

# PLESS: PSEUDO-LABEL ENHANCEMENT WITH SPREADING SCRIBBLES FOR WEAKLY SUPERVISED SEGMENTATION

Yeva Gabrielyan<sup>1</sup>, Varduhi Yeghiazaryan<sup>1</sup>, and Irina Voiculescu<sup>2</sup>

<sup>1</sup>Akian College of Science and Engineering, American University of Armenia, Yerevan, Armenia

<sup>2</sup>Department of Computer Science, University of Oxford, Oxford, UK

## ABSTRACT

Weakly supervised learning with scribble annotations uses sparse user-drawn strokes to indicate segmentation labels on a small subset of pixels. This annotation reduces the cost of dense pixel-wise labeling, but suffers inherently from noisy and incomplete supervision. Recent scribble-based approaches in medical image segmentation address this limitation using pseudo-label-based training; however, the quality of the pseudo-labels remains a key performance limit. We propose PLESS, a generic pseudo-label enhancement strategy which improves reliability and spatial consistency. It builds on a hierarchical partitioning of the image into a hierarchy of spatially coherent regions. PLESS propagates scribble information to refine pseudo-labels within semantically coherent regions. The framework is model-agnostic and easily integrates into existing pseudo-label methods. Experiments on two public cardiac MRI datasets (ACDC and MSCMRseg) across four scribble-supervised algorithms show consistent improvements in segmentation accuracy. Code will be made available on GitHub upon acceptance.

**Index Terms**— Image segmentation, weak supervision, scribble supervision, pseudo-labels, scribble spreading

## 1. INTRODUCTION

To reduce the cost of dense pixel-wise annotation, the segmentation community has increasingly turned to weakly supervised learning methods. These have been explored using sparse forms of annotation, including image-level labels, bounding boxes, key points, and scribbles [1, 2, 3, 4]. Among these, scribbles offer a balance between annotation efficiency and segmentation quality, and have also gained traction in medical image analysis [4, 5, 6].

Existing scribble-supervised methods expand sparse annotations using graphical-model-based models [4] or regularization strategies as loss functions [7], and more recent approaches leverage adversarial learning or teacher–student frameworks to improve structural consistency [4, 5, 7, 8, 9].

Mixup augmentation methods such as CycleMix [10] generate blended images, enhancing the training goal with consistency losses to improve segmentation reliability.

**Pseudo-label-based methods.** Pseudo-labels can be used to provide supervision for unlabeled images or pixels. During training, pseudo-labels can be generated from either model predictions or weak annotations, and treated as supervisory signals. The central idea is to transform imperfect or sparse annotations into higher-quality and more reliable training targets, enabling effective learning from limited supervision [11]. Recent studies have shown that semi-supervised and weakly supervised learning can benefit significantly from high-quality pseudo-labels [12, 13].

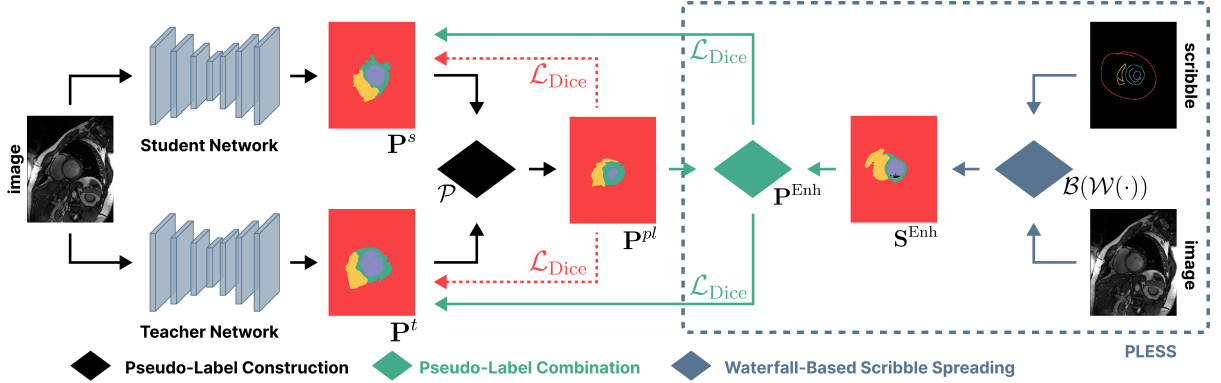
**Contributions.** This work proposes a generic pseudo-label enhancement with spreading scribbles (PLESS) to improve the effectiveness of pseudo-label-based training. It is designed as a model-agnostic framework that can be integrated with existing pseudo-label losses to strengthen supervision quality. The approach is evaluated on four representative scribble-supervised algorithms: DMPLS [14], DCDPL [15], ScribbleVC [16], and ScribbleVS [17]. Our method builds on a hierarchical image partitioning, obtained via watershed and waterfall transforms [18], that partitions the image into a hierarchy of regions and enables structured propagation of scribble information. We show how this structure can be further exploited to enhance pseudo-label quality in weakly supervised segmentation. Comprehensive experiments on two cardiac MRI datasets demonstrate that PLESS consistently improves pseudo-label-based methods, and that ScribbleVS enhanced with PLESS outperforms state-of-the-art (SOTA) approaches. Finally, extensive ablation studies analyze the impact of different components and configurations of PLESS.

