CheXseg: Combining Expert Annotations with DNN-generated Saliency Maps for X-ray Segmentation

Soham Gadgil* Mark Endo* Emily Wen* Andrew Y. Ng Pranav Rajpurkar Department of Computer Science, Stanford University SGADGIL@STANFORD.EDU MARKENDO@STANFORD.EDU EMILYWEN@STANFORD.EDU ANG@CS.STANFORD.EDU PRANAVSR@CS.STANFORD.EDU

Abstract

Medical image segmentation models are typically supervised by expert annotations at the pixel-level, which can be expensive to acquire. In this work, we propose a method that combines the high quality of pixel-level expert annotations with the scale of coarse DNN-generated saliency maps for training multi-label semantic segmentation models. We demonstrate the application of our semi-supervised method, which we call CheXseg, on multi-label chest X-ray interpretation. We find that CheXseg improves upon the performance (mIoU) of fully-supervised methods that use only pixel-level expert annotations by 9.7% and weakly-supervised methods that use only DNN-generated saliency maps by 73.1%. Our best method is able to match radiologist agreement on three out of ten pathologies and reduces the overall performance gap by 57.2% as compared to weakly-supervised methods. **Keywords:** Semi-Supervised Segmentation, Saliency Maps, Localization Performance

1. Introduction

The "black box" nature of neural networks represents a barrier to physicians' trust and model adoption in the clinical setting (Kelly et al., 2019). Saliency maps are a popular set of explanation methods that highlight regions of the image that are important for disease classification, but they have been shown to be untrustworthy for medical image interpretation (Eitel and Ritter, 2019; Crosby et al., 2020; Young et al., 2019; Arun et al., 2020). Segmentation models can produce more accurate pixel-level maps, but their training is typically limited by the time-consuming process of collecting expert annotations. The combination of saliency maps generated from widely available classification models and a limited amount of expert annotations for training medical image segmentation models may be able to provide higher quality segmentations at a lower cost, but this approach remains relatively unexplored.

In this work, we develop *CheXseg*, a semi-supervised method for multi-pathology segmentation that leverages both the pixel-level expert annotations and the saliency maps generated by image classification models. First, we find that CheXseg achieves a mean IoU of 0.270, outperforming both fully-supervised (mIoU of 0.246) and weakly-supervised (mIoU of 0.156) methods. Second, we find that initializing the segmentation encoder with weights

^{*} Contributed equally

learned from supervised classification of the same task leads to higher performance than using a self-supervised MoCo initialization (He et al., 2020) or ImageNet initialization (Deng et al., 2009). Third, CheXseg reduces the overall gap to radiologist localization performance (mIoU) by 57.2% compared to solely using DNN-generated saliency maps. We expect this method to be broadly useful for medical image segmentation, where classification models are widely available and expert annotations are expensive.

2. Related Work

2.1. Weakly-Supervised Semantic Segmentation

In this work, we focus on an approach in which classification models trained with image-level labels are used to create pixel-level pseudo-labels (Yao and Gong, 2020; Ciga and Martel, 2019; Ouyang et al., 2019). These pseudo-labels can then be utilized to train a segmentation model. Our paper is closely related to Viniavskyi et al. (2020), which proposes a deep CNN-based approach that generates pseudo-labels by applying an Inter-pixel Relation Network (IRNet) (Ahn et al., 2019) to improve Grad-CAM++ (Chattopadhay et al., 2018) generated activation maps. This approach is evaluated on the SIIM-ACR Pneumothorax dataset. In our work, we generate pseudo-labels with IRNet and extend the approach to the semi-supervised setting for a larger set of pathologies.

Several other methods also propagate class activation from areas of high confidence to similar adjacent regions (Kolesnikov and Lampert, 2016; Huang et al., 2018; Ahn and Kwak, 2018). We choose to build upon Viniavskyi et al. (2020) rather than these approaches because it has the best performance on the PASCAL VOC 2012 (Everingham et al., 2010) validation set (mIoU of 0.646) and was shown to perform well on a medical imaging task.

