



Weakly Supervised Volumetric Image Segmentation with Deformed Templates

Wickramasinghe, Udaranga; Jensen, Patrick; Shah, Mian; Yang, Jiancheng; Fua, Pascal

Published in:

Proceedings of the 25th International Conference of Medical Image Computing and Computer Assisted Intervention, MICCAI 2022

Link to article, DOI:

[10.1007/978-3-031-16443-9_41](https://doi.org/10.1007/978-3-031-16443-9_41)

Publication date:

2022

Document Version

Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):

Wickramasinghe, U., Jensen, P., Shah, M., Yang, J., & Fua, P. (2022). Weakly Supervised Volumetric Image Segmentation with Deformed Templates. In *Proceedings of the 25th International Conference of Medical Image Computing and Computer Assisted Intervention, MICCAI 2022* (Vol. 13435, pp. 422-432). Springer. https://doi.org/10.1007/978-3-031-16443-9_41

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Weakly Supervised Volumetric Image Segmentation with Deformed Templates

Udaranga Wickramasinghe¹, Patrick Jensen^{1,2}, Mian Shah¹, Jiancheng Yang^{1,3}, Pascal Fua¹

¹EPFL, Switzerland ²DTU, Denmark ³Shanghai Jiao Tong University, China

Abstract. There are many approaches to weakly-supervised training of networks to segment 2D images. By contrast, existing approaches to segmenting volumetric images rely on full-supervision of a subset of 2D slices of the 3D volume. We propose an approach to volume segmentation that is truly weakly-supervised in the sense that we only need to provide a sparse set of 3D points on the surface of target objects instead of detailed 2D masks. We use the 3D points to deform a 3D template so that it roughly matches the target object outlines and we introduce an architecture that exploits the supervision it provides to train a network to find accurate boundaries. We evaluate our approach on Computed Tomography (CT), Magnetic Resonance Imagery (MRI) and Electron Microscopy (EM) image datasets and show that it substantially reduces the required amount of effort.

1 Introduction

State-of-the-Art volumetric segmentation techniques rely on Convolutional Neural Networks (CNNs) operating on image volumes [3, 23]. However, their performance depends critically on obtaining enough annotated data, which itself requires expert knowledge and is both tedious and expensive.

Weakly-supervised image segmentation techniques can be used to mitigate this problem. They typically rely on tag annotations [10, 8] or coarse object annotations in the form of point annotations [20, 33], bounding box annotations [9], scribbles [33] or approximate target shapes [14]. However, these techniques have been mostly demonstrated in 2D and do not provide enough information when segmenting complex shapes, such as the liver or the hippocampus. For 3D volume segmentation, the dominant approach is to fully label a subset of 2D slices [3, 5]. This is often referred to as *weak supervision*, even though it requires full supervision within individual slices.

By contrast, we propose a truly weakly-supervised approach that only requires a sparse set of 3D points on the surface of the target objects instead of the usual 2D masks in selected slices. Given an appropriate user-interface, this is much faster and easier because it eliminates the need to painstakingly outline fine details, as shown in Fig. 1. To this end, we introduce the **Weak-Net** architecture depicted by Fig. 2. It comprises two U-Net-like networks. The first one produces a segmentation map that matches a rough model of the target object obtained




Fig. 1. Iterative annotation strategy. The annotator provides a few 3D points, which we fit to the template using an Active Surface Model (ASM) [28]. The result is rasterized and overlaid on the images. The annotator can then add more and deform the template again as needed.

from the 3D point annotations. The second one takes the map as input and uses it to reconstruct the original image, which forces the segmentation boundaries to be accurate even though those of the template are not.

We evaluate the performance of **Weak-Net** on Computed Tomography (CT), Magnetic Resonance Imagery (MRI) and Electron Microscopy (EM) datasets. We show that it outperforms the standard approach to weak-supervision in 3D at a reduced supervision cost. More specifically, we can deliver the same accuracy as when fully annotating 2D slices for less than a third of the annotation effort. This matters because annotators typically are experts whose time is both scarce and valuable. Furthermore, it creates the basis for interactive annotating strategies that deliver the full accuracy at a lower cost than full supervision.

