

Scribble2D5: Weakly-Supervised Volumetric Image Segmentation via Scribble Annotations

Qiuhui Chen¹ and Yi Hong^{1(✉)}

¹Department of Computer Science and Engineering, Shanghai Jiao Tong University
yi.hong@sjtu.edu.cn

Abstract. Image segmentation using weak annotations like scribbles has gained great attention, since such annotations are easier to obtain compared to time-consuming and labor-intensive labeling at the pixel/voxel level. However, scribbles lack structure information of the region of interest (ROI), thus existing scribble-based methods suffer from poor boundary localization. Moreover, current methods are mostly designed for 2D image segmentation, which do not fully leverage volumetric information. In this paper, we propose a scribble-based volumetric image segmentation, Scribble2D5, which tackles 3D anisotropic image segmentation and improves boundary predictions. To achieve this, we augment a 2.5D attention UNet with a proposed label propagation module to extend semantic information from scribbles and a combination of static and active boundary prediction to learn ROI’s boundaries and regularize its shape. Extensive experiments on three public datasets demonstrate Scribble2D5 significantly outperforms existing scribble-based methods and approaches the performance of fully-supervised ones. Our code is available at <https://github.com/Qybc/Scribble2D5>.

Keywords: Weakly-supervised Learning · Scribble Annotation · Volumetric Image Segmentation.

1 Introduction

Deep learning based methods have achieved impressive accuracy in many medical segmentation tasks, especially in a fully-supervised manner [19,30]. However, such segmentation methods typically require a large amount of dense annotations for pixels or voxels to train a deep model. While dense annotations sometimes are not easy to obtain in practice because annotating at the image pixel-/voxel-level is time-consuming and needs medical expertise to provide high-quality labels. On the other hand, fully-unsupervised segmentation methods [24,6] have shown promising results; however, their performance gap with respect to fully-supervised approaches is too large to make them practical. Therefore, weakly-supervised approaches by using weak annotations have gained great attention to greatly reduce the workload of manual annotations while having promising and comparable results compared to fully-supervised approaches.

Commonly-used weak annotations include image-level annotations [26,2], bounding boxes [15,19], scribbles [11,20,7,21], and extreme points [13,17], etc.

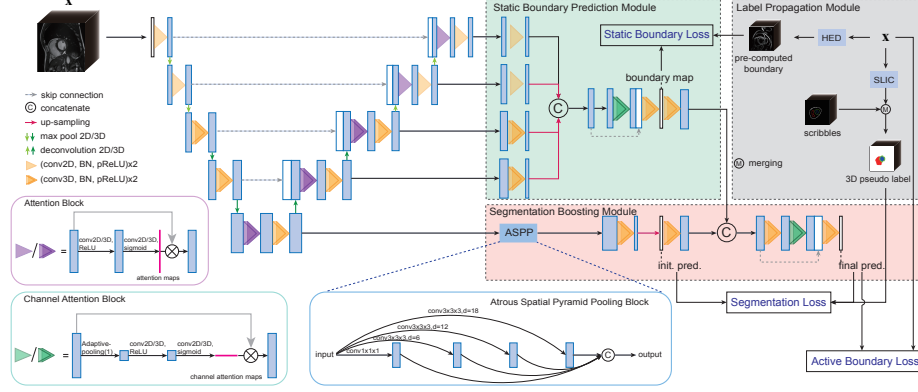


Fig. 1. Our Scribble2D5 architecture for volumetric image segmentation: 1) pseudo label propagation (gray box): generating pseudo 3D segmentation masks and pre-computed boundaries; 2) static boundary prediction (green box): incorporating object boundary information from the input image; 3) segmentation boosting (pink box): further considering active boundaries via an active boundary loss. (Best view in color)

Compared to image-level and bounding box annotations, scribbles provide rough positions of Region of Interests (ROIs) to allow a better location of ROI. Also, scribbles are more flexible than bounding boxes and extreme points when annotating, especially for ROIs with irregular shapes. In addition, extreme point annotations are more suitable for convex shapes and may not work well for non-convex ones. Therefore, we choose scribbles as our weak annotations. However, scribbles are often sparse with no structure information of ROIs; as a result, scribble-based methods have difficulty in accurately locating ROI boundaries [20]. Moreover, existing methods [11,23,28,21,12,29] are typically designed for segmenting 2D image slices, which do not fully leverage the whole image volume, with missing continuity between slices. Researchers attempt to alleviate such issue by regularizing the volume size of segmentation outputs [10]. Another solution [7] performs 3D segmentation with transfer learning, which learns with dense annotations in the source domain and with scribbles in the target domain.