2.2. Semi- and Fully-Supervised Semantic Segmentation

Semi-supervised methods use a combination of expert pixel-level annotations and pseudolabels to train semantic segmentation models. Some weakly-supervised methods have been extended to semi-supervised methods through the replacement of pseudo-labels. The weakly-supervised SGAN model (Yao and Gong, 2020) was adapted to a semi-supervised setting by replacing a subset of the saliency maps with the corresponding manually annotated ground truth labels. In this work, we use a similar idea of utilizing radiologist annotated labels in addition to saliency maps to train the segmentation model, extending Viniavskyi et al. (2020)'s weakly-supervised model.

Though weakly- and semi-supervised methods can perform well, fully-supervised methods are still considered the upper-bound (Chan et al., 2020). Many fully-supervised semantic segmentation approaches have been proposed for chest X-rays (Sirazitdinov et al., 2019; Jaiswal et al., 2019), but none of these extend their work to the semi-supervised setting.

3. Methods

3.1. Setup

The multi-label semantic segmentation task is to classify each pixel of a chest X-ray image into zero or more of 10 possible pathologies: Airspace Opacity, Atelectasis, Cardiomegaly, Consolidation, Edema, Enlarged Cardiomediastinum, Lung Lesion, Pleural Effusion, Pneumothorax, and Support Devices.

We utilize CheXpert (Irvin et al., 2019), an existing large dataset with 224,316 chest Xrays of 65240 patients. This dataset features image-level labels obtained using an automated labeler that detects the aforementioned pathologies from radiology reports. A subset of the dataset is hand-annotated by radiologists at the pixel level. In our work, we use a set of 200 radiologist-annotated chest X-rays to validate model performance of the weaklysupervised method. For the fully-supervised and semi-supervised methods, we use 150 of the radiologist-annotated labels as a train set and save 50 examples for a validation set. We selected this validation set to exclude scarce pathologies, as examples with those pathologies are most valuable in the training process. When evaluating performance on this validation set, we only look at the most common pathologies. For all methods, we use a test set of an additional 500 radiologist-annotated images.

Models are evaluated by their average performance on the semantic segmentation task across the ten pathologies of interest. For each of the pathologies, the IoU (Intersectionover-Union) score is computed. We report the mIoU (mean IoU) score, which is the average IoU score across all pathologies.

3.2. CheXseg

We develop *CheXseg*, a semi-supervised method for multi-pathology segmentation that leverages both the pixel-level expert annotations and the saliency maps generated by image classification models. In this method, a DenseNet121 (Huang et al., 2017) classification model, trained on the entire CheXpert train set, is first used to generate saliency maps using Grad-CAM (Selvaraju et al., 2017). This approach uses the classification model outputs to create a coarse localization map highlighting the image regions important for prediction. The saliency maps are further processed to create per-pixel segmentation masks, referred to as weak pseudo-labels, by using either a thresholding scheme or an Inter-Pixel Relation Network (IRNet) (Ahn et al., 2019). IRNet takes these generated CAMs and tries to improve the seeds by training two output branches, a displacement vector field and a class boundary map. Details about these methods are provided in Appendix A.

Once the pseudo-labels have been generated, we combine them with pixel-level expert annotations in a semi-supervised manner to train semantic segmentation models. Due to the scarcity of high-quality pixel-level expert annotations, we implement a sampling strategy of the different label types to allow for a high level of contribution from the radiologist annotations. For comparison, we train fully-supervised segmentation models (solely using pixel-level annotations) and weakly-supervised segmentation models (solely using pseudolabels).

All our methods utilize DeepLabv3+ as the core semantic segmentation model (Chen et al., 2018). We experiment with various encoder initializations to transfer knowledge from the classification task to the segmentation task. Our experiments utilize a ResNet encoder architecture (He et al., 2016). Figure 1 shows a visual representation of the different supervision strategies used.



Figure 1: Workflows of the different methods analysed for chest X-ray segmentation

3.3. Training Details

Here we describe the training details of our classification and segmentation models. The IRNet training details are available in Appendix B.

3.3.1. Classification Model

The pre-trained classification model used for generating the CAMs from image-level labels is a DenseNet121. We use the Adam optimizer with default β -parameters of $\beta_1 = 0.9$, $\beta_2 = 0.999$ and learning rate 1×10^{-4} which is fixed for the duration of the training. Batches are sampled using a fixed batch size of 16 images. We train for 3 epochs, saving checkpoints every 4800 iterations.