2 Related Work

We review current approaches to weak-supervision for 2D and 3D image segmentation. We then discuss using atlases to segment biomedical image volumes.

2.1 Weakly-Supervised Image Segmentation

Segmenting 2D Images. Tag and box annotations are among the weakest forms of annotations used to segment natural and medical images [6, 10, 9]. However, they rarely provide enough supervision for accurate results. By contrast, point annotations [20], scribbles [33], and approximate shape annotation [14] can be used to provide useful shape information.

The annotation process can be sped-up using dynamic programming [19] or deformable contours [12]. This makes it possible to mark only a subset of points along a contour and have the system refine it to match the target object boundary. Unfortunately, these algorithms are hard to deploy effectively in medical imagery because there are many contours besides those of interest and they can easily confuse these algorithms. This is addressed in [21] by introducing *deep deformable contours* that only need a simple approximate contour for initialization purposes. In [2, 15], the annotator is brought into the loop by giving corrective

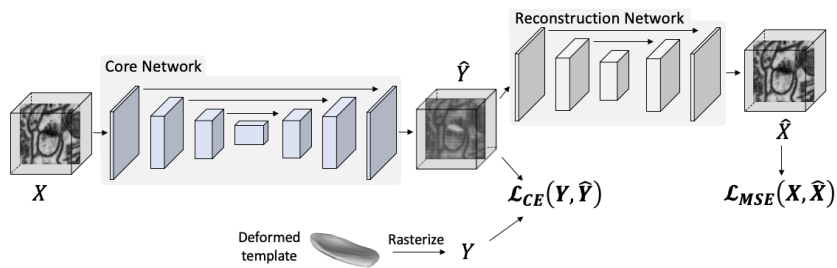


Fig. 2. Weak-Net architecture. A first U-Net takes the image X as input and outputs a segmentation \hat{Y} , which is in turn fed to a second U-Net that outputs a reconstructed image \hat{X} . Training is achieved by jointly minimizing L_{mse} and L_{ce} . This encourages \hat{X} to resemble the original image and \hat{Y} to be similar to Y , a roughly aligned version of a template.

clicks when necessary. However, these deep contours require fully labelled data for their own training.

In the result section, we show that supplying a few 3D points, as we do, is faster than using these techniques to delineate whole 2D contours in slices, as is often done.

Segmenting 3D Volumes. By far the most prevalent approach to weak-supervision for 3D segmentation, is to *fully* annotate subsets of slices from the training volumes [3]. Even though this reduces the segmentation effort, the annotator still has to carefully trace the object boundary in those slices. The effort can be reduced by using scribbles in individual slices [5] or any of the semi-automated techniques described above. Unfortunately, the effectiveness of those that do not require training is limited [19, 12] while the others [2, 15, 21] require full supervision and are not applicable in our scenario.

Part of the problem is that these annotation techniques operate purely in 2D without exploiting the 3D nature of the data. Because we do so, we only need a comparatively small number of point annotations for effective training.

2.2 Template Based Approaches

Shape priors have long been used for image segmentation [7] and much recent work use shape priors in conjunction with deep nets [29, 18]. For medical imaging, these priors are usually supplied in the form of sophisticated templates known as Probabilistic Atlases (PAs) that assign to each pixel or voxel a probability of belonging to a specific class. They are typically built by fusing multiple manually annotated images and used as auxiliary CNN inputs to provide localization priors that help the network find structures of interest. The PAs can be either very detailed, as in [25], to model structures that are known in detail or very rough, as in [27], to deal with 3D structures whose shape can vary significantly.

PAs built by annotating points have been used in medical imaging [11, 22] as a source of prior information. They are often referred to as seed layers and come in two main flavors, Gaussian priors [22] or binary seed layers [11]. These seed layers either indicate points inside the object [32, 11] or points on the boundary of the object [17, 26]. In the context of deep learning, PAs have been mostly used in fully-supervised approaches [11, 22, 25]. An exception is the work of [25]. However, they are only used for pre-training purposes. Another are the one-shot and few-shot learning-based segmentation algorithms of [4]. However, they require a few fully annotated target objects as the atlases, which we do not.

