

Learning to Segment Medical Images from Few-Shot Sparse Labels

Pedro H. T. Gama*, Hugo Oliveira*[†], Jefersson A. dos Santos*

*Department of Computer Science, Universidade Federal de Minas Gerais, Brazil

Email: {phtg, jefersson}@dcc.ufmg.br

[†]Institute of Mathematics and Statistics, University of São Paulo, Brazil

Email: oliveirahugo@vision.ime.usp.br

Abstract—In this paper, we propose a novel approach for few-shot semantic segmentation with sparse labeled images. We investigate the effectiveness of our method, which is based on the Model-Agnostic Meta-Learning (MAML) algorithm, in the medical scenario, where the use of sparse labeling and few-shot can alleviate the cost of producing new annotated datasets. Our method uses sparse labels in the meta-training and dense labels in the meta-test, thus making the model learn to predict dense labels from sparse ones. We conducted experiments with four Chest X-Ray datasets to evaluate two types of annotations (grid and points). The results show that our method is the most suitable when the target domain highly differs from source domains, achieving Jaccard scores comparable to dense labels, using less than 2% of the pixels of an image with labels in few-shot scenarios.

I. INTRODUCTION

Medical images are useful tools to assist doctors in multiple clinical scenarios and to plan for surgery. X-Ray, Magnetic Resonance Imaging (MRI), Computed Tomography (CT) and other imaging modalities are non-invasive methods that can help in diagnosis, pathology localization, anatomical studies, and other tasks [1].

Convolutional Neural Networks (CNNs) and their variants are the state of the art for object classification, detection, semantic segmentation and other Computer Vision problems. Classical convolutional networks are known for their large data requirements, often hindering their usage in scenarios where data availability is limited, as in medical imaging. Relatively few medical datasets are publicly available due to privacy and ethical concerns [2]–[8].

Even among the public datasets, properly curated labeled data is limited due to the need for specialized annotators (i.e. radiologists), severely hampering the creation of general models for medical image understanding. While many datasets contain image-level annotations indicating the presence or absence of a set of medical conditions, the creation of pixel-level labels that allow for the training of semantic segmentation models is much more laborious. Volumetric image modalities as MRIs or CT scans further compound these difficulties by requiring per-slice annotations, often followed by cross-axes analysis to detect inconsistencies, which can take hours for one single exam. Hence, there is a need for automatic and semi-automatic segmentation methods to assist physicians in the annotation of these images. One way to alleviate the

burden of medical professionals in labeling the exams is to improve the generalization capabilities of existing pretrained models. For instance, domain adaptation can be used to transfer knowledge from related medical imaging datasets to improve segmentation performance in unseen target tasks.

Scenarios with low amounts of data available, often called Few-shot, have been studied in recent years, and tasks such as few-shot classification [9], [10] are the most explored in the literature. As for the problem of pixel-level annotations, one efficient option is *sparse labeling*, that is, specifying the labels of a small number of pixels. Methods that can make efficient use of few-shot and sparse labels can solve semantic segmentation medical problems on datasets created in a labor-efficient manner. Thus, **our main contribution** is the proposal of a novel approach to few-shot semantic segmentation in medical images from sparse labels. For that, we introduce the **Weakly-supervised Segmentation Learning (WeaSeL)**, which extends the MAML [10] algorithm by introducing annotation sparsity directly to its meta-training stage.

II. RELATED WORK

Biomedical Image Segmentation: Automatic and semi-automatic segmentation of medical images has been studied for decades, as such methods can considerably alleviate the burden of physicians in labeling these data [11]. Medical image segmentation from sparse labels can be especially useful, since these images often have a small number of labeled samples due to data privacy restrictions and the lack of specialists for annotating the samples. The survey of Tajbakhsh *et. al* [12] reviews a collection of Deep Learning (DL) solutions on medical image segmentation, which includes a section for segmentation with noisy/sparse labels. The methods reviewed can be summarized in a selective loss with or without mask completion. Sparse segmentation methods either use a loss function able to ignore unlabeled pixels/voxels or employ some technique to augment the sparse annotations to resemble dense masks.

Meta-Learning: Few-shot learning has attracted considerable attention over the last years, mainly due to recent advances in Deep Learning for Self-Supervised learning [13] and Meta-Learning [9], [10]. Meta-Learning has become a proliferous research topic in Deep Learning, as the literature aims to improve on the generalization capabilities of Deep Neural

Networks (DNNs). Within Meta-Learning, two prominent and distinct methodologies gained attention: Gradient-based and Metric learning. Finn *et al.* [10] proposed the MAML framework, based on the trend of Gradient-based approaches. MAML uses multiple tasks – that is, multiple data/sample pairs and a loss function – during its meta-training to create a generalizable model able to perform quick adaptation and feature reuse to infer over unseen related tasks. Meta-Learning arose at first for object classification tasks, while tasks such as detection and segmentation still lacking development. Additionally, the intersection of few-shot learning and medical image segmentation from sparse labels has proven to be quite a challenging task, with very few methods being described in the literature [14].