## 2. PSEUDO-LABEL ENHANCEMENT

Following [19], a waterfall-based scribble spreading strategy is used to construct spatially enhanced labels. The original approach applies scribble spreading directly to the input scribble annotations to generate training supervision. In contrast, PLESS applies waterfall-based scribble spreading to the constructed pseudo-labels during training.

---

This work was supported by the Afeyan Family Foundation Seed Grants and the JACE Foundation Research Innovation Grant Program at AUA.



**Fig. 1.** Overview of the proposed PLESS framework, showing pseudo-label construction, waterfall-based enhancement, and integration into pseudo-label loss. The red dotted arrows indicate the original pseudo-label DSC loss formulation used prior to PLESS, which is replaced by the green-arrow loss terms introduced by the proposed enhancement strategy.

Each 2D slice is partitioned into a hierarchy of semantically coherent regions. Label classes can denote different anatomical structures of interest, or the background. The labels are then spread inside the coarse layers, then inside the finer layers of the partition hierarchy, as described in [19]. In each partitioning layer, regions containing a unique class candidate are assigned that label; regions with conflicting class assignments remain unlabeled. This hierarchical procedure propagates labels within homogeneous regions while avoiding propagation (‘leaks’) across ambiguous boundaries.

In order to avoid further label leakage into the background, we introduce an explicit background expansion step to refine the enhanced labels. This step identifies isolated unlabeled regions that are entirely surrounded by background or unlabeled pixels and assigns them to the background label; it also preserves unlabeled regions that touch foreground label boundaries. This operation removes isolated gaps and holes near the image borders and improves the spatial consistency of the enhanced supervision.

The overall PLESS framework is illustrated in Fig. 1, where pseudo-labels derived from student and teacher predictions are refined by the PLESS framework via waterfall-based spreading of the scribbles.

Mathematically, let the pseudo-label construction be defined as  $\mathbf{P}^{pl} = \mathcal{P}(\mathbf{P}^s, \mathbf{P}^t)$ , where  $\mathbf{P}^s$  and  $\mathbf{P}^t$  denote the outputs of the student and teacher networks, respectively, and  $\mathcal{P}(\cdot)$  is a network-specific fusion operator that combines the two predictions. The exact form of  $\mathcal{P}$  depends on the underlying method and may involve, for example, additive fusion, confidence-based selection, or some other task- or data-specific strategy defined by each algorithm.

The waterfall-based scribble spreading is defined as  $\mathbf{S}^w = \mathcal{W}(\mathbf{S}, \mathbf{I})$ , where  $\mathbf{S}$  denotes the original scribble annotations and  $\mathbf{I}$  is the input image, and  $\mathcal{W}(\cdot)$  denotes the hierarchical waterfall-based spreading operator. To take the enhancement even further, a background expansion operator is applied:  $\mathbf{S}^{Enh} = \mathcal{B}(\mathbf{S}^w)$ , where  $\mathcal{B}(\cdot)$  fills isolated unlabeled regions with the background label while preserving uncertain areas near foreground boundaries.