2.3 Image Reconstruction

Image reconstruction is used for semi-supervised [16] and unsupervised [31] image segmentation as an auxiliary task to improve the results. In this work, we demonstrate that this idea also applies in a weakly-supervised setting to improve the segmentations produced by the core-network trained with rough annotations. In our framework, the image reconstruction network helps refine the rough initial shapes we obtain from the annotations. There are alternative approaches to boundary refinement [1] but they are designed for 2D segmentation and extending them to 3D would be non-trivial.

3 Method

Our approach to training a network involves an annotator providing only a sparse set of 3D points on the surface of target objects for each image volume, as opposed to carefully annotating several individual slices in each. These points are used to deform a template, such as those shown in Fig. 1, using active surface models [28]. This provides a rough indication of where the target object boundary is. **Weak-Net** uses it to learn weights that yield accurate object boundaries. In short, we provide minimal human input at training time so that, at inference time, the trained network can be used without human intervention.

3.1 Network Architecture and Losses

Weak-Net is depicted by Fig. 2 and comprises two separate U-Net networks [3]. The first takes as input an image volume \mathbf{X} of $D \times H \times W$, where D , H , W stand for depth, height and width. It outputs a tensor $\hat{\mathbf{Y}}$ of dimension $D \times H \times W$ that stores the probabilities of each voxel belonging to the foreground. The second takes $\hat{\mathbf{Y}}$ as input and yields $\hat{\mathbf{X}}$, which is of the same dimension as \mathbf{X} . Ideally, $\hat{\mathbf{Y}}$ should be the desired segmentation and $\hat{\mathbf{X}}$ should be equal to \mathbf{X} .

To train **Weak-Net**, we minimize a weighted sum of two losses

$$\mathcal{L} = \mathcal{L}_{ce} + \lambda \mathcal{L}_{mse}, \quad (1)$$

$$\mathcal{L}_{ce} = - \sum_{i,j,k} \mathbf{Y}_{i,j,k} \log(\hat{\mathbf{Y}}_{i,j,k}), \quad \mathcal{L}_{mse} = \sum_{i,j,k} \mathbf{W}_{i,j,k}^{mse} (\mathbf{X}_{i,j,k} - \hat{\mathbf{X}}_{i,j,k})^2,$$




Fig. 3. Impact of the λ parameter in the loss of Eq. 1 and the thresholding parameter γ . (Col. 2-4) When $\lambda = 0$, the segmentation is very similar to the template. When λ is large, the segmentation follows boundaries that exist in the image but are not necessarily the right ones. In between, the boundaries are correct. (Col. 5-8) We set λ to 10^{-4} and vary γ .

where \mathbf{Y} is the rasterized template fitted to the target object and λ is a scalar that controls the influence of the second loss. \mathcal{L}_{ce} is a standard cross entropy loss whose minimization promotes similarity between the rasterized template \mathbf{Y} and the segmentation $\hat{\mathbf{Y}}$. \mathcal{L}_{mse} is a voxel-wise mean squared error in which the individual voxels are given weights $W_{i,j,k}^{mse}$ whose value is high within a distance d from boundaries in \mathbf{Y} and low elsewhere. At inference time, we only use the first U-Net, which we will refer to as our *core network*, and obtain the final segmentation by thresholding its output $\hat{\mathbf{Y}}$ using a threshold γ .

Minimizing \mathcal{L}_{ce} during training ensures that the segmentations will be roughly correct. However, because the template can only be expected to provide a coarse depiction of the object, this is not enough. We therefore also minimize \mathcal{L}_{mse} to force the network to yield accurate boundaries. Fig. 3 illustrates the influence of the γ and λ parameters on a real image. In practice, the results are insensitive to how they are chosen over a wide range. However, minimizing \mathcal{L}_{ce} alone ($\lambda=0$) produces boundaries in $\hat{\mathbf{Y}}$ that are exactly those of the template while minimizing \mathcal{L}_{mse} alone ($\lambda=1000$) yields boundaries that exist in the image but are not necessarily those we are looking for. Minimizing a properly weighted sum of the two yields segmentations that conform to the template while matching actual image boundaries.

