mou-Cheng Xu, Chen Jin, Joseph Jacob, Olga Ciccarelli, Frederik Barkhof, and Daniel C Alexander. Disentangling human error from the ground truth in segmentation of medical images. *arXiv preprint arXiv:2007.15963*, 2020c.
- Tianyi Zhang, Guosheng Lin, Weide Liu, Jianfei Cai, and Alex Kot. Splitting vs. merging: Mining object regions with discrepancy and intersection loss for weakly supervised semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, 2020d.
- 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*, 2017.
- 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*, 2016.
- Tianyi Zhou, Shengjie Wang, and Jeff Bilmes. Robust curriculum learning: From clean label detection to noisy label self-correction. In *International Conference on Learning Representations*, 2020.

Appendix for “Adaptive Early-Learning Correction for Segmentation from Noisy Annotations”

The appendix is organized as follows.

- In Appendix A, we include additional figures illustrating the early-learning and memorization phenomena in PASCAL VOC 2012 and SegTHOR. We also include additional results for these two datasets.
- In Appendix B, we describe the implementation details of our proposed method ADELE, including a description of the hyperparameters, experiment settings, and other technical details.
- In Appendix C, we report ablation studies on the influence of different components and hyperparameters of ADELE.

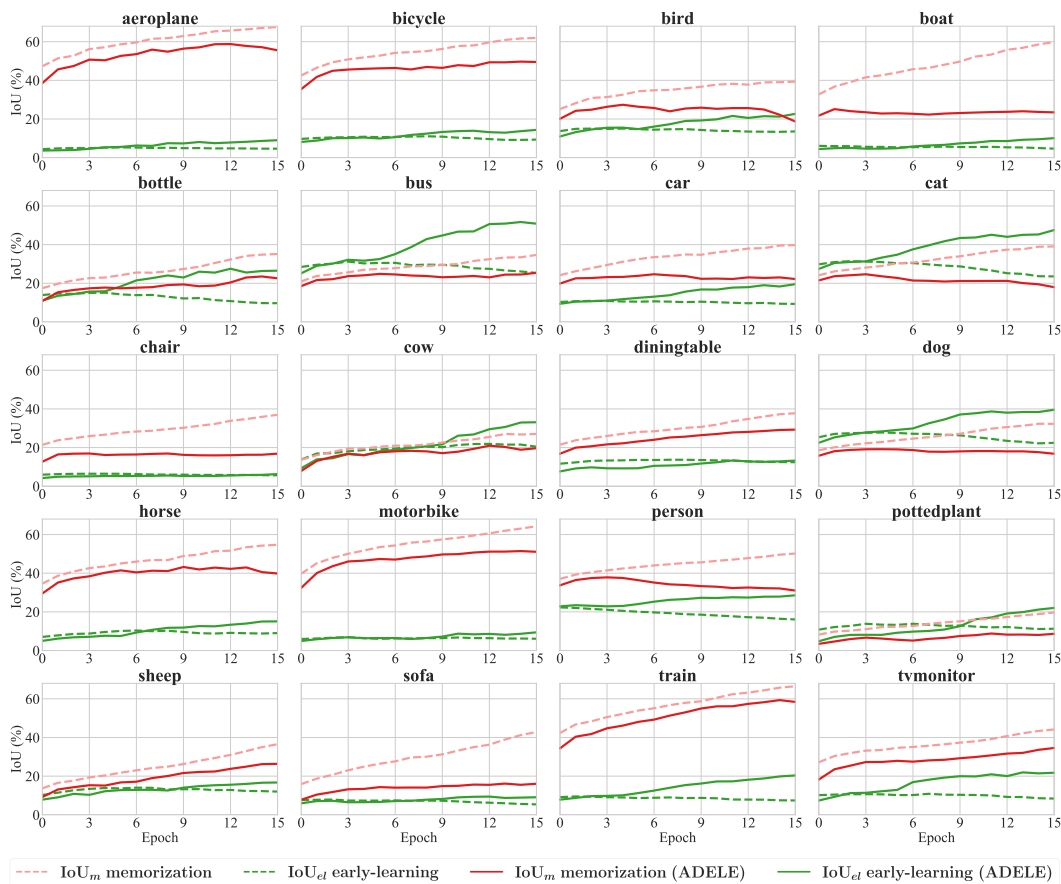


Figure 8: We visualize the effect of early learning (IoU_{el} , green curves) and memorization (IoU_m , red curves) on segmentation models trained with (solid lines) and without (dashed lines) ADELE for each category of the WSSS dataset VOC 2012 (Everingham et al., 2015). The WSSS model is a standard DeepLab-v1 network trained with annotations obtained from SEAM (Wang et al., 2020b). IoU_{el} is the IOU between the model output and the ground truth computed over the incorrectly-labeled pixels. IoU_m is the IOU between the model output and the incorrect annotations. For all classes, IoU_m increases substantially as training proceeds because the model gradually memorizes the incorrect annotations. This again occurs at different speeds for different categories. In contrast, IoU_{el} first increases during an early-learning stage where the model learns to correctly segment the incorrectly-labeled pixels, but eventually decreases as memorization occurs (the phenomenon is more evident when we zoom in, as shown in Figure 9 in the Appendix). Like memorization, early-learning also happens at varying speeds for the different semantic categories.




Figure 9: Zoomed-in illustration of early learning for the different semantic categories on PASCAL VOC 2012. IoU_{el} first increases during an early-learning stage where the model learns to correctly segment the incorrectly-labeled pixels, but eventually decreases as memorization occurs. Early learning happens at varying speeds for the different semantic categories. The experimental setting is the same as Figure 2.




Figure 10: Illustration of the proposed curve fitting method to decide when to begin label correction in ADELE (Results on Pascal VOC). On the left, we plot the IoU between the model predictions and the initial noisy annotations for the same model used in Figures 9 and 3 and the corresponding fit with the parametric model in Equation 1. The label correction beginning iteration is based on the relative slope change of the fitted curve. The center image shows the label correction times for different semantic categories, showing that they are quite different. On the right graph, the green line shows the IoU_{el} for a given category. The IoU_{el} equals the IoU between the model output and the ground truth computed over the incorrectly-labeled pixels, and therefore quantifies early-learning. The label correction begins close to the end of the early-learning phase, as desired.