Few-shot Semantic Segmentation: Few-shot segmentation became a relevant topic only recently. Several methods [15]–[19] were recently proposed to make use of the small subset of labeled images, often called *support* set, to segment a query image. However, the vast majority of few-shot segmentation methods in the literature do not explore sparsely labeled images in the support set, only dealing with densely labeled samples. Rakelly *et al.* [20] proposed the first algorithm for few-shot sparse segmentation: Guided Networks. Guided Nets use a pretrained CNN backbone to extract features of both the support set and the query image. The weakly labeled support samples and sparse segmentation masks are combined, pooled, and subsequently used as weights to segment the query. The vast majority of few-shot segmentation methods [15]–[20] were originally proposed to RGB images as the ones seen in datasets such as the Pascal VOC [21], thus, most of the literature on this topic rely on pretrained CNN or even Fully Convolutional Networks (FCNs) backbones.

III. METHODOLOGY

We define a segmentation task $\mathcal{S} = \{\mathcal{D}^{sup}, \mathcal{D}^{qry}, c\}$, where \mathcal{D} is a dataset with partitions \mathcal{D}^{qry} (or query set) and \mathcal{D}^{sup} (or support set) such that, $\mathcal{D}^{qry} \cap \mathcal{D}^{sup} = \emptyset$. The class c is the positive/foreground class for the task. A dataset \mathcal{D} is a set of pairs (\mathbf{x}, \mathbf{y}) , where \mathbf{x} is an image and \mathbf{y} is the respective semantic label, with \mathbf{y} being a dense mask for images in \mathcal{D}^{qry} and a sparsely annotated mask for the \mathcal{D}^{sup} set. In particular, we define a few-shot segmentation task \mathcal{F} as a task wherein \mathcal{D}^{sup} has a small amount of labeled samples (e.g. 20 or less) and \mathcal{D}^{qry} labels are absent or unknown. We refer to a few-shot task as k -shot when $k = |\mathcal{D}^{sup}|$, that is, the number of samples in its support set is k .

Thus, given a set of segmentation tasks $\mathcal{S} = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_n\}$, and a target few-shot task \mathcal{F} , we want to segment the images from $\mathcal{D}_{\mathcal{F}}^{qry}$ using information from both \mathcal{S} and $\mathcal{D}_{\mathcal{F}}^{sup}$. Also, it holds that no pair of image/semantic label of \mathcal{F} is present in any task \mathcal{S}_i in either *qry* or *sup* partition.

A. Gradient-based Segmentation from Sparse Labels

As previously mentioned, we propose WeaSeL, an gradient-based approach for semantic segmentation derived from

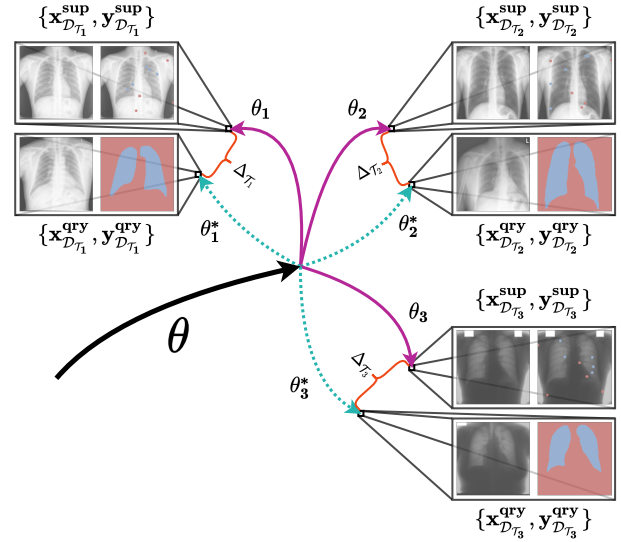


Fig. 1. Visualization of the proposed approach. The parameters θ_i are optimized using sparse labels from support sets. The optimal θ_i^* would be obtained if dense labels were presented in meta-training, as the ones in the query set used to compute the task outer loss. The model learns to intrinsically minimize this difference Δ between parameters, and thus fastly adapt to the few-shot task.

MAML [10]. A graphical representation of our method can be seen in Figure 1.

We define a meta task $\mathcal{T} = \{\mathcal{L}, \mathcal{D}^{sup}, \mathcal{D}^{qry}, c\}$ comprised of a loss function \mathcal{L} and a pair of datasets \mathcal{D}^{sup} and \mathcal{D}^{qry} – the meta-train and meta-test sets, respectively – and a class c . This is an extension of the segmentation task definition from Section III. We assume a distribution over tasks $p(\mathcal{T})$ that our model f_θ (parametrized by θ) is desired to be able to adapt to. A description of this supervised training adaptation of MAML can be seen in Algorithm 1.