Inspired by two recent works [23,28], we propose a volumetric segmentation network based on scribble annotations, called scribble2D5. As shown in Fig. 1, we adopt a 2.5D attention UNet [19] to handle anisotropic medical volumes with different voxel spacings, which is very common in practice like our datasets. To amplify the influence of sparse scribbles and suit for volumetric segmentation, we use a label propagation module based on supervoxels to generate 3D pseudo masks from scribbles for supervision. To address the boundary localization issue, we propose to learn both static and active boundaries via predicting edges in 3D and optimizing an active boundary loss in 3D based on active contour model [5]. Different from existing solutions, our scribble2D5 tackles 3D anisotropic image inputs directly, and needs scribbles only for training. At the inference stage, segmentation operates automatically, with no need of scribble inputs.

We evaluate our methods on three datasets, including ACDC dataset [4] for cardiac segmentation, VS dataset [18] for tumor segmentation, and CHAOS dataset [9] for abdominal organ segmentation. For both ACDC and CHAOS datasets, our method outperforms the current state-of-the-art (SOTA) by large margins on three different evaluation metrics; and on the VS dataset, our method achieves better performance in Dice compared to SOTA. Our method reduces the performance gap between weakly-supervised and fully-supervised approaches, which makes it more practical to be used in the future.

Overall, our contributions in the paper are summarized as follows:

- We propose a scribble2D5 network for segmenting medical image volumes with scribbles for training only. Our method is compared to five baselines and significantly outperforms scribble-based methods on three datasets.
- We propose a label propagation module for 3D pseudo mask generation and an active boundary loss to regularize 3D segmentation results. These modules are general and could be used in other segmentation networks.

2 Scribble2D5: Scribbles-based Volumetric Segmentation

Figure 1 presents the framework of our Scribble2D5, a weakly-supervised image segmentation network based on scribble annotations. Scribble2D5 has a 2.5D attention UNet [19] as the backbone network, which is augmented by three modules, i.e., a label propagation module for generating 3D pseudo masks and boundaries, a static boundary prediction module for incorporating object boundary information from images, and a segmentation boosting module for further considering active boundaries via an active boundary loss.

2.1 3D Pseudo Label Generation via Label Propagation

Scribble annotations are often sparse, which cover a small amount of pixels on each slice of an image volume. As a result, the supervision information from scribbles is not strong enough to produce good guidance, like UNet_{PCE} [20]. To address this issue, we propose to magnify the effects of scribble annotations in 3D by leveraging supervoxels. That is, we adopt SLIC [1], which generates supervoxels from images using an adaptive k-means clustering by considering both image intensity and distance similarities. We collect the supervoxels that scribbles pass through, resulting in 3D pseudo segmentation masks for ROIs.

Except for the pseudo mask we generate from the scribble annotations, we generate the pseudo static boundary of ROI from an image volume by stacking 2D edges detected on each slice. This boundary is static since it is pre-computed from the image and keeps unchanged during training, which is different from the active boundary we will discuss later. To obtain 2D edges, we directly use an existing method, HED [25], which is pretrained on the generic edges of BSDS500 [3]. In this way, we have a Label Propagation Module (LPM) to generate 3D pseudo labels from scribbles and images for ROI segmentation and pre-computed boundary for static boundary prediction, respectively.