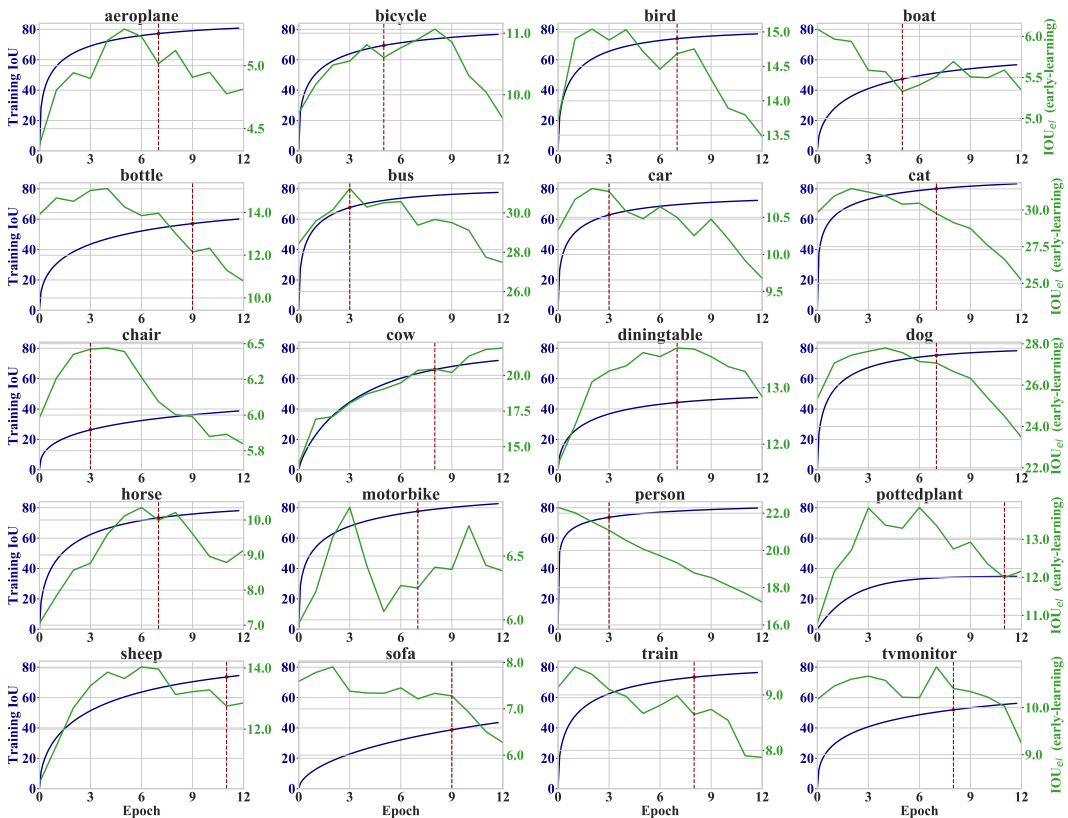


Figure 11: Illustration of the proposed curve fitting method to decide when to begin label correction in ADELE. The blue line shows the parametric model in Equation 1 fit to the IoU between the model predictions and the initial noisy annotations for the same WSSS model used in Figures 2 and 3. The green line shows the IoU_{el} for a given category. The IoU_{el} equals the IOU between the model output and the ground truth computed over the incorrectly-labeled pixels, and therefore quantifies early-learning. The label correction begins close to the end of the early-learning phase for most of the categories with a few exceptions (boat and motorbike).




Figure 12: Multi-scale consistency regularization leads to more accurate corrected annotations (result on Pascal VOC).




Figure 13: Additional segmentation results of SEAM and SEAM+ADELE for several examples on the validation set of PASCAL VOC 2012. We set the background color to gray for ease of visualization. Supplementary for Figure 7.