3.3.2. Segmentation Model

The semantic segmentation model is trained using a class average dice loss and Adam optimizer. We use a learning rate of 0.001 when training a small amount of data and we decrease it to 0.0001 when training a large amount of data. We train on up to four Nvidia GTX 1070s using a batch size of 8.

4. Experiments

4.1. Combining Weak and Full Supervision

We investigate the segmentation performance of combining DNN-generated saliency maps and expert annotations with various sampling ratios. We use 100 saliency maps in combination with 200 annotated pixel-level labels and explore different weightings between the two types of labels. We vary the probability of selecting an expert annotation in a single batch during training, $p \in 0, 0.2, 0.4, 0.6, 0.8.0.85, 0.9, 1$. Thus, p determines the expected fraction of images with expert annotations in a single batch. For each value of p, we perform three trials containing different sets of the 100 saliency maps. The results reported are the mIoU scores obtained by averaging across these 3 trials. The segmentation model is initialized with CheXpert encoder weights since it performs the best as observed in experiment 4.2. We also compare the performance of this semi-supervised model with the weakly-supervised and fully-supervised models.

Results We find that for both Grad-CAM and IRNet, there is an inverted U-shape trend in performance as we increase p. There is a sharp increase in performance as p increases from 0 (mIoU score of 0.156), and then the curves remain relatively flat before dropping off when p = 1. Specifically, p = 0.9 (*CheXseg*) gives the best mIoU performance of 0.270 ± 0.00872 and 0.267 ± 0.00993 (95% CI) for CAM and IRNet respectively. This high weighting of pixel-level labels takes advantage of the more accurate information encoded within these labels as compared to the saliency maps. The reduced performance for the fully-supervised case (p = 1, mIoU score of 0.246 ± 0.01837 , 95% CI) is likely attributed to the weak labels no longer being utilized in training. The size of the train set shrinks, and the model does not benefit from the variation and the scale provided by the weak pseudolabels. For the weakly-supervised case (p = 0, mIoU score of 0.156), the poor performance can be attributed to the absence of any pixel-level expert annotations to guide the model predictions. Detailed results are shown in Figure 2.



Figure 2: IoU scores of semi-supervised segmentation models (averaged across 3 trials) using either Grad-CAM or IRNet to generate weak labels. The DeepLabV3+ and ResNet18 setup is used with CheXpert encoder initializion. p is the probability of selecting a pixel-level labeled training sample in the current batch. (a) is the average IoU score calculated across all pathologies. (b), (c), and (d) are IoU scores for Pleural, Consolidation, and Edema respectively. p = 0 represents the weakly-supervised case while p = 1 represents the fully-supervised case. Full results in Table 1.

4.2. Comparing Encoder Initializations

We investigate the impact of using various encoder initializations on segmentation performance. In the fully-supervised and weakly-supervised methods, we initialize the encoder weights to either a CheXpert classification model (Irvin et al., 2019), MoCo-CXR (Sowrirajan et al., 2020), ImageNet (Deng et al., 2009), or random.

Results We find that for all methods, the best models are initialized with CheXpert encoder weights. This may be expected since CheXpert weights are learned from supervised

learning on the same dataset that we use for segmentation. The models initialized with MoCo-CXR weights have similar performance to the models with ImageNet encoder initialization. Figure 3 shows the detailed results for fully-supervised and weakly-supervised encoder initializations.



segmentation models with DeepLabV3+ and ResNet18 initialized with various weights. CheXpert encoder initialization results in best performance (average 0.246).

(a) mIoU of fully-supervised (b) mIoU of weakly-supervised segmentation models with DeepLabV3+ and ResNet18 setup using either Grad-CAM or IRNet pseudo-labels and various encoder initializations. IRNet outperforms CAMs when using CheXpert encoder initialization (0.156 vs 0.142), but underperforms when using other initializations (0.111 vs 0.136)with Random, 0.128 vs 0.139 with MoCo-CXR, and 0.124 vs 0.139 with ImageNet). Full comparisons with confidence intervals in Table 2.

Figure 3: Performance of fully-supervised and weakly-supervised methods

4.3. Comparison to Radiologists

Compared to our best weakly-supervised method, CheXseg reduces the overall performance (mIoU) gap with radiologists by 57.2%. CheXseg outperforms radiologists in terms of IoU score on the segmentation of Atelectasis (156% higher), Airspace Opacity (70% higher), and Pleural Effusion (30% higher), while performing worse on the remaining pathologies. Detailed results are shown in Figure 5.