2.2 Scribble2D5 Network

Backbone. The image volumes studied in the experiments have different voxel spacings. Roughly, the in-plane resolution within a slice is about four times the thickness of a slice. Since 2D CNNs ignore the important correlations among slices and 3D CNNs typically handle isotropic image volumes, we choose a 2.5D neural network that considers the anisotropic properties of an image volume. In particular, we adopt an attention UNet2D5 [19] as our backbone network, which augments UNet2D5 by adding a simple attention block at each deconvolutional layer, as shown in Fig. 1. Specifically, at the top two layers of both encoder and decoder branches, we have 2D convolutional operations; while at other layers, the feature maps are isotropic, which are suitable for 3D convolutions. The attention blocks are colored in purple in Fig. 1. Their attention maps are estimated via two layers of convolutions, i.e., one with ReLU and the other with a Sigmoid activation function. This 2.5D network suits for all images in our experiments.

Static Boundary Prediction Module (SBPM). This module encourages the backbone network to extract image features with rich boundary structures at different scales. Following [28], we collect feature maps from different layers of the network decoder, and concatenate them at different resolutions right after one convolutional layer with a filter of size $1 \times 1 \times 1$. To fuse these features, we feed them to a residual channel attention block (as shown by a green triangle in Fig. 1) and a $1 \times 1 \times 1$ convolutional layer to produce a boundary map b in 3D. Under the supervision of the previously generated 3D pseudo boundary B , the network is trained with a cross-entropy loss: $\mathcal{L}_{bry}(b, B) = -\sum_{c=1}^N B_c \log(b_c)$. Here, N is the total number of classes for segmenting in a dataset.

Segmentation Boosting Module (SBM). This module performs segmentation under the supervision of the pseudo mask generated with supervoxels and a regularization on segmentation output. The module includes an initial segmentation and a final one with further considering static and active boundaries. To predict a preliminary mask, we employ a dense atrous spatial pyramid pooling (DenseASPP) block [27] right after the bottom layer of the backbone network in Fig. 1, which enlarges its receptive fields by utilizing different dilation rates. In this block, the convolutional layers are connected in a dense way to cover a larger scale range without significantly increase the model size.

To generate the initial segmentation mask M^{init} , we adopt two additional 3D convolutional layers followed by a $1 \times 1 \times 1$ convolution, resulting in the initial prediction supervised by the generated pseudo mask M^{pseudo} . Considering the oversegment nature of supervoxels, one supervoxel may be selected by multiple different classes. To avoid this confusion, we only consider those supervoxels with a unique label, which are set as 1 in the mask M^{voxel} with others being zeros. We use the partial cross entropy to supervise the initial segmentation: $\mathcal{L}_{seg}(M^{init}, M^{pseudo}, M^{voxel}) = -\sum_{c=1}^N M_c^{voxel} \cdot M_c^{pseudo} \log(M_c^{init})$. This loss function allows early feedback to fasten the network convergence.

To refine the initial estimation and obtain a boundary-preserving mask for a final prediction, we merge SBPM outputs with those from the initial mask prediction for a refinement. These feature maps are fed to a residual channel

Table 1. Quantitative comparison among baselines and our method for volumetric segmentation on three datasets. Mean and standard deviation (subscript) are reported. The upper bounds are colored in blue, and the best results by using scribbles are marked in bold. [†]P is short for Point, indicating extreme points. such annotations are available only in the VS dataset. *These numbers are taken from [8]. (Best viewed in color)