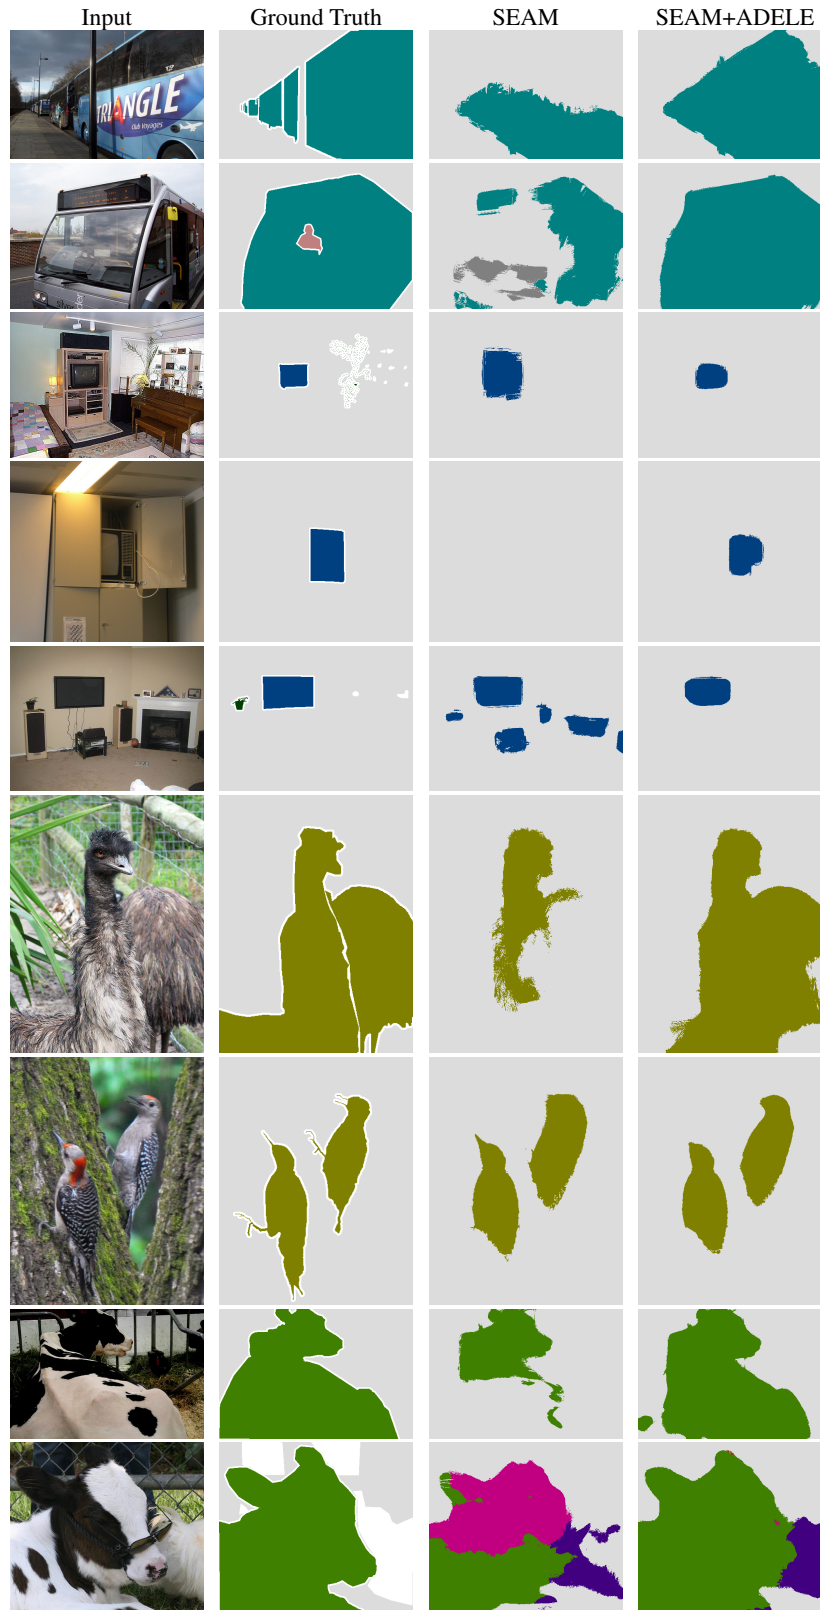


Figure 14: Additional segmentation results of SEAM and SEAM+ADELE for several examples on the validation set of PASCAL VOC 2012. We set the background color to gray for ease of visualization. Supplementary for Figure 7.




Figure 15: Additional segmentation results of SEAM and SEAM+ADELE for several examples on the validation set of PASCAL VOC 2012. We set the background color to gray for ease of visualization. Supplementary materials for Figure 7.