3.2 Template Deformation

The template \mathbf{Y} of Eq. 1 should approximately match the target structure. Hence, the annotator should supply points that are distributed across the object surface. These points are then used to deform the template. In practice, For structures of genus 0, we start from a simple spherical template but more complex ones are possible. As we increase the number of points, we get increasingly refined templates, as shown in Fig. 4.

To perform this deformation interactively, we developed a GUI that relies on Active Surface Models (ASMs) [28] implemented as a MITK [30] plugin. It lets the annotator supply a few points by clicking on 2D cross sections of the input image volume. The ASM then deforms the template in real-time and overlays it on the image data, both as 2D cross sections and 3D surface renderings.




Fig. 4. Deformed templates. (a,b,c,d,e) Deformed template using $N = 25, 50, 125, 250,$ and 3661 user-supplied 3D points when segmenting a liver. $N = 3661$ corresponds to full annotations in all slices.

The annotator can then add more points wherever the deformed template is too far from the target organ’s boundary and iterate as often as necessary. This effectively puts the human in the loop in a painless and practical way. We illustrate this in a video that can be found in the supp. material.

4 Experiments

4.1 Datasets, Metrics and Baseline

We use three datasets acquired using MRI, FIB-SEM, and CT. MRI dataset consists of 260 labeled MRI image cubes of hippocampuses from the Medical Segmentation Decathlon [24]. CT image dataset consists of 20 labeled CT image cubes of the liver from the CHAOS challenge [13]. EM Dataset consist of 26 image cubes of Synaptic Junctions from a $500 \times 500 \times 200$ FIB-SEM image stack of a mouse cortex [29].

Our aim is to produce segmentations of these image values such that they are above a given quality threshold using as few annotations as possible. To quantify this goal, we use two metrics, one for quality and the other for annotation effort. We use the standard IoU metric [3] to quantify segmentation quality. We use the number of points provided by the annotator as a proxy for amount of effort. When the annotator provides individual points, this is clearly proportional to the time spent. When the annotator outlines contours on 2D slices, this becomes the sum of the contour lengths in each slice. It could be argued that this is an overestimate because the task is easier. However, in our experience it is not because precisely outlining a contour requires deliberation.

As discussed in Section 2.1, the generally accepted way to provide weak-supervision for 3D image segmentation is to fully annotate a few 2D slices [3]. To provide a baseline, we therefore use this approach to train a 3D U-Net, as described in [3]. For a fair comparison, we use the same one as in **Weak-Net**.

4.2 Comparative Results

We exploit the real-time performance of active surface models to enable the annotator to provide a few points, deform the template accordingly, and then add more points where the deformed shape is not satisfactory. We first benchmark

Table 1. Performance given human annotations. Values are given as mean \pm std. **Weak-Net** consistently outperforms the baseline in terms of both IoU and annotation effort required to achieve it.

		Hippocampus	Liver	Syn. Junction
Baseline	IoU (%)	71.2 \pm 0.9	81.9 \pm 0.7	68.2 \pm 0.9
	Annotation Effort (%)	12.4%	6.9%	10.8%
Weak-Net	IoU (%)	74.2 \pm 0.7	84.2 \pm 1.0	71.5 \pm 1.4
	Annotation Effort (%)	10.1%	6.8%	9.2%
Full Annotation		79.3 \pm 0.4	87.3 \pm 0.3	73.3 \pm 0.6

Table 2. Annotation time. Values are given as mean \pm std. Times are in minutes.

	Hippocampus	Liver
Manual Contouring	6.7 \pm 0.7 min.	125 \pm 19.3 min.
3D Region Growing	5.4 \pm 1.2 min.	75.4 \pm 30.5 min.
3D Fast Marching	5.7 \pm 1.4 min.	86.1 \pm 25.3 min.
Ours	4.6 \pm 0.9 min.	12.3 \pm 1.6 min.

this scenario using human annotations and then provide results using simulated annotation on the object surfaces to eliminate the subjective element it contains.

4.3 Human Annotations

Weak vs Full Supervision. We compare **Weak-Net** and the baseline against providing full supervision and report the results in Tab. 1. **Weak-Net** delivers higher IoU numbers than the baseline on all three datasets. It delivers 92 to 97% of the accuracy that can be achieved with full supervision for only 7 to 10% of the annotation effort, as defined above.