Approach		ACDC			VS			CHAOS		
		Dice (%, [†])	HD95 (mm, [↓])	Precision (%, [†])	Dice (%, [†])	HD95 (mm, [↓])	Precision (%, [†])	Dice (%, [†])	HD95 (mm, [↓])	Precision (%, [†])
Supervision Type	Scribble									
	UNet _{PCE} [20]	79.0 ₀₆	6.9 ₀₄	77.3 ₀₆	44.6 ₀₈	6.5 ₀₃	43.8 ₀₅	34.4 ₀₆	9.4 ₀₃	36.6 ₀₅
	MAAG [21]	83.4 ₀₄	8.6 ₀₄	78.5 ₀₅	69.4 ₀₆	5.9 ₀₅	56.8 ₀₅	66.4 ₀₅	3.8 ₀₅	57.2 ₀₆
	Ours w/o LPM	83.2 ₀₅	7.7 ₀₃	84.1 ₀₅	78.8 ₀₅	4.6 ₀₁	77.6 ₀₅	81.2 ₀₇	5.8 ₀₈	82.0 ₀₆
	Ours w/o SBPM	85.6 ₀₅	4.6 ₀₄	85.5 ₀₄	80.6 ₀₅	7.1 ₀₃	81.6₀₄	84.6 ₀₅	5.5 ₀₅	83.1 ₀₅
Mask	Ours w/o ABL	88.7 ₀₄	5.1 ₀₈	86.0₀₅	81.0 ₀₃	4.8 ₀₁	80.1 ₀₅	85.6 ₀₄	4.8 ₀₅	81.3 ₀₂
	Scribble2D5(ours)	90.6₀₃	2.3₀₅	84.7 ₀₅	82.6₀₇	4.7 ₀₄	81.5 ₀₆	86.0₀₄	2.9₀₂	88.2₀₃
	InExtremIS [8]	-	-	-	81.9 ₀₃ [*]	3.7₀₃[*]	92.9₀₂[*]	-	-	-
Mask	2D UNet [16]	93.0 ₀₅	3.5 ₁₅	90.2 ₀₇	80.4 ₀₃	7.3 ₀₄	81.2 ₀₃	82.3 ₀₄	3.3 ₀₁	81.7 ₀₅
	2.5D UNet [19]	96.1₀₃	0.3₀₀	95.3₀₄	87.3₀₂	6.8 ₀₄	84.7 ₀₃	90.8₀₃	1.1₀₀	91.4₀₅

attention block, followed by a $1 \times 1 \times 1$ convolutional layer to predict the final mask M^{final} . Similarly, we use the partial cross-entropy loss to predict the final mask under the supervision of the generated pseudo mask M^{pseudo} .

Active Boundary (AB) Loss. The pseudo masks are imperfect because supervoxels are coarse segmentation masks of ROIs and have oversegment issues, resulting in a potential of having many false positives. To mitigate this issue, we propose regularizing the surface and volume of the 3D segmentation region by upgrading the active contour loss [5] to a 3D version. We apply an AB loss as following: $\mathcal{L}_{AB} = Surface + \lambda_1 \cdot Volume_{In} + \lambda_2 \cdot Volume_{Out}$, where $Surface = \int_S |\nabla u| ds$ and u is the prediction; $Volume_{In} = \int_V (c_1 - v)^2 u dx$, c_1 is the mean image intensity inside of interested regions V , and v is the input image; $Volume_{Out} = \int_{\bar{V}} (c_2 - v)^2 u dx$ and c_2 is the mean image intensity outside of the region. These items are balanced by two hyper-parameters λ_1 and λ_2 . In the experiments, we set $\lambda_1 = 1$ and $\lambda_2 = 0.1$, to emphasis more on the inside region of the volume. This new loss function considers the shape and intensity of an image in 3D, which regularizes ROI's shapes and reduces false positives.

The final loss function is $\mathcal{L}_{total} = \beta_1 \mathcal{L}_{bry}(b, B) + \mathcal{L}_{seg}(M^{init}, M^{pseudo}, M^{voxel}) + \mathcal{L}_{seg}(M^{final}, M^{pseudo}, M^{voxel}) + \beta_2 \mathcal{L}_{AB}$. Here, β_1 and β_2 are weights for balancing loss terms, which are both set as 0.3.

3 Experiments

3.1 Datasets and Experimental Settings

ACDC Dataset [4]. This dataset consists of Cine MR images collected from 100 patients. Manual segmentation masks of the left and right ventricles and




Fig. 2. Qualitative comparison among scribble-based (UNet_{PCE} and MAAG), mask-based (2D and 2.5D UNets), and our methods. (Best viewed in color)

myocardium are provided at the end-diastolic and end-systolic cardiac phases. The slice size is 256×208 with the pixel spacing varying from 1.37 to $1.68mm$. The number of slices is between 28 and 40, and the slice thickness is $5mm$ or $8mm$. We subject-wisely divide the ACDC dataset into sets of 70%, 15% and 15% for training, validation, and test, respectively.