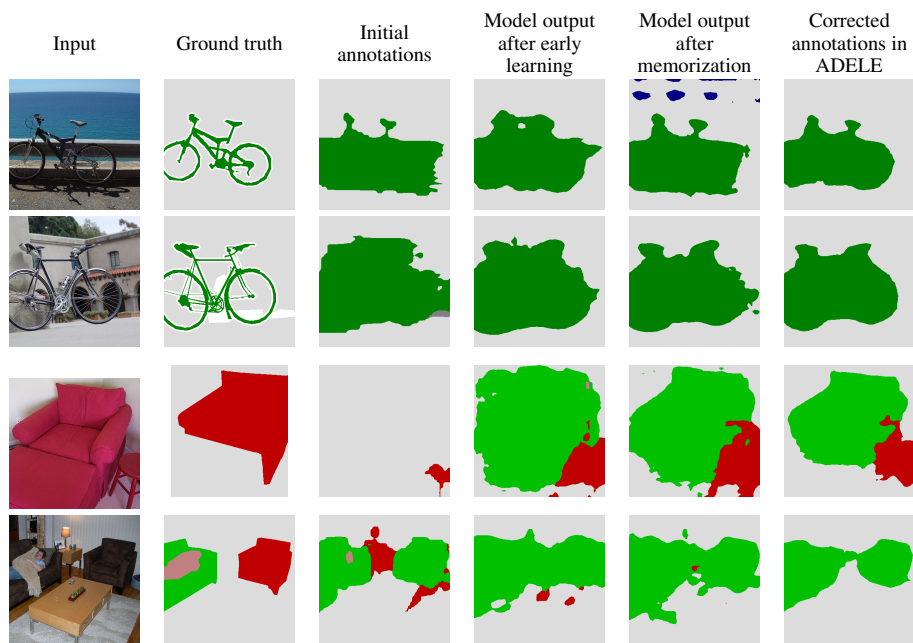


Figure 16: Illustration of a possible limitation of the proposed label-correction approach on PASCAL VOC 2012. The initial annotations coarsely segment the bike and misclassify the chair as a sofa consistently in several training examples. This highly structured annotation noise could potentially prevent early learning from happening, and therefore from being exploited for label correction. We set the background color to light gray for ease of visualization.