Based on the constructed pseudo-labels  $\mathbf{P}^{pl}$  and the enhanced scribbles  $\mathbf{S}^{Enh}$ , PLESS defines an enhanced pseudo-label  $\mathbf{P}^{Enh}$  as

$$\mathbf{P}^{Enh} = \begin{cases} \mathbf{P}^{pl} \odot (1 - \mathbf{M}) + \mathbf{S}^{Enh} \odot \mathbf{M} & \text{if } e \leq \tau E_{max} \\ \mathbf{P}^{pl} & \text{otherwise} \end{cases} \quad (1)$$

where  $\mathbf{M} = \mathbf{I}[\mathbf{S}^{Enh}]$  is a binary mask indicating pixels with

**Table 1.** Comparison of PLESS against the baseline for the ACDC dataset, using DSC, HD95 (mm) and ASD (mm). Bold indicates the best result for each method, and underlining denotes the overall best.

	Original						PLESS					
	DSC				HD95	ASD	DSC				HD95	ASD
	LV	RV	MC	Avg	Avg	Avg	LV	RV	MC	Avg	Avg	Avg
DMPLS [14]	.881	.847	.914	.881	9.37	2.53	.886	.851	.923	<b>.887</b>	<b>7.10</b>	<b>2.01</b>
	$\pm .002$	$\pm .005$	$\pm .006$	$\pm .003$	$\pm 2.94$	$\pm 0.53$	$\pm .004$	$\pm .005$	$\pm .005$	$\pm .004$	$\pm 1.63$	$\pm 0.53$
DCDPL [15]	.889	.857	.916	<b>.887</b>	8.36	2.45	.863	.823	.898	.861	<b>6.26</b>	<b>1.73</b>
	$\pm .003$	$\pm .007$	$\pm .004$	$\pm .004$	$\pm 1.83$	$\pm 0.8$	$\pm .011$	$\pm .006$	$\pm .010$	$\pm .008$	$\pm 1.22$	$\pm 0.37$
ScribbleVC [16]	.838	.839	.904	.860	5.94	1.42	.858	.845	.907	<b>.870</b>	<b>4.95</b>	<b>1.16</b>
	$\pm .025$	$\pm .006$	$\pm .010$	$\pm .009$	$\pm 0.57$	$\pm 0.21$	$\pm .012$	$\pm .007$	$\pm .012$	$\pm .007$	$\pm 1.01$	$\pm 0.27$
ScribbleVS [17]	.881	.854	.917	.884	9.77	2.71	.878	.861	.925	<b>.888</b>	<b>6.39</b>	<b>1.66</b>
	$\pm .015$	$\pm .016$	$\pm .014$	$\pm .015$	$\pm 6.86$	$\pm 1.92$	$\pm .014$	$\pm .011$	$\pm .008$	$\pm .009$	$\pm 3.13$	$\pm 0.83$

**Table 2.** Comparison of PLESS against the baseline for the MSCMRseg dataset, using DSC, HD95 (mm) and ASD (mm). Bold indicates the best result for each method, and underlining denotes the overall best.

	Original						PLESS					
	DSC				HD95	ASD	DSC				HD95	ASD
	LV	RV	MC	Avg	Avg	Avg	LV	RV	MC	Avg	Avg	Avg
DMPLS [14]	.888	.845	.924	<b>.886</b>	8.81	2.69	.881	.841	.925	.882	<b>5.40</b>	<b>1.67</b>
	$\pm .003$	$\pm .003$	$\pm .003$	$\pm .002$	$\pm 0.99$	$\pm 0.60$	$\pm .003$	$\pm .002$	$\pm .002$	$\pm .001$	$\pm 0.69$	$\pm 0.46$
DCDPL [15]	.853	.814	.921	<b>.862</b>	8.21	2.78	.818	.764	.891	.825	<b>7.07</b>	<b>2.05</b>
	$\pm .008$	$\pm .004$	$\pm .003$	$\pm .002$	$\pm 2.52$	$\pm 0.49$	$\pm .028$	$\pm .034$	$\pm .010$	$\pm .024$	$\pm 1.53$	$\pm 0.47$
ScribbleVC [16]	.822	.812	.919	<b>.851</b>	<b>6.58</b>	<b>1.69</b>	.757	.747	.912	.805	7.54	1.71
	$\pm .051$	$\pm .017$	$\pm .003$	$\pm .022$	$\pm 0.95$	$\pm 0.20$	$\pm .031$	$\pm .050$	$\pm .011$	$\pm .029$	$\pm 1.62$	$\pm 0.32$
ScribbleVS [17]	.876	.844	.925	<b>.882</b>	7.61	2.07	.871	.848	.928	<b>.882</b>	<b>5.53</b>	<b>1.33</b>
	$\pm .007$	$\pm .002$	$\pm .004$	$\pm .003$	$\pm 4.35$	$\pm 1.16$	$\pm .007$	$\pm .005$	$\pm .003$	$\pm .004$	$\pm 1.09$	$\pm 0.23$