VS Dataset [18]. This dataset collects T2-weighted MRIs from 242 patients with a single sporadic vestibular schwannoma (VS) tumor. The size of an image slice is 384×384 or 448×448 , with a pixel spacing of $0.5 \times 0.5mm^2$. The number of slices varies from 19 to 118, with a thickness of $1.5mm$. The VS tumor masks are manually annotated by neurosurgeons and physicists. The dataset is subject-wisely split into 172 for training, 20 for validation, and 46 for test.

CHAOS Dataset [9]. This dataset has abdominal T1-weighted MR images collected from 20 subjects and the corresponding segmentation masks for liver, kidneys, and spleen. The image slice size is 256×256 with a resolution of $1.36 - 1.89mm$ (average $1.61mm$). The number of slices is between 26 and 50 (average 36) with the slice thickness varying from 5.5 to $9mm$ (average $7.84mm$). We also subject-wisely divide this dataset into sets of 70%, 15% and 15% for training, validation, and test, respectively.

Scribble Generation and Other Settings. For the ACDC dataset, we use the scribbles provided in [21], which are manually drawn by experts at both end-diastolic and end-systolic phases. For both VS and CHAOS datasets, following [14], we simulate scribbles by an iterative morphological erosion and closing of segmentation masks, which results in a one-pixel skeleton for each object. Since the resulting background scribble is winding, we use ITK-Snap to annotate background with 1-pixel width curves.

For all datasets, we randomly crop an image volume and obtain patches of size $224 \times 224 \times 32$ as the network inputs. An image volume is padded with zeros if its size is smaller than the input size.

We train all models for 200 epochs with early stopping. The weights of the network are initialized by following a normal distribution with a mean of 0 and variance of 0.01. We use Adam optimizer with a weight decay 10^{-7} and an initial learning rate $1e-4$. The whole training takes about 6 hours with a batch size of 4 on one NVIDIA GeForce RTX 3090 GPU.

Baselines and Evaluation Metrics. To demonstrate the effectiveness of our methods, we select three groups of baselines, including fully-supervised methods (i.e., 2D UNet [16] and 2.5D UNet [19]), weakly-supervised methods using scribbles (i.e., UNet_{PCE} [20] and MAAG [21]) and a weakly-supervised method using extreme points [8]. To evaluate the segmentation performance, we use the Dice score to calculate the overlap between our segmentation and the ground truth (GT), the 95th percentile of the Hausdorff Distance (HD95) to measure the distance between our ROI boundary and GT, and the precision to check the purity of the positively-segmented voxels.

3.2 Experimental Results

Table 1 presents our experimental results on three datasets with comparison to five baselines. For all datasets, the upper bounds of the segmentation performance are mainly provided by the 2.5D UNet, which are colored in blue in Table 1. Compared to the scribble-based SOTA method on ACDC and CHAOS datasets, i.e., MAAG [21], scribble2D5 improves the Dice score by 7% and 19.5%, reduces the HD95 by 6.3mm and 1.8mm, and improves the precision by 6.2% and 29.5%, respectively. In addition, our method outperforms the most recent two methods, i.e., ScribbleSeg [12] and CycleMix [29], on ACDC Dataset, with our 0.903 mean dice vs 0.872 in [12] (using the same 5-fold cross validation), our 0.896 mean dice vs 0.848 in [29] (similarly, using 35 subjects for training). Compared to the extreme-point-based SOTA method on the VS dataset, i.e., InExtremeIS [8], although our method has a lower precision and HD95 value, it improves the Dice score by 0.7%. We do not report InExtremeIS’ results on ACDC and CHAOS datasets because extreme points for these two datasets are not available or easy to be generated.

Figure 2 demonstrates the sampled qualitative results of our method compared to the baselines. Overall, we have fewer false positives compared to scribble-based methods, i.e., UNet_{PCE} and MAAG, and better boundary localization with more accurate boundary prediction for each ROI. Regarding the comparison with mask-based methods, our method sometimes generates even better masks than 2D UNet, while needs improvements at details compared to 2.5D UNet.