5. Qualitative Results

Figure 4 shows the qualitative results for two pathologies - Cardiomegaly and Airspace Opeacity - for the best weakly-supervised (IRNet) and semi-supervised (CheXseg) methods. CheXseg gives better visualizations that are closer to ground truth as compared to the best weakly-supervised approach.



(b) Qualitative results for Airspace Opacity

Figure 4: Qualitative Results for Cardomegaly and Airspace Opacity. Column (i) represents the Ground Truth segmentation map. Column (ii) represents the segmentation map obtained from CheXseg. Column (iii) represents the segmentation map obtained from using the best weakly supervised method (IRNet).



Figure 5: IoUs of radiologists, our best semi-supervised method (CheXseg), and our best weakly-supervised method. CheXseg uses Grad-CAM while the weakly-supervised method uses IRNet.

6. Discussion

In this work, we develop *CheXseg*, a semi-supervised method for multi-pathology segmentation that leverages the benefits of both available medical image classification models and expert pixel-level annotations.

How does CheXseg performance compare to fully-supervised and weakly-supervised methods? We find that with a weighted sampling of saliency maps and expert annotations, our proposed method outperforms both the fully-supervised and weakly-supervised methods alone. An expert annotation to saliency map sample ratio of 0.9 (CheXseg) gives the best mIoU score of 0.270, compared to 0.246 for fully-supervised and 0.156 for weakly-supervised. This suggests a tradeoff between emphasizing the accurate information of expert pixel-level annotations and incorporating additional but noisy cues from weak pseudo-labels.

How do segmentation models initialized with CheXpert, ImageNet, MoCo-CXR, and random encoder weights compare? We find that CheXpert encoder initialization achieves the highest performance, followed by self-supervised (MoCo-CXR) initialization and ImageNet initialization. Random encoder initialization has the worst performance, as no transfer learning occurs. Since CheXpert weights are pre-trained using image-level labels for the same tasks, it is expected that this knowledge transfers well to segmentation on the same dataset and the same set of tasks. Whereas classification models with MoCo-CXR encoder initialization have been found to outperform classification models with ImageNet encoder initialization (Sowrirajan et al., 2020), we find that the two initializations have approximately the same performance for segmentation.

In closing, our work proposes a simple semi-supervised method that combines the benefits of the wide availability of classification models with the quality of expert annotations. We expect our method may be broadly useful, able to lower the cost of development and improve the performance of medical image segmentation models.

References

- 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, pages 4981–4990, 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, pages 2209–2218, 2019.
- Nishanth Arun, Nathan Gaw, Praveer Singh, Ken Chang, Mehak Aggarwal, Bryan Chen, Katharina Hoebel, Sharut Gupta, Jay Patel, Mishka Gidwani, et al. Assessing the (un) trustworthiness of saliency maps for localizing abnormalities in medical imaging. arXiv preprint arXiv:2008.02766, 2020.
- Lyndon Chan, Mahdi S Hosseini, and Konstantinos N Plataniotis. A comprehensive analysis of weakly-supervised semantic segmentation in different image domains. *International Journal of Computer Vision*, pages 1–24, 2020.
- Aditya Chattopadhay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In 2018 IEEE Winter Conference on Applications of Computer Vision (WACV), pages 839–847. IEEE, 2018.
- 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 (ECCV), pages 801–818, 2018.
- Ozan Ciga and Anne L Martel. Learning to segment images with classification labels. arXiv preprint arXiv:1912.12533, 2019.
- Jennie Crosby, Sophia Chen, Feng Li, Heber MacMahon, and Maryellen Giger. Network output visualization to uncover limitations of deep learning detection of pneumothorax. In Frank W. Samuelson and Sian Taylor-Phillips, editors, *Medical Imaging 2020: Image Perception, Observer Performance, and Technology Assessment*, volume 11316, pages 125 – 128. International Society for Optics and Photonics, SPIE, 2020. doi: 10.1117/12. 2550066. URL https://doi.org/10.1117/12.2550066.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A largescale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009.
- Fabian Eitel and Kerstin Ritter. Testing the robustness of attribution methods for convolutional neural networks in mri-based alzheimer's disease classification, 2019.
- Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer* vision, 88(2):303–338, 2010.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 9729–9738, 2020.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2261–2269, 2017. doi: 10.1109/CVPR.2017.243.
- Zilong Huang, Xinggang Wang, Jiasi Wang, Wenyu Liu, and Jingdong Wang. Weaklysupervised semantic segmentation network with deep seeded region growing. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 7014–7023, 2018.
- Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 33, pages 590–597, 2019.
- Amit Kumar Jaiswal, Prayag Tiwari, Sachin Kumar, Deepak Gupta, Ashish Khanna, and Joel JPC Rodrigues. Identifying pneumonia in chest x-rays: A deep learning approach. *Measurement*, 145:511–518, 2019.
- Christopher J Kelly, Alan Karthikesalingam, Mustafa Suleyman, Greg Corrado, and Dominic King. Key challenges for delivering clinical impact with artificial intelligence. BMC medicine, 17(1):1–9, 2019.
- 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.
- Xi Ouyang, Zhong Xue, Yiqiang Zhan, Xiang Sean Zhou, Qingfeng Wang, Ying Zhou, Qian Wang, and Jie-Zhi Cheng. Weakly supervised segmentation framework with uncertainty: A study on pneumothorax segmentation in chest x-ray. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 613–621. Springer, 2019.
- 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.
- Ilyas Sirazitdinov, Maksym Kholiavchenko, Tamerlan Mustafaev, Yuan Yixuan, Ramil Kuleev, and Bulat Ibragimov. Deep neural network ensemble for pneumonia localization from a large-scale chest x-ray database. *Computers & Electrical Engineering*, 78: 388–399, 2019.