Algorithm 1 Model-Agnostic Meta Learning: Weakly Supervised

Require: $p(\mathcal{T})$: distributions over tasks

Require: α, β : step size hyperparameters

Randomly initialize θ

while not done **do**

 Sample batch of tasks $\mathcal{T}_i \sim p(\mathcal{T})$

for all \mathcal{T}_i **do**

 Sample batch of datapoints $S_i = \{(\mathbf{x}, \mathbf{y})\}$ from $\mathcal{D}_{\mathcal{T}_i}^{sup}$

 Compute $\nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(f_{\theta})$ using S_i and $\mathcal{L}_{\mathcal{T}_i}$

 Update adapted parameters using gradient descent:

$\theta_i = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(f_{\theta})$

 Sample batch of datapoints $Q_i = \{(\mathbf{x}, \mathbf{y})\}$ from $\mathcal{D}_{\mathcal{T}_i}^{qry}$

end for

 Update $\theta \leftarrow \theta - \beta \nabla_{\theta} \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(f_{\theta_i})$ using each Q_i and $\mathcal{L}_{\mathcal{T}_i}$

end while

During meta-training, in contrast with MAML, the inner loss is computed using the sparse labels from samples of \mathcal{D}^{sup} , while the outer loss component takes into account the dense

labels from samples of \mathcal{D}^{qry} . This strategy directly encourages the model to generate dense segmentation from the sparse labels fed on the tuning phase.

Given that our problem is semantic segmentation, we define the loss function $\mathcal{L}_{\mathcal{T}_i}$ for all tasks $\mathcal{T}_i \sim p(\mathcal{T})$ as the Cross-Entropy loss, ignoring the pixels with unknown labels. This is achieved by computing a weighted Cross-Entropy loss in a pixel-wise fashion, with the caveat that unknown pixels have weight 0.

When the meta-training process is finished, we fine-tune the model on the target \mathcal{F} task, through a conventional supervised training on the $\mathcal{D}_{\mathcal{F}}^{sup}$ set, again, with a weighted cross-entropy loss.

IV. EXPERIMENTAL SETUP

A. Datasets

As MAML requires a bundle of tasks to properly learn to learn from few-shot examples, we constructed a Meta-Dataset of radiological image segmentation tasks from publicly available datasets. More specifically, we build this Meta-Dataset using six Chest X-Ray (CXR) [3], [22]–[25], two Mammographic X-Ray (MXR) [5], [8] and two Dental X-Ray (DXR) [6], [7] datasets. Some of these datasets contain segmentation masks for multiple organs, which were all included in the Meta-Dataset.

In Table I, we list all datasets used in the experiments. Some datasets present more than one class, and in these cases the datasets were binarized to construct the tasks. To define a task, we select a dataset and one class as foreground. The pixels of all remaining classes are treated as background.

In order to assess the performance of WeaSeL in a certain setting, we employ a Leave-One-Task-Out methodology. That is, all tasks but the pair (*dataset, class*) chosen as the Few-shot task (\mathcal{F}) are used in the Meta-Dataset, reserving \mathcal{F} for the tuning/testing phase. This strategy serves to simultaneously hide the target task from the meta-training, while also allowing the experiments to evaluate the proposed algorithm and baselines in a myriad of scenarios. Such scenarios include: 1) target tasks with both large (e.g. JSRT [2]) and small (e.g. Montgomery [3]) domain shifts compared to the ones in the meta-training set; 2) \mathcal{F} tasks with image samples seen in other tasks used during meta-training, but with different target classes (e.g. JSRT [2]); and 3) \mathcal{F} tasks with foreground class absent in all other tasks used during meta-training (e.g. JSRT heart segmentation).

B. Architecture and Hyperparameters

Due to, mostly, the large memory use necessary for second-order optimization, we propose an architecture called miniUNet¹. It is a adaptation of the usual U-Net architecture with minor changes. The network is comprised of three encoder blocks, a center block, three decoder blocks and a 1×1 convolutional layer that works as a pixel-classification layer.

¹The code for WeaSeL, including the miniUNet architecture, will be available on GitHub upon the publication of this paper.

TABLE I
LIST OF DATASETS INCLUDED IN THE META-DATASET.

Dataset	Image Type	# of Images	Classes
JSRT [22]	X-rays	247	Lungs, Clavicles and Hearts
Montgomery [3]	X-rays	138	Lungs
Shenzhen [3]	X-rays	662	Lungs
NIH-labeled [24]	X-rays	100	Lungs
OpenIST [23]	X-rays	225	Lungs
LIDC-IDRI-DRR [25]	CT-scans	835	Ribs
MIAS [5]	Mammograms	322	Pectoral Muscle, Breasts
INbreast [8]	Mammograms	410	Pectoral Muscle, Breasts
Panoramic [7]	X-rays	116	Mandibles
UFBA-UESC [6]	X-rays	1500	Teeths

Similar to the U-Net architecture skip connections between symmetric layers are present in miniUNet, with each decoder block receiving the concatenation of the last block output and its corresponding encoder output as inputs.

C. Evaluation Protocol and Metrics

We use a 5-fold cross validation protocol in the experiments, wherein datasets were divided in training and validation sets for each fold. Support sets are obtained from the training partition, while the query sets are the entire validation partition. All images and labels are resized to 128×128 prior to being fed to the models in order to standardize the input size and minimize the memory footprint of MAML on high-dimensional outputs. The metric within a fold is computed for all images in the query set according to the dense labels, resulting in the final values reported in Section V, which are computed by averaging the performance across all folds. The metric used is the Jaccard score (or Intersection over Union – IoU) of the validation images, a common metric for semantic segmentation.

During training, we use the Adam optimizer with learning rate 0.001, weight decay 0.0005, and momentum 0.9. Our batch size was set to 5. The number of tasks sampled for the inner loop of WeaSeL was set to 6 in the experiments.