enhanced scribble labels,  $\odot$  denotes element-wise multiplication,  $e$  is the current training epoch,  $E_{\max}$  is the total number of training epochs, and  $\tau$  is a tolerance parameter controlling the fraction of training during which pseudo-label scribble spreading is applied.

When using the enhanced pseudo-labels, the Dice-based pseudo-label loss is defined as

$$\mathcal{L}_{\text{pl}} = \frac{1}{2} (\mathcal{L}_{\text{DSC}}(\mathbf{P}^{\text{Enh}}, \mathbf{P}^{\text{s}}) + \mathcal{L}_{\text{DSC}}(\mathbf{P}^{\text{Enh}}, \mathbf{P}^{\text{t}})). \quad (2)$$

### 3. EXPERIMENTAL SETUP

**Dataset.** We evaluate our pseudo-label scribble spreading framework on two scribble-annotated cardiac MRI datasets. (a) The ACDC dataset contains cine-MRI scans from 150 patients. Scribble annotations are available for 100 cases, which are used in five-fold cross-validation with 80 scans for training and 20 for validation. The remaining 50 cases without scribbles are used for testing. (b) The MSCMRseg dataset includes late gadolinium enhancement MRI scans from 45 cardiomyopathy patients. The data are split into 25 training, 5 validation, and 15 test scans. All results are averaged over five runs.

Both datasets are annotated for three classes: left ventricle (LV), right ventricle (RV), and myocardium (MYO). All results are reported as averages over these classes.

**Evaluation Measures.** During testing, 2D slice predictions are stacked to form 3D volumes. The model with the

best validation performance is used for inference. Segmentation quality is evaluated using 3D Dice (DSC), 95th percentile Hausdorff Distance (HD95), and Average Surface Distance (ASD) [23], measured in mm. Higher DSC and lower HD95 and ASD indicate better performance.

**Implementation Details.** We follow the protocol of [14]. Slices are normalized to  $[0, 1]$ , and standard augmentations (rotation, flipping, noise) are applied. Batch sizes are 12 for DCDPL and ScribbleVS, 6 for DMPLS and ScribbleVC. Models are trained for 60K iterations in PyTorch on NVIDIA RTX 4090 GPUs. Aside from the pseudo-label loss, all other losses, e.g., cross-entropy, remain unchanged.


### 4. RESULTS

Performance comparisons between baseline and PLESS are reported in Table 1 for the ACDC dataset and Table 2 for the MSCMRseg dataset. On ACDC, the PLESS setup improves performance across all three measures for all algorithms, except for DCDPL, where improvements are observed for HD95 and ASD only. On MSCMRseg, PLESS leads to consistent improvements in HD95 and ASD for all methods except ScribbleVC.

For ACDC, the best DSC of .888 is achieved by ScribbleVS with the PLESS setup, while the lowest HD95 of 4.95 mm and ASD of 1.16 mm are obtained by ScribbleVC with PLESS. On MSCMRseg, although the highest DSC of .886 is achieved using the original setup, the best HD95