- Hari Sowrirajan, Jingbo Yang, Andrew Y Ng, and Pranav Rajpurkar. Moco pretraining improves representation and transferability of chest x-ray models. *arXiv preprint arXiv:2010.05352*, 2020.
- Ostap Viniavskyi, Mariia Dobko, and Oles Dobosevych. Weakly-supervised segmentation for disease localization in chest x-ray images. In *International Conference on Artificial Intelligence in Medicine*, pages 249–259. Springer, 2020.
- Qi Yao and Xiaojin Gong. Saliency guided self-attention network for weakly and semisupervised semantic segmentation. *IEEE Access*, 8:14413–14423, 2020.
- Kyle Young, Gareth Booth, Becks Simpson, Reuben Dutton, and Sally Shrapnel. Deep neural network or dermatologist? *Lecture Notes in Computer Science*, page 48–55, 2019. ISSN 1611-3349. doi: 10.1007/978-3-030-33850-3_6. URL http://dx.doi.org/10.1007/ 978-3-030-33850-3_6.

Appendix A. Methods for Generating Weak Pseudo-Labels

A.1. Grad-CAM

Grad-CAM is used to obtain saliency maps from the model predictions to highlight the areas that the model focuses on during classification. It makes use of the gradients of the output classes flowing into the last convolutional layer to make low-resolution heatmaps (3 × 3 in case of ResNets). Specifically, the gradients flowing back are global-average-pooled over the width and height dimensions (indexed by i and j respectively) to obtain the importance of the *k*th feature map for target class c, α_k^c :

$$\alpha_k^c = \frac{1}{Z} \sum_i \sum_j \frac{\partial y^c}{A_{ij}^k} \tag{1}$$

Here, y^c is the score (before softmax) of the class c, Z is the normalization factor, and A^k is the kth feature map activation.

After this, a weighted sum of the final feature maps followed by a ReLU is performed:

$$L^{c}_{Grad-CAM} = ReLU\left(\sum_{k} \alpha^{c}_{k} A^{k}\right)$$
⁽²⁾

A thresholding scheme is then used to convert the heatmap for each pathology into a segmentation map to use as pseudo-labels. The probability threshold is determined per pathology by maximizing the mIoU on the CheXpert train set.

A.2. Inter-Pixel Relation Network (IRNet)

We follow the method IRNet, which takes the previously generated CAMs and tries to improve these seeds by training two output branches. The first branch predicts a displacement vector field in which each pixel is represented by a 2D vector pointing to the centroid of the instance that the pixel is a part of. This displacement field is converted to a class-agnostic instance map by grouping together pixels whose vectors point to the same centroid. The second branch is used to detect class boundaries by computing pairwise semantic affinities, which is a confidence score for class equivalence between a pair of pixels. The instance-wise CAMs obtained from the first branch are enhanced by propagating their attention scores to relevant areas using the computed affinities between neighboring pixels. Finally, pseudolabels are generated independently for each pathology with positive pixels being the ones with higher class attention scores.