D. Baselines and Sparsity Modalities

We use two baselines: 1) *From Scratch* training on the sparse labels; and 2) *Fine-Tuning* a pretrained model from a source dataset with *dense* labels. During the meta-training phase of WeaSeL, the tuning phase of *Fine-Tuning*, and *From Scratch* training, the sparse labels are simulated for each sample from their the dense mask.

Although being a few-shot sparse segmentation method, we do not present the Guided Nets [20] as a baseline. Even to the best of our efforts, the episodic training of the original model was not able to converge to a usable model with the same medical Meta-Dataset used by WeaSeL. Thus, it did not seem fair to compare the Guided Nets to our approach.

We evaluate two types of sparse labeling, namely, *points* and *grid*, as well as their comparison to the performance of the models trained on the full masks. The *points* labels are a

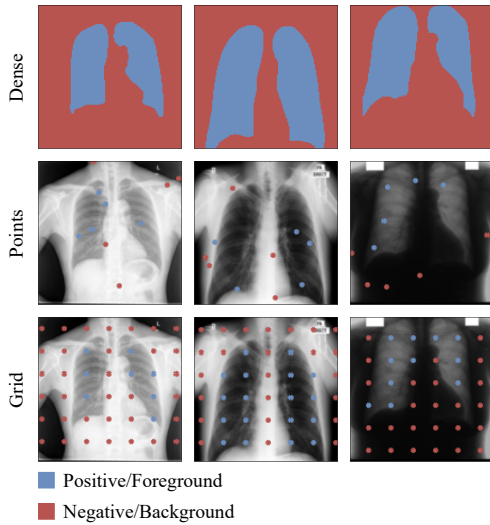


Fig. 2. Example of the evaluated type of annotations. **Top Row**: the ground truth labels of all pixels. **Middle Row**: the *points* annotation, $n(5)$ pixels of background/foreground are labeled. **Bottom Row**: the *grid* annotation, pixels spaced by $s(20)$ are selected and properly labeled.

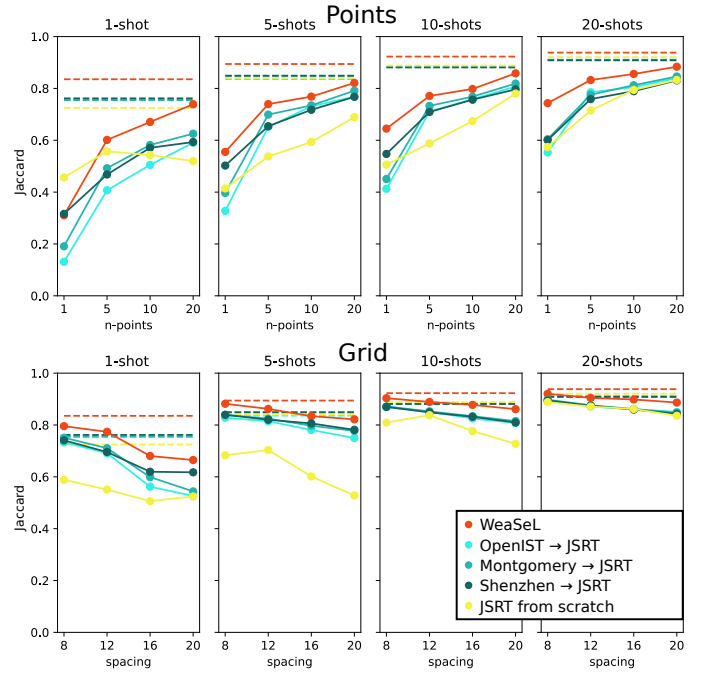
selection from pixels of the image, where the annotator alternately chooses pixels from the foreground and background. In the *grid* labeling the annotator receives a pre-selected group of pixels and change the class of the ones they consider positive. These pixels are disposed in a grid manner in the image. A visualization of these styles are presented in Figure 2.

We simulate this two types of annotations from ground truth labels. For the *points* labels, given a parameter n , we randomly select n pixels from the foreground class, and n from the background. For the *grid* annotation, given a parameter s , we choose pixels spaced, horizontally and vertically, by s , starting from a random distance in the range $(0-s)$ from the upper-left corner. For a consistent evaluation, a random seed for the Few-Shot \mathcal{F} task labels is fixed for each image, ensuring that all methods use the same sparse labels for the fine-tuning process. In experiments, we vary the parameters n and s , and analyze the impact of the sparsity of the labels.