Annotation Time. In the results of Tab. 1, we use the number of points supplied by the annotators to gauge the annotation effort. To complement this, we asked the annotators to manually annotate some images using other MITK tools [30], slice-by-slice *manual contouring* of the borders, *3D Region Growing*, and *3D Fast Marching*. *3D Region Growing* and *3D Fast Marching* produce many false positive and negative regions, as shown in the supp. material. Therefore, we had to perform slice-by-slice corrections to obtain final segmentation using these two tools. In Tab. 2, we report the average time it took to fully annotate a single sample using each tool. As the Hippocampus volumes are small, point annotation is only \sim 1.5x faster than annotating the full volume. For the large Liver volumes however, point annotation is \sim 10x faster, which is significant.

Simulated Annotations To eliminate subjectivity from our experiments, we use the fact that our datasets are fully annotated to simulate the annotation

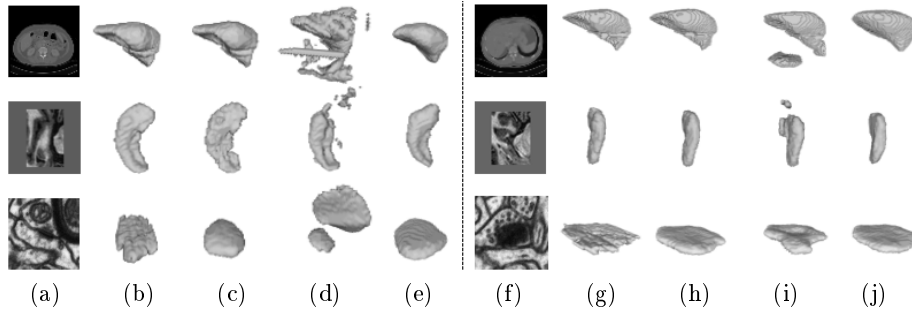


Fig. 5. Segmentation results. (a+f) A slice from the input volume (b+g) Ground truth (c+h) U-Net with full-supervision (d+i) Corresponding Baseline (e) **Weak-Net** trained with 25 points per sample. (f) **Weak-Net** trained with 125 points per sample.

process using the algorithm described in the supp. material. It is driven by two numbers, N the number of points per sample and P the number of samples we annotate. To keep the number of experiments within a manageable range, we vary both N and P when experimenting on the Hippocampus dataset and only N for the other two. To provide a baseline, we randomly pick a number of slices from three image planes to be annotated and use the ground-truth annotations for these slices. When selecting them, we check that they contain the target object. When evaluating the baseline, we vary the number of slices we use and of samples we annotate.

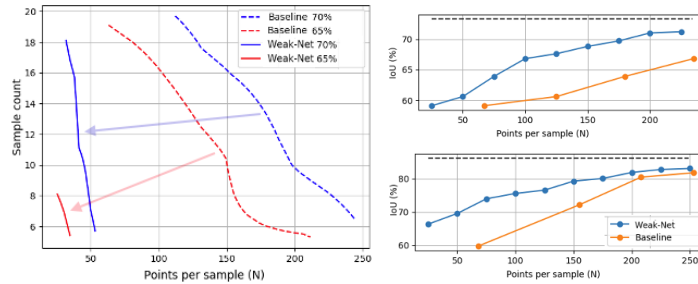


Fig. 6. Performance given synthetic annotations (left) Annotation effort required to achieve 65% and 70% IoU on the hippocampus. The blue and red arrows indicate the effort reduction our approach delivers. (right) IoU as function of the number of points used (N) for the synaptic junction and liver dataset sets. The dashed line denotes fully supervised results.

We present qualitative results of the experiments in Fig. 5. On the left of Fig. 6, we plot the total annotation effort, that is, the product of the number of points per sample and the number of samples, required to attain a target IoU—either 65% or 70%—in the Hippocampus dataset. As there are many ways

to achieve a given IoU by increasing one number while decreasing the other, we draw *iso-IoU* curves. The baseline ones are dashed and ours are full and clearly to the left of the dashed ones. In other words, we need significantly less effort to achieve a similar result. On the right of Fig. 6, we report our results on the other two datasets that contain fewer samples. Hence we used all training samples and plot the IoU as a function of the number of points per sample. For the same number of points, our approach consistently outperforms the baseline. When using it, we cannot annotate less than one slice and, hence, reduce the annotation effort below a certain level, which is why the blue curves extend further to the left than the yellow ones.