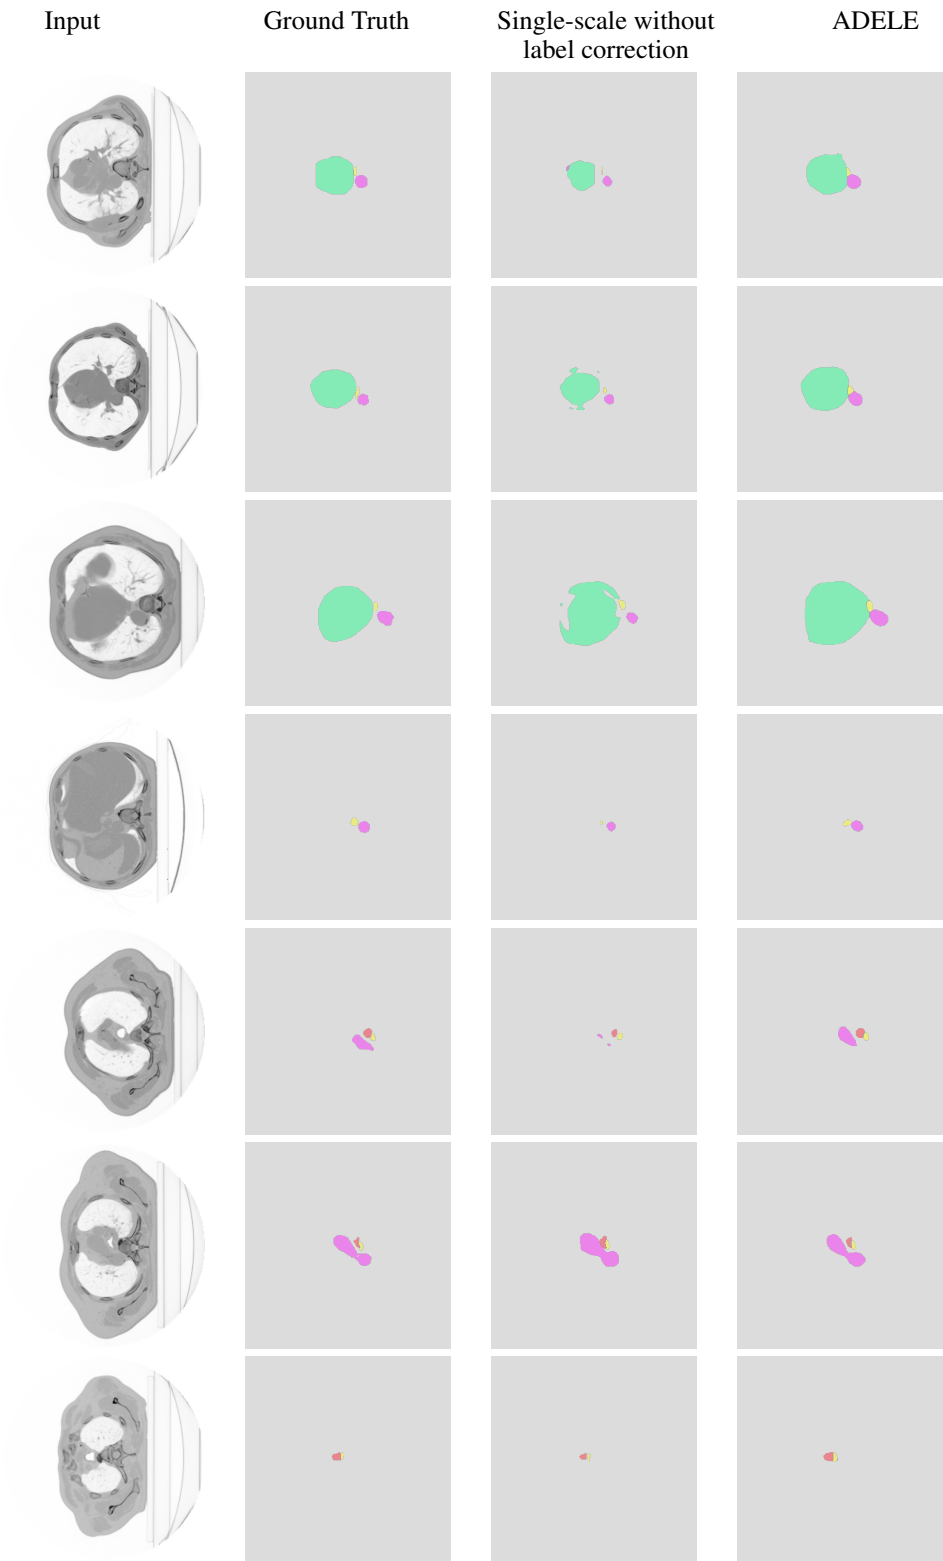


Figure 17: Visualization of the segmentation results of baseline and baseline+ADELE for several examples on the *validation* set of SegTHOR.