**Table 3.** Comparison with SOTA on the individual features LV, MYO, RV, and their average, evaluated with the Avg DSC.

	ACDC				MSCMRseg			
	LV	MYO	RV	Avg	LV	MYO	RV	Avg
CutMix [20]	.641	.734	.740	.705	.578	.622	.761	.654
	$\pm .359$	$\pm .144$	$\pm .216$		$\pm .063$	$\pm .121$	$\pm .105$	
Puzzle Mix [21]	.663	.650	.559	.624	.061	.634	.028	.241
	$\pm .333$	$\pm .231$	$\pm .343$		$\pm .021$	$\pm .084$	$\pm .012$	
Co-Mixup [22]	.622	.621	.702	.648	.356	.343	.053	.251
	$\pm .304$	$\pm .214$	$\pm .211$		$\pm .075$	$\pm .067$	$\pm .022$	
CycleMix [10]	<b>.883</b>	.798	.863	.848	.870	.739	.791	.800
	$\pm .095$	$\pm .075$	$\pm .073$		$\pm .061$	$\pm .049$	$\pm .072$	
ScribbleVS PLESS	.878	<b>.861</b>	<b>.925</b>	<b>.888</b>	<b>.871</b>	<b>.848</b>	<b>.928</b>	<b>.882</b>
	$\pm .014$	$\pm .011$	$\pm .008$	$\pm .009$	$\pm .007$	$\pm .005$	$\pm .003$	$\pm .004$



**Fig. 2.** Comparison of segmentation performance using the baseline and PLESS on three scans (a, b, c) from the ACDC dataset. Row 1 displays the original MRI slice and ground truth (GT), while subsequent rows show segmentation results from different models: DMPLS [14], DCDPL [15], ScribbleVC [16], and ScribbleVS [17]. Reported DSC, HD95, and ASD scores refer to the whole 3D scan. Label colors: background ■, right ventricle (RV) ■, myocardium (MYO) ■, and left ventricle (LV) ■.

(5.53 mm) and ASD (1.33 mm) are still obtained with PLESS. These results demonstrate that PLESS consistently improves boundary accuracy and surface agreement, leading to more reliable segmentations across multiple architectures.

Further comparison with SOTA methods that do not employ pseudo-label losses is reported in Table 3. On both the ACDC and MSCMRseg datasets, the highest average DSC is achieved by the ScribbleVS algorithm using PLESS. It is important to note that the SOTA methods in Table 3 follow a different experimental protocol, where only the training split of the dataset is used and further divided into separate training,

validation, and test sets. In contrast, our experiments follow the standard dataset splits commonly used for fully supervised segmentation, and we use the original training, validation, and test partitions provided by each dataset.

A qualitative comparison of the performance of all four segmentation methods on the ACDC dataset is presented in Fig. 2. For each image and each method, results for both original and PLESS setups are shown. In most cases, PLESS leads to improvement in boundary delineation, evident by the improved ASD and HD95 scores.

The three examples in Fig. 2 illustrate representative fail-

**Table 4.** Ablation study of the different PLESS setups for the ScribbleVS network on the ACDC dataset.

	25%			50%			75%			100%		
	DSC	HD95	ASD	DSC	HD95	ASD	DSC	HD95	ASD	DSC	HD95	ASD
enh	<b>.888</b>	6.68	1.73	.885	6.41	1.89	.884	<b>6.33</b>	1.88	.869	10.78	2.93
	$\pm .012$	$\pm 3.51$	$\pm 0.82$	$\pm .010$	$\pm 1.02$	$\pm 0.21$	$\pm .015$	$\pm 2.20$	$\pm 0.62$	$\pm .007$	$\pm 3.86$	$\pm 1.20$
enh+bg	<b>.888</b>	6.39	<b>1.66</b>	.885	8.97	2.18	.886	6.42	1.71	.873	6.74	1.84
	$\pm .009$	$\pm 3.13$	$\pm 0.83$	$\pm .011$	$\pm 4.64$	$\pm 1.11$	$\pm .009$	$\pm 2.19$	$\pm 0.63$	$\pm .011$	$\pm 2.73$	$\pm 0.97$
enh+bg+prop	.887	6.90	1.93	.886	6.96	1.93	.886	7.26	1.96	.873	8.91	2.51
	$\pm .009$	$\pm 1.69$	$\pm 0.50$	$\pm .013$	$\pm 2.88$	$\pm 0.98$	$\pm .013$	$\pm 2.23$	$\pm 0.82$	$\pm .007$	$\pm 2.22$	$\pm 0.78$

ure cases of the original setup that are mitigated by PLESS. In example (a), for slices where a particular anatomical structure is not present, the original setup for DCDPL, ScribbleVC, and ScribbleVS frequently produces false positive predictions, most notably for the myocardium (MYO). In contrast, with the PLESS setup these errors are fixed.