Ablation Study. To check the effectiveness of each module in our method, we perform an ablation study with three variants: a) **Ours w/o LPM**: Scribble2D5 without the label propagation module (LPM); b) **Ours w/o SBPM**: Scribble2D5 without the static boundary prediction model (SBPM), including SBPM and active boundary loss; and c) **Ours w/o ABL**: Scribble2D5 without the active boundary loss (ABL). Take the ACDC dataset as an example, as shown in Table 1, without LPM but all others, the Dice score reduces from 90.4% to 80.6%. With LPM but without SBPM, the Dice score is 85.6%; then




Fig. 3. Visualization of our intermediate and final results on image samples from ACDC dataset. The ground truth (GT) is colored in blue, like the blue region in the first column and the blue contours in other images, while our predictions are colored in red. Yellow arrows show the effect of the active boundary loss (ABL). (Best viewed in color)

the static boundary prediction module contributes an improvement of 3.1%, and the active boundary loss contributes an additional improvement of 1.7% in Dice score. Figure 3 visualizes two samples from the ACDC dataset with our intermediate and final prediction results. Without LPM, our method suffers from false positives far away from the ROI; without SBPM, our method has oversegment issues of the ROI; by adding the boundary map and active boundary regularization, our method adjusts the prediction based on the image edge and texture information, resulting the closest results compared to GT.

4 Conclusion and Discussion

In this paper, we proposed a weakly-supervised volumetric image segmentation network, Scribble2D5, which significantly outperforms existing scribble-based methods and reduces the performance gap between weakly-supervised and full-supervised segmentation methods. One limitation of our method is that our pseudo boundary labels are not purely 3D, which will be explored in the future. Currently, we do not consider the case of missing scribble annotations on some slices or further reducing the manual work via an adaptive annotation. One potential solution for this is using the watershed technique [22] in 3D. Also, We observe that the shape and location of scribbles would affect the segmentation accuracy, summarizing a way to make scribble annotations for different shapes of ROIs will be useful in practice by providing some rules for users to make annotations, which will benefit the segmentation and be left as the future work.

Acknowledgment. This work was supported by Shanghai Municipal Science and Technology Major Project 2021SHZDZX0102.

References

1. Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Süsstrunk, S.: Slic superpixels compared to state-of-the-art superpixel methods. *IEEE transactions on pattern analysis and machine intelligence* **34**(11), 2274–2282 (2012)
2. Ahn, J., Kwak, S.: 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*. pp. 4981–4990 (2018)
3. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence* **33**(5), 898–916 (2010)
4. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Ballester, M.A.G., et al.: Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE transactions on medical imaging* **37**(11), 2514–2525 (2018)
5. Chen, X., Williams, B.M., Vallabhaneni, S.R., Czanner, G., Williams, R., Zheng, Y.: Learning active contour models for medical image segmentation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11632–11640 (2019)
6. Dey, R., Hong, Y.: Asc-net: Adversarial-based selective network for unsupervised anomaly segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 236–247. Springer (2021)
7. Dorent, R., Joutard, S., Shapey, J., Bisdas, S., Kitchen, N., Bradford, R., Saeed, S., Modat, M., Ourselin, S., Vercauteren, T.: Scribble-based domain adaptation via co-segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 479–489. Springer (2020)
8. Dorent, R., Joutard, S., Shapey, J., Kujawa, A., Modat, M., Ourselin, S., Vercauteren, T.: Inter extreme points geodesics for end-to-end weakly supervised image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 615–624. Springer (2021)
9. Kavur, A.E., Gezer, N.S., Barış, M., Aslan, S., Conze, P.H., Groza, V., Pham, D.D., Chatterjee, S., Ernst, P., Özkan, S., et al.: Chaos challenge-combined (ctmr) healthy abdominal organ segmentation. *Medical Image Analysis* **69**, 101950 (2021)
10. Kervadec, H., Dolz, J., Tang, M., Granger, E., Boykov, Y., Ayed, I.B.: Constrained-cnn losses for weakly supervised segmentation. *Medical image analysis* **54**, 88–99 (2019)
11. Lin, D., Dai, J., Jia, J., He, K., Sun, J.: Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3159–3167 (2016)
12. Luo, X., Hu, M., Liao, W., Zhai, S., Song, T., Wang, G., Zhang, S.: Scribble-supervised medical image segmentation via dual-branch network and dynamically mixed pseudo labels supervision. *arXiv preprint arXiv:2203.02106* (2022)
13. Maninis, K.K., Caelles, S., Pont-Tuset, J., Van Gool, L.: Deep extreme cut: From extreme points to object segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 616–625 (2018)
14. Rajchl, M., Koch, L.M., Ledig, C., Passerat-Palmbach, J., Misawa, K., Mori, K., Rueckert, D.: Employing weak annotations for medical image analysis problems. *arXiv preprint arXiv:1708.06297* (2017)