A ADDITIONAL EXPERIMENTAL RESULTS

In this section, we include additional examples (A.1) and results (A.2).

A.1 ADDITIONAL EXAMPLES

WSSS dataset (PASCAL VOC 2012)

- *Early-learning.* (Supplementary material for Figure 2 and 4 in Section 2) Figure 8 demonstrates that early-learning and memorization on a segmentation CNN trained with noisy annotations generated by a classification model for the standard WSSS dataset (PASCAL VOC2012). In Figure 9 we show the zoomed-in version of the early-learning IoU_{el} curves. Early-learning happens for most of the classes at different speeds. Figure 10 show the curve fitting illustration for PASCAL VOC. Figure 11 shows that the label correction begins close to the end of the early-learning phase for most of the semantic categories for PASCAL VOC.
- *Multiscale consistency.* Figure 12 shows that multi-scale consistency regularization leads to more accurate corrected annotations for PASCAL VOC.
- *Superior performance.* (Supplementary material for Figure 7 in Section 5) We provide more visualization examples on the validation set of PASCAL VOC in Figure 13, 14 and 15 comparing ADELE to the baseline method SEAM. ADELE reduces false positives for some examples (e.g. boat, person, sofa, bottle, tv etc.) and produces a more complete segmentation of other examples (e.g. cat, dog, bird, bus, bottle, tv, horse etc.).
- *Highly structured annotation noise.* (Supplementary material for Figure 7) Figure 16 shows training examples with highly structured noise, which may prevent early learning from happening, and therefore from being used for label correction.

More visual examples for SegTHOR

- *Superior performance.* (Supplementary material for Section 4). We show some visual examples that indicate ADELE improves segmentation performance on the validation set of SegTHOR dataset in Figure 17.

A.2 ADDITIONAL RESULTS

ADELE outperforms methods using external information. (Supplementary material for Section 5) ADELE using initial labels generated from SEAM (Wang et al., 2020b) and ICD (Fan et al., 2020) achieves state-of-the-art performance without using external saliency models. This performance is on par with or even slightly better than methods that rely on external saliency models (Jiang et al., 2019; Zhang et al., 2020d; Sun et al., 2020; Yao et al., 2021).

Method	Supervision	Backbone	val	test
DSRG (Huang et al., 2018)	I.+ S.	ResNet-101	61.4	63.2
SPLIT MERGE (Zhang et al., 2020d)	I.+S.	ResNet-50	66.6	66.7
MCIS (Sun et al., 2020)	I.+S.	ResNet-101	66.2	66.9
NSROM (Yao et al., 2021)	I.+S.	ResNet-101	68.2	68.5
SEAM + Ours	I.	ResNet-38	69.3	68.8
ICD + Ours	I.	ResNet-38	68.6	68.9

Table 4: Comparison with previous works that use external saliency models (Jiang et al., 2013) on the PASCAL VOC 2012 dataset (Everingham et al., 2015). I. stands for image-level labels, S. stands for external saliency models.

ADELE improves performance for most categories. (Supplementary material for Figure 7 in Section 5) Table 5 shows that ADELE performs substantially better than the baseline method SEAM on most of the semantic categories, as well as on average.

ADELE improves performance across different noise levels. We provided the full result of Figure 6 in Table 6 for SegTHOR dataset, which shows that ADELE can improve the model performance across different noise level.

Method	bkg	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbk	person	plant	sheep	sofa	train	tv	mIoU
SEAM	88.8	68.5	33.3	85.7	40.4	67.3	78.9	76.3	81.9	29.1	75.5	48.1	79.9	73.8	71.4	75.2	48.9	79.8	40.9	58.2	53.0	64.5
SEAM + Ours	91.1	77.6	33.0	88.9	67.1	71.7	88.8	82.5	89.0	26.6	83.8	44.6	84.4	77.8	74.8	78.5	43.8	84.8	44.6	56.1	65.3	69.3

Table 5: Category-wise comparison of the IoU (%) of SEAM (Wang et al., 2020b) and SEAM combined with the proposed method ADELE on the validation set of PASCAL VOC 2012.