5 Conclusion

We have presented a weakly supervised approach to segmenting 3D image volumes that outperforms more traditional approaches that rely on fully annotating individual 2D slices.

It relies on deforming simple spherical templates that incorporate no shape prior. In future work, to further reduce the annotation burden, we will develop more sophisticated templates that are parameterized in terms of low-dimensional latent vectors and can therefore be deformed by specifying even fewer 3D points than we do now.

Acknowledgment. This work was supported in part by a Swiss National Science Foundation grant.

Bibliography

- [1] Acuna, D., Kar, A., Fidler, S.: Devil is in the Edges: Learning Semantic Boundaries from Noisy Annotations. In: Conference on Computer Vision and Pattern Recognition (2019) [4](#)
- [2] Akuna, D., Ling, H., Kar, A., Fidler, S.: Efficient Interactive Annotation of Segmentation Datasets with Polygon-RNN++. In: Conference on Computer Vision and Pattern Recognition (2018) [2](#), [3](#)
- [3] Çiçek, Ö., Abdulkadir, A., Lienkamp, S., Brox, T., Ronneberger, O.: 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In: Conference on Medical Image Computing and Computer Assisted Intervention. pp. 424–432 (2016) [1](#), [3](#), [4](#), [6](#)
- [4] Dalca, A., Yu, E., Golland, P., Fischl, B., Sabuncu, M., Iglesias, J.: Unsupervised Deep Learning for Bayesian Brain MRI Segmentation. In: Conference on Medical Image Computing and Computer Assisted Intervention (2019) [4](#)
- [5] 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: Conference on Medical Image Computing and Computer Assisted Intervention (2020) [1](#), [3](#)
- [6] Feng, X., Yang, J., Laine, A., Angelini, E.: Discriminative Localization in CNNs for Weakly-Supervised Segmentation of Pulmonary Nodules. In: Conference on Medical Image Computing and Computer Assisted Intervention. pp. 568–576 (2017) [2](#)
- [7] Freedman, D., Zhang, T.: Interactive Graph-Cut Based Segmentation with Shape Priors. In: Conference on Computer Vision and Pattern Recognition. pp. 755–62 (2005) [3](#)
- [8] Ge, W., Yanga, S., Yu, Y.: Multi-Evidence Filtering and Fusion for Multi-Label Classification, Object Detection and Semantic Segmentation Based on Weakly Supervised Learning. In: Conference on Computer Vision and Pattern Recognition (2018) [1](#)
- [9] Hsu, C., Hsu, K., Tsai, C., Lin, Y., Chuang, Y.: Weakly Supervised Instance Segmentation Using the Bounding Box Tightness Prior. In: Advances in Neural Information Processing Systems (2019) [1](#), [2](#)
- [10] Huang, Z., Wang, X., Wang, J., Liu, W., Wang, J.: Weakly-Supervised Semantic Segmentation Network with Deep Seeded Region Growing. In: Conference on Computer Vision and Pattern Recognition (2018) [1](#), [2](#)
- [11] Januszewski, M., Jain, V.: High-Precision Automated Reconstruction of Neurons with Flood-Filling Networks. *Nature Methods* (2018) [4](#)
- [12] Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active Contour Models. *International Journal of Computer Vision* **1**(4), 321–331 (1988) [2](#), [3](#)
- [13] Kavur, A., Selver, M.: CHAOS Challenge - Combined (CT-MR) Healthy Abdominal Organ Segmentation. In: arXiv Preprint (2020) [6](#)
- [14] Khoreva, A., Benenson, R., Hosang, J., Hein, M., Schiele, B.: Simple Does It: Weakly Supervised Instance and Semantic Segmentation. In: Conference on Computer Vision and Pattern Recognition. pp. 1665–1674 (2017) [1](#), [2](#)
- [15] Ling, H., Gao, J., Kar, A., Chen, W., Fidler, S.: Fast Interactive Object Annotation with Curve-Gcn. In: Conference on Computer Vision and Pattern Recognition. pp. 5257–5266 (2019) [2](#), [3](#)