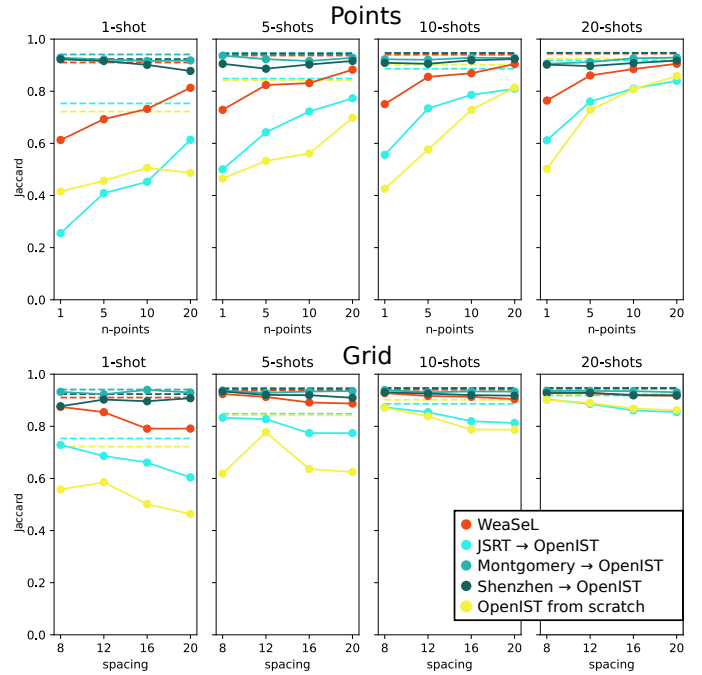
V. RESULTS

We perform experiments to access the performance of WeaSeL in all CXR datasets. That is, we evaluate the Few-shot task segmentation of lungs in all five datasets, and hearts and clavicles segmentation in the JSRT dataset. For brevity, we include only a subset of the results.

In Figure 3 we see the results of experiments in the lungs class. The results of Montgomery and Shenzhen datasets are similar to the OpenIST, thus we only present the later. WeaSeL shows better metrics than the *From Scratch* baseline in all cases. As expected, increasing the amount of annotated pixels (e.g. by increasing n , or decreasing the grid spacing s), as well as having more training samples (larger k -shots), have a direct impact in IoU. For JSRT (Figure 3a), WeaSeL yields better performance than *Fine-Tuning* from all datasets, in contrast with OpenIST (Figure 3b) where the source dataset



(a) JSRT Lungs



(b) OpenIST Lungs

Fig. 3. Jaccard results for lung segmentation in two target datasets: JSRT (a) and OpenIST (b). Solid lines indicate the performance of the methods using sparse label (*Points* on the top and *Grid* on the bottom), while the dashed line presents the performance of the methods trained with the dense masks.

was decisive in the performance of *Fine-Tuning*. The networks pretrained on JSRT achieved performances lower than WeaSeL, while the pretraining from Shenzhen or Montgomery yielded better results than our approach. This discrepancy can be explained by the different domains, as the JSRT dataset is the most singular within the five CXR datasets – that is, the

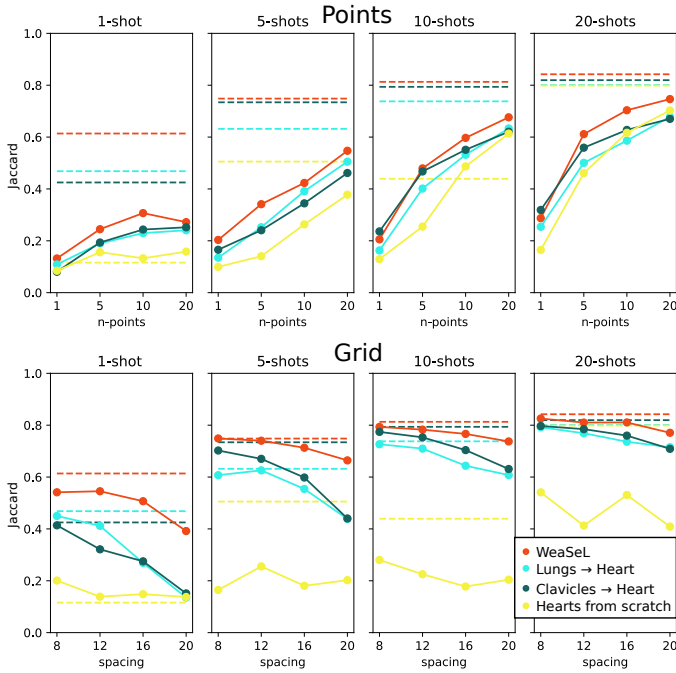


Fig. 4. Jaccard results for heart segmentation in JSRT dataset.

domain shift between JSRT and the other datasets is larger than the shift between the other four. It should be noticed that real-world scenarios may not allow for a fair evaluation of the domain shift between the source and target datasets, making WeaSeL the safer choice in this case.

In Figure 4 we see the results for the heart class in JSRT. This, and the clavicles class, are only present in the JSRT dataset. Since these classes are only present in JSRT, for the *Fine-Tuning* baselines we choose the pairs of JSRT dataset and one of the remaining classes as source tasks (e.g., for hearts we use the pairs $(JSRT, lungs)$, and $(JSRT, clavicles)$). Again, the *From Scratch* baseline is the worst performer, with the WeaSeL being the superior in most scenarios. The clavicles, display a similar tendency, although their IoU score, in majority of the cases, is lower than the heart class scores.