Iteration of erosion/dilation	Annotation mIoU	Multiscale input augmentation (baseline)	Multiscale label correction	Multiscale consistency regularization	ADELE
0	1.00	0.745	0.743	0.762	0.766
1	0.91	0.743	0.726	0.759	0.757
2	0.73	0.702	0.710	0.714	0.734
3	0.61	0.646	0.658	0.647	0.711
4	0.52	0.556	0.651	0.606	0.666
6	0.39	0.446	0.556	0.514	0.564
8	0.28	0.416	0.407	0.423	0.481

Table 6: The mIoU (%) comparison of the baseline and ADELE on the test set of SegTHOR (Lambert et al., 2020). We report the test mIoU of the model that achieves best mIoU on the validation set. It can be seen that ADELE stably improve the model performance across different noise levels.

B IMPLEMENTATION DETAILS

B.1 HYPERPARAMETERS

Table 7 provides a complete list of hyperparameter values for our experiments. We use almost identical hyperparameters of ADELE on PASCAL VOC 2012 and SegTHOR. λ (*consistency strength*) controls the strength of the consistency regularization added to the cross entropy loss. ρ (*consistency confidence threshold*) is the threshold that determines when multiscale consistency regularization is applied (when the maximum prediction probability for any category in any pixel of the average q is above ρ). r (*curve fitting threshold*) is the threshold that controls label correction for each semantic category (see Equation equation 2). τ (*label correction confidence threshold*) is the threshold that determines which pixels are corrected by the model prediction. Only the pixels with confidence (maximum prediction probability for any category) above τ are corrected. Note that we report ablation studies for most of these hyperparameters in Section C.

Dataset	λ	ρ	r	τ
PASCAL VOC 2012	1	0.8	0.9	0.8
SegThor	1	0.8	0.9	0.7

Table 7: Complete list of ADELE hyperparameters for PASCAL VOC 2012 (Everingham et al., 2015) and SegTHOR (Lambert et al., 2020).

B.2 MEDICAL SEGMENTATION WITH SIMULATED NOISE: SEGTHOR

Here we provide implementation details for our experiments on SegTHOR (Supplementary material for Section 4 in the main paper).

Details of Noise Synthesis. We apply dilation and erosion on the ground-truth segmentation masks to simulate “over-annotation” and “under-annotation” labels respectively. In particular, “over-

annotation” labels tend to assign background pixels surrounding a target organ to the class of this organ. On the contrary, “under-annotation” labels tend to assign a foreground pixel on the edge of a target organ to the background class. These two noise patterns have been previously utilized to simulate common errors that human annotators make during manual segmentation (Zhang et al., 2020c). In our experiments, we randomly choose the type and degree of synthetic noise that will be applied to each example.

Training. For SegTHOR (Lambert et al., 2020), we use one NVIDIA V100 GPU to train the model. We train the UNet (Ronneberger et al., 2015) using SGD optimizer. To optimize the hyper-parameters, we search for the learning rate in $\{0.1, 0.01, 0.001, 0.0001, 0.00001\}$ and set to 0.01. λ , ρ , r and τ are set according to Table 7. We trained our model for 100 epochs with a batch size of 5.

Inference during testing. We conduct the single-scale evaluation using the input without augmentations.

Generating model prediction to conduct label correction. For the medical dataset, since we do not apply any augmentation to the input images, we directly use the outputs during training at every iteration to correct the labels. This is in contrast to PASCAL VOC where we compute the outputs at the end of each epoch, as explained in Section C.1.

B.3 WEAKLY SUPERVISED SEMANTIC SEGMENTATION: PASCAL VOC 2012

We provide the implementation details of the model for PASCAL VOC 2012. (Supplementary material for Section 5).

Training. For PASCAL VOC 2012 (Everingham et al., 2015), we use two NVIDIA Quadro RTX 8000 GPUs to train the model. We use the official code of AffinityNet (Ahn & Kwak, 2018), SEAM (Wang et al., 2020b) and ICD (Fan et al., 2020) to generate initial pixel-level annotations that are used for training the segmentation network. The experimental settings to train the segmentation network follow (Ahn & Kwak, 2018; Zhang et al., 2020a; Shimoda & Yanai, 2019; Wang et al., 2020b) in which DeepLabv1 (Chen et al., 2014) is adopted. The model uses ResNet38 (Wu et al., 2019) as a backbone, with the initial weight load from ImageNet (Deng et al., 2009) pretrained classification model. Following the same settings as SEAM (Wang et al., 2020b), we use a SGD optimizer with momentum 0.9 and weight decay $5e^{-4}$. The initial learning rate lr_{init} is set to 0.001, and reduced following a polynomial function of the iteration number itr : $lr_{itr} = lr_{init} \left(1 - \frac{itr}{max_itr}\right)^\gamma$ with $\gamma = 0.9$. We train our segmentation network for 20000 iterations ($max_itr = 20000$) with a batch size of 10. The input images are randomly scaled and then randomly cropped to 448×448 .