In example (b), the original setup incorrectly segments regions that visually resemble the target structure, even though they are anatomically unrelated. This is observed for DMPLS, DCDPL, and ScribbleVS, where regions with similar appearance or geometry are mistakenly labeled. The PLESS setup substantially reduces these errors, suggesting that the propagated labels provide stronger spatial and contextual information, which helps the algorithms differentiate between similar-looking structures.

In example (c), when only two of the three cardiac structures are present in a slice, the original setup often hallucinates the missing third class. This behavior is observed across all four methods and reflects a common limitation of weak supervision, where models tend to over-predict plausible anatomical classes. With PLESS, these false positive segmentations are reduced. These qualitative results show that PLESS improves not only segmentation accuracy, but also the robustness of predictions in challenging and ambiguous cases.

#### 4.1. Ablation Studies


To identify the most effective PLESS configuration, an ablation study is conducted on two key components. Dif-

ferent strategies for scribble spreading are first analyzed, considering three progressive variants. The baseline scribble spreading (*enh*) corresponds to the original waterfall-based region expansion proposed in [19] and *enh+bg* is the introduced background expansion variant. Finally, full region propagation (*enh+bg+prop*) is applied, where the remaining unlabeled pixels are iteratively assigned labels from neighboring regions until full image coverage is achieved. The three different stages of PLESS are visualized in Fig. 3.

The second key component is the tolerance level. Table 4 reports the performance of the ScribbleVS network under different PLESS configurations and tolerance levels, ranging from 25% to 100%. The *enh+bg* variant achieves the best performance across two of the measures (DSC and ASD). Performance is strongest at 25% and 75%, while a degradation is observed at 100%. This suggests that aggressive scribble spreading at full tolerance can introduce noise, as the enhanced annotations may contain inaccuracies and ambiguities. At later training stages, pseudo-labels generated by the teacher and student networks can be more reliable than the propagated enhancements themselves. This leads to better performance when scribble spreading is not applied in the later episodes. In most cases, the variant without full propagation (*enh+bg*) outperforms *enh+bg+prop*, indicating that complete label propagation introduces incorrect labels in uncertain regions, which degrades performance.

## 5. CONCLUSION

The PLESS framework exploits spatial structure to improve the reliability of pseudo-label supervision in weakly supervised segmentation. Rather than introducing additional model complexity, PLESS operates at the level of image structure. Since its supervision is fairly generic, it can be integrated into existing training pipelines with minimal overhead. The observed performance gains across multiple scribble-supervised methods suggest that explicitly encoding spatial coherence during pseudo-label refinement is a powerful direction related to traditional edge detection and image segmentation. Beyond scribble-based supervision, the proposed strategy is applicable to other forms of sparse or noisy annotation, making it a flexible tool for a broad range of weakly supervised segmentation settings, including those outside the medical domain.