As seen in both experiments, the performance of the *grid* annotation is higher than the *points* annotation, in virtually all cases. In order to assess the efficiency of this annotations we constructed the graph in Figure 5. The x axis represents the average number of inputs per image, and was computed considering each positive labeled pixel in an image as an user input, and then average across all images. In the figure, we confirm the overall better performance of the *grid* annotation, and a better efficiency as the score of the method with the same number of inputs is higher using this annotation over the *point* annotation. One explanation is that even in the larger spacing scenario ($s = 20$) the number of annotated pixels is greater than the the $2n$ pixels in any *points* labeling case. Thus, even with a likely class imbalance in the *grid* scenario — the foreground objects is usually smaller than the background, and with a fixed grid will probably have fewer labeled pixels —,

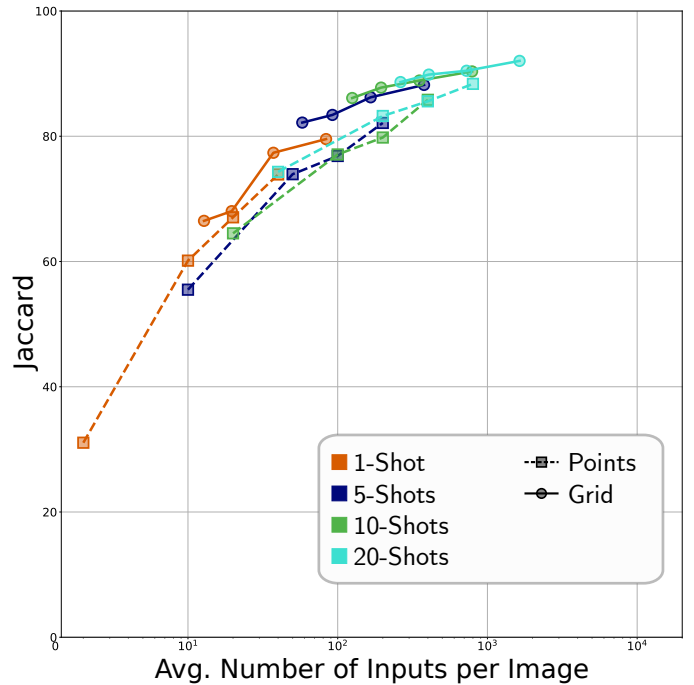


Fig. 5. Jaccard results of the WeaSeL method in the JSRT Lungs task by the average number of inputs.

the simple presence of more data increases the metric score for this type of annotation.

Sparse annotation show competitive results with dense labels, given the much smaller labeled data. The case with most labeled data is the *grid* annotation with $s = 20$, that have approximately 2% of annotated pixels. In some cases, the methods using sparse labels are equivalent to using dense labels ground truths, specially with the *grid* annotation. In the lungs segmentation in the OpenIST (and Montgomery/Shenzhen) dataset (Figure 3b), the *Fine-tuning* baselines from similar datasets seem indifferent to sparse or dense labels. As aforementioned, the small domain shifts between these datasets and the pretraining with dense labels explain this results.

VI. QUALITATIVE ANALYSIS

In Figures 6-8 we show visual segmentation examples of WeaSeL on three tasks: JSRT Lungs, OpenIST Lungs and JSRT Heart, respectively. In each example, lines represent the number of shots in the experiment, while columns vary the sparsity parameters: n for points (first four columns) and s for grids (last four columns). These sparsity scenarios were used to simulate the sparse labels of the support set on the target few-shot task.

As one can observe in these visual results, increasing the annotation density — that is, increasing the number of points n or using a smaller spacing s — yields a larger improvement in the predictions than increasing the number of labeled images; that is, the number of shots. Even the 1-shot scenario achieved acceptable results for the lung tasks (Figures 6, 7) when given

a sufficient number of labeled pixels, e.g. at least 20 labeled pixels for each class. For heart segmentation (Figure 8), we observe a great difference between the points and grid annotations. Since this is a unique task – i.e. the only task in the meta-dataset with *heart* as a target class, which is only present in the JSRT dataset – the model requires more support samples/labels to adapt to this task. This requirement is not fully achieved by the points annotation with one and five shots, but is obtained by less sparse grid annotations in all observed number of shots. Again, the total number of labeled pixels in the grid annotation seems to compensate the lack of images for training.

It is clear that having a larger number of annotated images impacts the final results, especially for harder, or more distinct, datasets or tasks as is the case of JSRT. Consequently, a middle ground in relation to the number of images and the sparsity of annotations seems to be a good compromise to achieve reasonable results with limited human intervention. That is, annotating five or more images with at least five pixels labeled per class can lead to good results with a very low labeling burden. However, if the target task is truly few-shot, increasing the number of samples in the support set may be costly or even impossible. In this case, one can instead increase the per sample label density.

VII. CONCLUSION

We evaluated WeaSeL, our proposed adaptation of MAML, in multiple few-shot segmentation tasks with sparse labels. WeaSeL showed to be a viable solution to most tasks, and a suitable option when there are very discrepant target and source tasks. Since in real-world scenarios is difficult to access the domain shift between the datasets, WeaSeL seems to be the safer choice in these cases. The method have some limitations, like the cost of computing second-derivatives in the outer loop. This limited the size of training images to 128×128 , which can hinder task with small target objects (e.g. clavicles).

Future works include the use of approximations to the second derivative and how this affects the performance, as well as analyzing different types of sparse labels annotations also in 3D volumes. Additionally, we intend to adapt other first- and second-order Meta-Learning methods for sparsely labeled semantic segmentation.