Appendix B. Training Details

B.1. IRNet

The two branches of IRNet share the same ResNet50 backbone and are jointly trained by minimizing the sum of three losses:

1. Loss for displacement field prediction, which consists of two losses:

(a) L1 loss between the image coordinate displacement between a pair of nearby foreground pixels (i, j), denoted by $\hat{\delta}(i, j) = x_j - x_i$, and the displacement obtained from their vector fields, \mathcal{D} , denoted by $\delta(i, j) = \mathcal{D}(x_j) - \mathcal{D}(x_i)$:

$$\mathcal{L}_{\rm fg}^{\mathcal{D}} = \frac{1}{|\mathcal{P}_{\rm fg}^+|} \sum_{(i,j)\in\mathcal{P}_{\rm fg}^+} \left| \hat{\delta}(i,j) - \delta(i,j) \right| \tag{3}$$

Here, \mathcal{P}_{fg}^+ refers to the set of neighboring foreground pixel pairs with the same pseudo label (pixels with attention scores larger than 0.3).

(b) Loss for background pixels as a normalized sum of their image coordinate displacements:

$$\mathcal{L}_{\rm bg}^{\mathcal{D}} = \frac{1}{|\mathcal{P}_{\rm bg}^+|} \sum_{(i,j)\in\mathcal{P}_{\rm bg}^+} |\delta(i,j)| \tag{4}$$

Here, \mathcal{P}_{bg}^+ refers to the set of neighboring background pixel pairs with the same pseudo label (pixels with attention scores less than 0.05).

2. Loss for Class Boundary Detection, which makes use of the semantic affinities between a pair of pixels x_i and x_j , a_{ij} :

$$a_{ij} = 1 - \max_{k \in \Pi_{ij}} \mathcal{B}(x_k) \tag{5}$$

Here, Π_{ij} is a set of pixels on the line between x_i and x_j and $\mathcal{B} \in [0, 1]^{w \times h}$ is the output. Then, the loss is the cross-entropy loss between the binary affinity label, with value 1 for the same pseudo-class labels and 0 otherwise, and the predicted affinity of two pixels:

$$\mathcal{L}^{\mathcal{B}} = -\sum_{(i,j)\in\mathcal{P}_{fg}^{+}} \frac{\log(a_{ij})}{2\mid\mathcal{P}_{fg}^{+}\mid} - \sum_{(i,j)\in\mathcal{P}_{bg}^{+}} \frac{\log(a_{ij})}{2\mid\mathcal{P}_{bg}^{+}\mid} - \sum_{(i,j)\in\mathcal{P}^{-}} \frac{\log(1-a_{ij})}{\mid\mathcal{P}^{-}\mid}$$
(6)

Here, \mathcal{P}^- represents the set of pixel pairs with different pseudo labels.

The two branches are jointly trained by minimizing all three losses at the same time:

$$\mathcal{L} = \mathcal{L}_{\rm fg}^{\mathcal{D}} + \mathcal{L}_{\rm bg}^{\mathcal{D}} + \mathcal{L}^{\mathcal{B}}$$
(7)

The model is trained with stochastic gradient descent using a learning rate of 0.1 with polynomial decay and a batch size of 16. The segmentation maps created are then used in the semantic segmentation model as weak pseudo-labels.

Appendix C. Semi-Supervised Results

Table 1: IoU scores for semi-supervised segmentation using IRNet and CAMs by weighting of pixel-level and weak pseudo-labels. p represents the probability of picking a expert annotated training example in the current batch.