**Fig. 3.** Qualitative comparison among enhancement results with different PLESS setups on a sample from ACDC.

## 6. REFERENCES

- [1] Deepak Pathak, Evan Shelhamer, Jonathan Long, and Trevor Darrell, “Fully convolutional multi-class multiple instance learning,” *arXiv preprint*, 2014.
- [2] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille, “Weakly- and semi-supervised learning of a deep convolutional network for semantic image segmentation,” in *ICCV*, 2015, pp. 1742–1750.
- [3] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei, “What’s the point: Semantic segmentation with point supervision,” in *ECCV*, 2016, pp. 549–565.
- [4] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun, “ScribbleSup: Scribble-supervised convolutional networks for semantic segmentation,” in *CVPR*, 2016, pp. 3159–3167.
- [5] Gabriele Valvano, Andrea Leo, and Sotirios A Tsaftaris, “Learning to segment from scribbles using multi-scale adversarial attention gates,” *IEEE TMI*, vol. 40, no. 8, pp. 1990–2001, 2021.
- [6] Lisa Margret Koch, Martin Rajchl, Wenjia Bai, Christian Frederik Baumgartner, Tong Tong, Jonathan Passerat-Palmbach, Paul Aljabar, and Daniel Rueckert, “Multi-atlas segmentation using partially annotated data: methods and annotation strategies,” *IEEE TPAMI*, vol. 40, no. 7, pp. 1683–1696, 2017.
- [7] Meng Tang, Federico Perazzi, Abdelaziz Djelouah, Ismail Ben Ayed, Christopher Schroers, and Yuri Boykov, “On regularized losses for weakly-supervised CNN segmentation,” in *ECCV*, 2018, pp. 507–522.
- [8] Yigit B Can, Krishna Chaitanya, Basil Mustafa, Lisa M Koch, Ender Konukoglu, and Christian F Baumgartner, “Learning to segment medical images with scribble-supervision alone,” in *DLMIA/ML-CDS*, 2018, pp. 236–244.
- [9] Xiaoming Liu, Quan Yuan, Yaozong Gao, Kelei He, Shuo Wang, Xiao Tang, Jinshan Tang, and Dinggang Shen, “Weakly supervised segmentation of COVID19 infection with scribble annotation on CT images,” *Pattern Recognition*, vol. 122, pp. 108341, 2022.
- [10] Ke Zhang and Xiahai Zhuang, “CycleMix: A holistic strategy for medical image segmentation from scribble supervision,” in *CVPR*, 2022, pp. 11656–11665.
- [11] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang, “Semi-supervised semantic segmentation with cross pseudo supervision,” in *CVPR*, 2021, pp. 2613–2622.
- [12] Wenfeng Luo and Meng Yang, “Semi-supervised semantic segmentation via strong-weak dual-branch network,” in *ECCV*, 2020, pp. 784–800.
- [13] Yicheng Wu, Minfeng Xu, Zongyuan Ge, Jianfei Cai, and Lei Zhang, “Semi-supervised left atrium segmentation with mutual consistency training,” in *MICCAI*, 2021, pp. 297–306.
- [14] Xiangde Luo, Minhao Hu, Wenjun Liao, Shuwei Zhai, Tao Song, Guotai Wang, and Shaoting Zhang, “Scribble-supervised medical image segmentation via dual-branch network and dynamically mixed pseudo labels supervision,” in *MICCAI*, 2022, pp. 528–538.
- [15] Ziyang Wang and Irina Voiculescu, “Weakly supervised medical image segmentation through dense combinations of dense pseudo-labels,” in *MICCAI*, 2023, pp. 1–10.
- [16] Zihan Li, Yuan Zheng, Xiangde Luo, Dandan Shan, and Qingqi Hong, “ScribbleVC: Scribble-supervised medical image segmentation with vision-class embedding,” in *ACM Multimedia*, 2023, pp. 3384–3393.
- [17] Tao Wang, Xinlin Zhang, Yuanbin Chen, Yuanbo Zhou, Longxuan Zhao, Tao Tan, and Tong Tong, “ScribbleVS: Scribble-supervised medical image segmentation via dynamic competitive pseudo label selection,” *arXiv preprint*, 2024.
- [18] Varduhi Yeghiazaryan, Yeva Gabrielyan, and Irina Voiculescu, “Parallel watershed partitioning: GPU-based hierarchical image segmentation,” *J. Parallel Distrib. Comput.*, vol. 205, pp. 105140, 2025.
- [19] Yeva Gabrielyan, Varduhi Yeghiazaryan, and Irina Voiculescu, “Semantic segmentation with spreading scribbles,” in *MIUA*, 2025, pp. 73–87.
- [20] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo, “CutMix: Regularization strategy to train strong classifiers with localizable features,” in *ICCV*, 2019, pp. 6023–6032.
- [21] Jang-Hyun Kim, Wonho Choo, and Hyun Oh Song, “Puzzle Mix: Exploiting saliency and local statistics for optimal mixup,” in *ICML*, 2020, pp. 5275–5285.
- [22] Jang-Hyun Kim, Wonho Choo, Hosan Jeong, and Hyun Oh Song, “Co-Mixup: Saliency guided joint mixup with supermodular diversity,” *arXiv preprint*, 2021.
- [23] Varduhi Yeghiazaryan and Irina Voiculescu, “Family of boundary overlap metrics for the evaluation of medical image segmentation,” *J. Med. Imaging*, vol. 5, no. 1, pp. 015006, 2018.