ACKNOWLEDGMENTS

This work was supported by the Serrapilheira Institute (grant number Serra – R-2011-37776). This work was supported in part by the Minas Gerais Research Funding Foundation (FAPEMIG) under Grant APQ-00449-17, by the National Council for Scientific and Technological Development (CNPq) under Grant 311395/2018-0 and Grant 424700/2018-2, and by the *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES)* – Finance Code 001. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

REFERENCES

- [1] N. Sharma and L. M. Aggarwal, “Automated medical image segmentation techniques,” *Journal of Medical Physics/Association of Medical Physicists of India*, vol. 35, no. 1, p. 3, 2010.
- [2] J. Shiraishi, S. Katsuragawa, J. Ikezoe, T. Matsumoto, T. Kobayashi, K.-i. Komatsu, M. Matsui, H. Fujita, Y. Kodera, and K. Doi, “Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists’ detection of pulmonary nodules,” *American Journal of Roentgenology*, vol. 174, no. 1, pp. 71–74, 2000.
- [3] S. Jaeger, S. Candemir, S. Antani, Y.-X. J. Wang, P.-X. Lu, and G. Thoma, “Two public chest x-ray datasets for computer-aided screening of pulmonary diseases,” *Quantitative Imaging in Medicine and Surgery*, vol. 4, no. 6, p. 475, 2014.
- [4] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, “Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases,” in *CVPR*, 2017, pp. 2097–2106.
- [5] J. Suckling *et al.*, “The mammographic image analysis society digital mammogram database digital mammography ed ag gale, sm astley, dr dance and ay cairns,” 1994.
- [6] G. Silva, L. Oliveira, and M. Pithon, “Automatic segmenting teeth in x-ray images: Trends, a novel data set, benchmarking and future perspectives,” *Expert Systems with Applications*, vol. 107, pp. 15–31, 2018.
- [7] A. H. Abdi, S. Kasaei, and M. Mehdizadeh, “Automatic segmentation of mandible in panoramic x-ray,” *Journal of Medical Imaging*, vol. 2, no. 4, p. 044003, 2015.
- [8] I. C. Moreira, I. Amaral, I. Domingues, A. Cardoso, M. J. Cardoso, and J. S. Cardoso, “Inbreast: toward a full-field digital mammographic database,” *Academic Radiology*, vol. 19, no. 2, pp. 236–248, 2012.
- [9] J. Snell, K. Swersky, and R. Zemel, “Prototypical networks for few-shot learning,” in *NIPS*, 2017, pp. 4077–4087.
- [10] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *ICML*. JMLR. org, 2017, pp. 1126–1135.
- [11] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *MICCAI*. Springer, 2015, pp. 234–241.
- [12] N. Tajbakhsh, L. Jeyaseelan, Q. Li, J. N. Chiang, Z. Wu, and X. Ding, “Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation,” *Medical Image Analysis*, p. 101693, 2020.
- [13] P. Goyal, D. Mahajan, A. Gupta, and I. Misra, “Scaling and benchmarking self-supervised visual representation learning,” in *ICCV*, 2019, pp. 6391–6400.
- [14] G. Wang, W. Li, M. A. Zuluaga, R. Pratt, P. A. Patel, M. Aertsen, T. Doel, A. L. David, J. Deprest, S. Ourselin *et al.*, “Interactive medical image segmentation using deep learning with image-specific fine tuning,” *TMI*, vol. 37, no. 7, pp. 1562–1573, 2018.
- [15] A. Shaban, S. Bansal, Z. Liu, I. Essa, and B. Boots, “One-shot learning for semantic segmentation,” in *Proceedings of the British Machine Vision Conference (BMVC)*, September 2017, pp. 167.1–167.13.
- [16] N. Dong and E. Xing, “Few-shot semantic segmentation with prototype learning,” in *BMVC*, vol. 3, no. 4, 2018.
- [17] T. Hu, P. Yang, C. Zhang, G. Yu, Y. Mu, and C. G. Snoek, “Attention-based multi-context guiding for few-shot semantic segmentation,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 8441–8448.
- [18] X. Zhang, Y. Wei, Y. Yang, and T. S. Huang, “Sg-one: Similarity guidance network for one-shot semantic segmentation,” *IEEE Transactions on Cybernetics*, vol. 50, no. 9, pp. 3855–3865, 2020.
- [19] K. Wang, J. H. Liew, Y. Zou, D. Zhou, and J. Feng, “Panet: Few-shot image semantic segmentation with prototype alignment,” in *ICCV*, 2019, pp. 9197–9206.
- [20] K. Rakelly, E. Shelhamer, T. Darrell, A. A. Efros, and S. Levine, “Few-shot segmentation propagation with guided networks,” *CoRR*, vol. abs/1806.07373, 2018.
- [21] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, “The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results,” 2012.
- [22] J. Shiraishi, S. Katsuragawa, J. Ikezoe, T. Matsumoto, T. Kobayashi, K.-i. Komatsu, M. Matsui, H. Fujita, Y. Kodera, and K. Doi, “Development of a digital image database for chest radiographs with and without a




Fig. 6. Visual segmentation examples for the JSRT Lungs task.