Task	Method	p=0	p = 0.2	p = 0.4	p=0.6	p = 0.8	p = 0.85	p=0.9	p=1
Mean IoU	CAM IRNet	0.142 0.156	0.238 0.216	0.246 0.228	0.248 0.236	0.255 0.252	0.261 0.254	0.270 0.267	0.246
Airspace Opacity	CAM IRNet	0.161 0.225	0.364 0.334	0.381 0.355	0.382 0.362	0.376 0.374	0.376 0.377	0.398 0.403	0.388
Atelectasis	CAM IRNet	0.099 0.182	0.299 0.269	0.295 0.278	0.297 0.287	0.306 0.316	0.299 0.308	0.310 0.323	0.307
Cardiomegaly	CAM IRNet	0.326 0.280	0.420 0.393	0.438 0.401	0.434 0.425	0.440 0.428	0.442 0.440	0.450 0.461	0.445
Consolidation	CAM IRNet	0.056 0.076	0.097 0.099	0.099 0.104	0.105 0.107	0.109 0.108	0.108 0.110	0.114 0.110	0.103
Edema	CAM IRNet	0.170 0.149	0.231 0.198	0.253 0.219	0.241 0.239	0.250 0.237	0.261 0.242	0.270 0.257	0.225
Enlarged Cardiomediastinum	CAM IRNet	0.236 0.327	0.516 0.466	0.494 0.480	0.494 0.489	0.514 0.528	0.526 0.531	0.543 0.535	0.549
Lung Lesion	CAM IRNet	0.002 0.003	0.004 0.002	0.003 0.002	0.003 0.003	0.005 0.014	0.008 0.014	0.008 0.007	0.002
Pleural Effusion	CAM IRNet	0.150 0.171	0.254 0.182	0.256 0.192	0.263 0.191	0.264 0.243	0.266 0.243	0.272 0.273	0.214
Pneumothorax	CAM IRNet	0.057 0.024	0.021 0.018	0.025 0.029	0.035 0.033	0.044 0.040	0.074 0.050	0.077 0.053	0.017
Support Devices	CAM IRNet	0.159 0.123	0.203 0.202	0.223 0.217	0.231 0.226	0.241 0.229	0.249 0.221	0.257 0.248	0.246

Appendix D. Weakly-Supervised Results

Table 2: IoU scores of weakly-supervised segmentation models using either CAM or IRNet pseudo-labels and varying encoder architectures and train set sizes. Confidence intervals are calculated using $\alpha = 0.05$.

Encoder Architecture	Train Set Size	Test mIoU		
	100	0.074 ± 0.00783		
ResNet18	2000	0.126 ± 0.00275		
	20000	0.136 ± 0.00990		
	100	0.082 ± 0.00396		
ResNet18-ImageNet	2000	0.132 ± 0.00687		
	20000	0.139 ± 0.00095		
	100	0.084 ± 0.00238		
ResNet18-MoCo-CXR	2000	0.133 ± 0.00002		
	20000	0.139 ± 0.00076		
	100	0.129 ± 0.00036		
ResNet 18-CheXpert	2000	0.141 ± 0.00003		
	20000	0.142 ± 0.00095		
	100	0.076 ± 0.00133		
ResNet18	2000	0.105 ± 0.00292		
	20000	0.111 ± 0.00103		
	100	0.081 ± 0.00245		
ResNet18-ImageNet	2000	0.117 ± 0.00576		
	20000	0.124 ± 0.00001		
	100	0.082 ± 0.00007		
ResNet18-MoCo-CXR	2000	0.120 ± 0.00010		
	20000	0.128 ± 0.00001		
	100	0.125 ± 0.00031		
ResNet18-CheXpert	2000	0.146 ± 0.00002		
	20000	0.156 ± 0.00000		
	Encoder Architecture ResNet18 ResNet18-ImageNet ResNet18-MoCo-CXR ResNet18-CheXpert ResNet18-ImageNet ResNet18-MoCo-CXR ResNet18-MoCo-CXR	Encoder ArchitectureTrain Set SizeResNet18100ResNet1820000ResNet18-ImageNet2000ResNet18-ImageNet2000ResNet18-MoCo-CXR2000ResNet18-MoCo-CXR2000ResNet18-MoCo-CXR2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet182000ResNet182000ResNet182000ResNet182000ResNet18-ImageNet2000ResNet18-ImageNet2000ResNet18-MoCo-CXR2000ResNet18-ImageNet2000ResNet18-MoCo-CXR2000ResNet18-MoCo-CXR2000ResNet18-MoCo-CXR2000ResNet18-MoCo-CXR2000ResNet18-MoCo-CXR2000ResNet18-MoCo-CXR2000ResNet18-MoCo-CXR2000ResNet18-MoCo-CXR2000ResNet18-MoCo-CXR2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet18-CheXpert2000ResNet1		