- [16] Liu, X., Tsaftaris, S.: Semi-Supervised Meta-Learning with Disentanglement for Domain-Generalised Medical Image Segmentation. In: Conference on Medical Image Computing and Computer Assisted Intervention (2021) 4
- [17] Maninis, K., Caelles, S., Pont-Tuset, J., Gool, L.: Deep Extreme Cut: from Extreme Points to Object Segmentation. In: Conference on Computer Vision and Pattern Recognition (2018) 4
- [18] Mirikharaji, Z., Hamarneh, G.: Star Shape Prior in Fully Convolutional Networks for Skin Lesion Segmentation. In: Conference on Medical Image Computing and Computer Assisted Intervention. pp. 737–745 (2018) 3
- [19] Mortensen, E., Barrett, W.: Intelligent Scissors for Image Composition. In: ACM SIGGRAPH. pp. 191–198 (August 1995) 2, 3
- [20] and O. Russakovsky, A.B., Ferrari, V., Fei-Fei, L.: What's the Point: Semantic Segmentation with Point Supervision. In: European Conference on Computer Vision (2016) 1, 2
- [21] Peng, S., Jiang, W., Pi, H., Li, X., Bao, H., Zhou, X.: Deep Snake for Real-Time Instance Segmentation. In: Conference on Computer Vision and Pattern Recognition (2020) 2, 3
- [22] Roth, H., Zhang, L., Yang, D., Milletari, F., Xu, Z., Wang, X., Xu, D.: Weakly Supervised Segmentation from Extreme Points. In: Conference on Medical Image Computing and Computer Assisted Intervention (2019) 4
- [23] Shvets, A., Iglovikov, V.: Automatic Instrument Segmentation in Robot-Assisted Surgery Using Deep Learning. In: arXiv Preprint (2018) 1
- [24] Simpson, A., Menze, B.: A Large Annotated Medical Image Dataset for the Development and Evaluation of Segmentation Algorithms. In: arXiv Preprint (2019) 6
- [25] Spitzer, H., Kiwitz, K., Amunts, K., Harmeling, S., Dickscheid, T.: Improving Cytoarchitectonic Segmentation of Human Brain Areas with Self-Supervised Siamese Networks. In: Conference on Medical Image Computing and Computer Assisted Intervention (2018) 3, 4
- [26] Wang, Z., Acuna, D., Ling, H., Kar, A., Fidler, S.: Object Instance Annotation with Deep Extreme Level Set Evolution. In: European Conference on Computer Vision (2020) 4
- [27] Wickramasinghe, U., Knott, G., Fua, P.: Probabilistic Atlases to Enforce Topological Constraints. In: Conference on Medical Image Computing and Computer Assisted Intervention (2019) 3
- [28] Wickramasinghe, U., Knott, G., Fua, P.: Deep Active Surface Models. In: Conference on Computer Vision and Pattern Recognition (2021) 2, 4, 5
- [29] Wickramasinghe, U., Remelli, E., Knott, G., Fua, P.: Voxel2mesh: 3D Mesh Model Generation from Volumetric Data. In: Conference on Medical Image Computing and Computer Assisted Intervention (2020) 3, 6
- [30] Wolf, I., Vetter, M., Wegner, I., Nolden, M., Bötger, T., Hastenteufel, M., Schobinger, M., Kunert, T., Meinzer, H.: The Medical Imaging Interaction Toolkit (MITK): A Toolkit Facilitating the Creation of Interactive Software by Extending VTK and ITK. In: IEEE Transactions on Medical Imaging (2004) 5, 7
- [31] Xia, X., Kulis, B.: W-Net: A Deep Model for Fully Unsupervised Image Segmentation. In: arXiv Preprint (2017) 4
- [32] Yang, L., Wang, Y., Xiong, X., Yang, J., Katsaggelos, A.: Efficient Video Object Segmentation via Network Modulation. In: Conference on Computer Vision and Pattern Recognition (2018) 4
- [33] Zhao, T., Yin, Z.: Weakly Supervised Cell Segmentation by Point Annotation. In: IEEE Transactions on Medical Imaging (2020) 1, 2