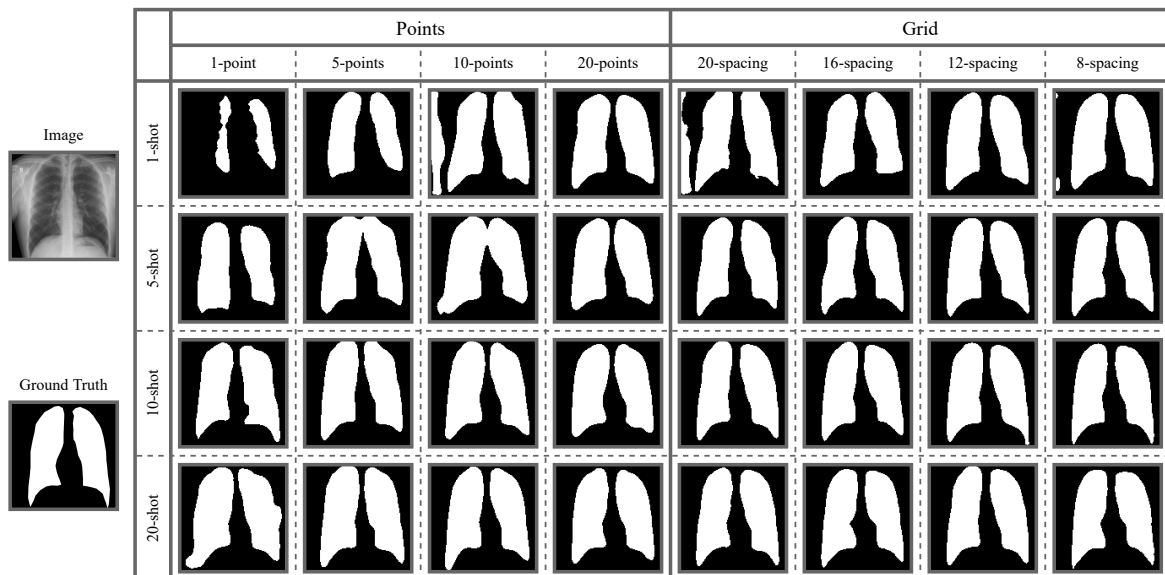


Fig. 7. Visual segmentation examples for the OpenIST Lungs task.

lung nodule: receiver operating characteristic analysis of radiologists' detection of pulmonary nodules," *American Journal of Roentgenology*, vol. 174, no. 1, pp. 71–74, 2000.

- [23] A. Taranov, "Openist: A set of open source tools (c classes and cmd-utils) for image segmentation and classification," <https://github.com/pinull-mezon/OpenIST/tree/master/Datasets>.
- [24] Y. Tang, Y. Tang, J. Xiao, and R. M. Summers, "Xlsor: A robust and accurate lung segmentor on chest x-rays using criss-cross attention and customized radiorealistic abnormalities generation," 2019.
- [25] H. Oliveira, V. Mota, A. M. Machado, and J. A. dos Santos, "From 3d to 2d: Transferring knowledge for rib segmentation in chest x-rays," *Pattern Recognition Letters*, vol. 140, pp. 10–17, 2020.

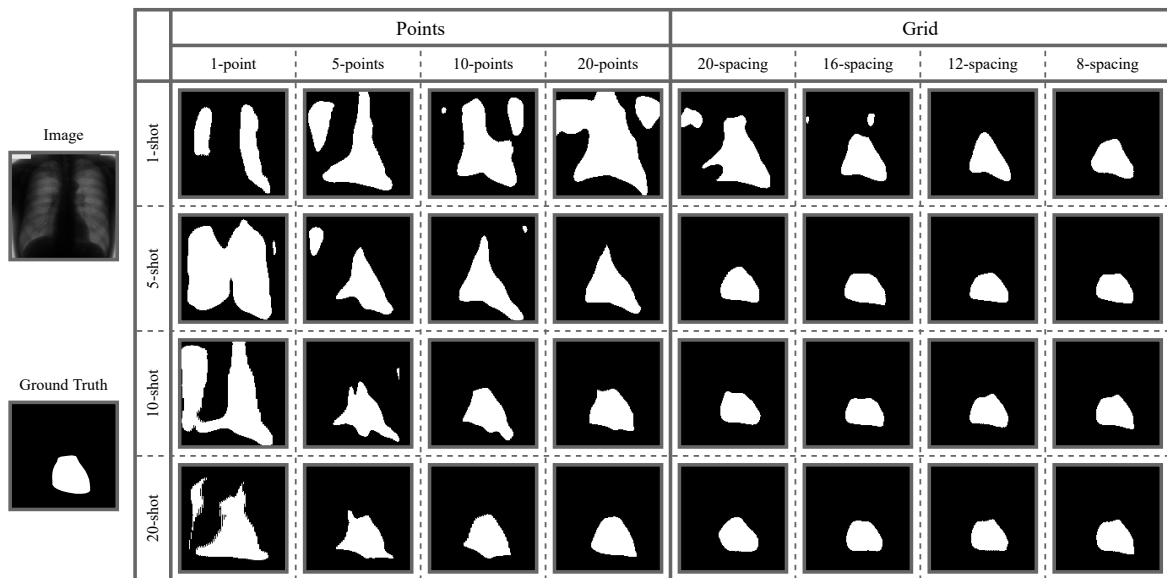


Fig. 8. Visual segmentation examples for the JSRT Heart task.