Inference during testing. During testing, we use the same inference pipeline as SEAM (Wang et al., 2020b), which includes multi-scale inference (Ahn & Kwak, 2018; Wang et al., 2020b; Fan et al., 2020; Zhang et al., 2020d) and CRF (Krähenbühl & Koltun, 2011).

Generating model predictions to conduct label correction. Updating the model using the model predictions after processing each batch would be difficult for the PASCAL VOC 2012 dataset. The reason is that random data augmentations (*e.g.* rescaling, cropping, random changing contrast in the image, *etc.*) are often applied (Ahn & Kwak, 2018; Wang et al., 2020b; Fan et al., 2020; Zhang et al., 2020d) and these augmentations would need to be inverted in order to use the output for label correction (otherwise the predictions would be inconsistent across training epochs). In fact, some augmentations are not invertible at all, *e.g.* cropping cuts the object out thus is not invertible, rescaling might result in loss of information thus is not invertible as well. In order to avoid this issue on the PASCAL VOC dataset, we evaluate the model at the end of each training epoch on inputs without any random augmentation, and then use the model outputs to perform label correction.




Figure 18: Ablation study for r on SegTHOR. We fixed the other hyperparameters to the default settings in Table 7. We report the test mIoU (%) at the best validation epoch on the **Left**, the full result is shown on the **Right**.




Figure 19: Ablation study for τ on SegTHOR. We fixed the other hyperparameters to the default settings in Table 7. We report the test mIoU (%) at the best validation epoch on the **Left**, the full result is shown on the **Right**. The result shows that ADELE is not very sensitive to the value of τ .

C ADDITIONAL ABLATION STUDIES

C.1 MEDICAL SEGMENTATION WITH SIMULATED NOISE: SEGTHOR

Here we report additional ablation studies for different components in the proposed label correction method on SegTHOR dataset (Supplementary material for Section 4).

Different options for label correction As described in Section 2.2, ADELE uses the model outputs to correct labels. In this section, we compare two options for computing these outputs.

- *Iteration*: the outputs are computed after each training iteration.
- *Epoch*: the outputs are computed at the end of each training epoch.

Table 8 shows that on SegTHOR *Iteration* performs slightly better on the best test mIoU than *Epoch*, and outperforms it substantially on the mIoU at the last epoch, suggesting that *Iteration* is more effective in preventing memorization. Both two options are improving results with respect to the baseline.

Method	Best val	Last epoch	Max test
ADELE (<i>Iteration</i>)	71.1 ± 0.7	70.8 ± 0.7	71.2 ± 0.6
ADELE (<i>Epoch</i>)	69.6 ± 0.8	64.1 ± 1.3	70.2 ± 0.5

Table 8: mIoU(%) of ADELE on SegTHOR when label correction is based on model output at the end of each iteration or each epoch. The former achieves better results. **r and τ for label correction.** We report the ablation for r and τ for ADELE on SegTHOR dataset in Figure 18 and 19 respectively. For the r value, we observe similar results as in the PASCAL VOC dataset. If the r value is too small (e.g. 0, 0.3) and label correction is conducted too early, the network has not been trained well. This degrades the label correction quality, which hurts the generalization of the model. If the r value is too large (e.g. 0.99) then the network barely conduct label correction for any class before stopping training. Results are again very robust to the choice of τ value.




Figure 20: Ablation study for r and τ on PASCAL VOC. We fixed the other hyperparameters to the default settings in Table 7. We report the validation mIoU (%) at the last training epoch.

C.2 WEAKLY SUPERVISED SEMANTIC SEGMENTATION: PASCAL VOC

r and τ for label correction. We report the ablation result for r and τ of ADELE on PASCAL VOC dataset in Figure 20. Small r values encourage the model to conduct label correction earlier, larger r values delay correction. If the r value is too small (*e.g.* 0.3) and label correction is conducted too early, the network has not been trained well. This degrades the label correction quality, which hurts the generalization of the model. If the r value is too large (*e.g.* 0.99), then the method barely conducts label correction for any class before the end of training, which results in a performance similar to the case without label correction. We observe that our model is very robust to the choice of τ .