15. Rajchl, M., Lee, M.C., Oktay, O., Kamnitsas, K., Passerat-Palmbach, J., Bai, W., Damodaram, M., Rutherford, M.A., Hajnal, J.V., Kainz, B., et al.: Deepcut: Object segmentation from bounding box annotations using convolutional neural networks. *IEEE transactions on medical imaging* **36**(2), 674–683 (2016)
16. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. pp. 234–241. Springer (2015)
17. Roth, H.R., Yang, D., Xu, Z., Wang, X., Xu, D.: Going to extremes: weakly supervised medical image segmentation. *Machine Learning and Knowledge Extraction* **3**(2), 507–524 (2021)
18. Shapey, J., Kujawa, A., Dorent, R., Wang, G., Bisdas, S., Dimitriadis, A., Grishchuck, D., Paddick, I., Kitchen, N., Bradford, R., et al.: Segmentation of vestibular schwannoma from magnetic resonance imaging: An open annotated dataset and baseline algorithm. *The Cancer Imaging Archive* (2021)
19. Shapey, J., Wang, G., Dorent, R., Dimitriadis, A., Li, W., Paddick, I., Kitchen, N., Bisdas, S., Saeed, S.R., Ourselin, S., et al.: An artificial intelligence framework for automatic segmentation and volumetry of vestibular schwannomas from contrast-enhanced t1-weighted and high-resolution t2-weighted mri. *Journal of neurosurgery* **134**(1), 171–179 (2019)
20. Tang, M., Djelouah, A., Perazzi, F., Boykov, Y., Schroers, C.: Normalized cut loss for weakly-supervised cnn segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1818–1827 (2018)
21. Valvano, G., Leo, A., Tsaftaris, S.A.: Learning to segment from scribbles using multi-scale adversarial attention gates. *IEEE Transactions on Medical Imaging* **40**(8), 1990–2001 (2021)
22. Vincent, L., Soille, P.: Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis & Machine Intelligence* **13**(06), 583–598 (1991)
23. Wang, B., Qi, G., Tang, S., Zhang, T., Wei, Y., Li, L., Zhang, Y.: Boundary perception guidance: A scribble-supervised semantic segmentation approach. In: *IJCAI International joint conference on artificial intelligence* (2019)
24. Xia, X., Kulis, B.: W-net: A deep model for fully unsupervised image segmentation. *arXiv preprint arXiv:1711.08506* (2017)
25. Xie, S., Tu, Z.: Holistically-nested edge detection. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1395–1403 (2015)
26. Xu, J., Schwing, A.G., Urtasun, R.: Learning to segment under various forms of weak supervision. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3781–3790 (2015)
27. Yang, M., Yu, K., Zhang, C., Li, Z., Yang, K.: Denseaspp for semantic segmentation in street scenes. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3684–3692 (2018)
28. Zhang, J., Yu, X., Li, A., Song, P., Liu, B., Dai, Y.: Weakly-supervised salient object detection via scribble annotations. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 12546–12555 (2020)
29. Zhang, K., Zhuang, X.: Cyclemix: A holistic strategy for medical image segmentation from scribble supervision. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11656–11665 (2022)
30. Zhou, Z., Rahman Siddiquee, M.M., Tajbakhsh, N., Liang, J.: Unet++: A nested u-net architecture for medical image segmentation. In: *Deep learning in medical image analysis and multimodal learning for clinical decision support*, pp. 3–11. Springer (2018)