



Leveraging Labeling Representations in Uncertainty-Based Semi-supervised Segmentation

Sukesh Adiga Vasudeva^(✉), Jose Dolz^{}, and Herve Lombaert^{}

ETS Montreal, Montreal, Canada
sukesh.adiga-vasudeva.1@ens.etsmtl.ca

Abstract. Semi-supervised segmentation tackles the scarcity of annotations by leveraging unlabeled data with a small amount of labeled data. A prominent way to utilize the unlabeled data is by consistency training which commonly uses a teacher-student network, where a teacher guides a student segmentation. The predictions of unlabeled data are not reliable, therefore, uncertainty-aware methods have been proposed to gradually learn from meaningful and reliable predictions. Uncertainty estimation, however, relies on multiple inferences from model predictions that need to be computed for each training step, which is computationally expensive. This work proposes a novel method to estimate the pixel-level uncertainty by leveraging the labeling representation of segmentation masks. On the one hand, a labeling representation is learnt to represent the available segmentation masks. The learnt labeling representation is used to map the prediction of the segmentation into a set of plausible masks. Such a reconstructed segmentation mask aids in estimating the pixel-level uncertainty guiding the segmentation network. The proposed method estimates the uncertainty with a single inference from the labeling representation, thereby reducing the total computation. We evaluate our method on the 3D segmentation of left atrium in MRI, and we show that our uncertainty estimates from our labeling representation improve the segmentation accuracy over state-of-the-art methods. Code is released at [GitHub](#).

Keywords: Semi-supervised learning · Segmentation · Labeling representation · Uncertainty

1 Introduction

Segmentation of organs or abnormal regions is a fundamental task in clinical applications, such as diagnosis, intervention and treatment planning. Deep learning techniques are driving progress in automating the segmentation task under the full-supervision paradigm [5, 16]. Training these models, however, relies on a large amount of pixel-level annotations, which require expensive clinical expertise [4].

Semi-supervised learning techniques alleviate the annotation scarcity by leveraging unlabeled data with a small amount of labeled data. Current semi-supervised

segmentation methods typically utilize the unlabeled data either in the form of pseudo labels [1, 32], regularization [6, 17, 19] or knowledge priors [8, 31]. For instance, self-training methods [1] generate pseudo labels from unlabeled data, which are used to retrain the network iteratively. A wide range of regularization-based methods has been explored for semi-supervised segmentation using adversarial learning [3, 17], consistency learning [2, 12, 14, 30], or co-training [19, 25, 27]. Adversarial methods encourage the segmentation of unlabeled images to be closer to those of the labeled images. In contrast, consistency and co-training methods encourage two or more segmentation predictions, either from the same or different networks, to be consistent under different perturbations of the input data. Such consistency-based methods are popular in semi-supervision due to their simplicity. Consequently, self-ensembling [9] and mean teacher-based [22] methods are often used in semi-supervised segmentation of medical images [2, 6, 12]. However, their generated predictions from the unlabeled images may not always be reliable. To alleviate this issue, uncertainty-aware regularization methods [15, 21, 24, 26, 30] are proposed to gradually add reliable target regions in predictions. This uncertainty scheme is also employed in co-training [27] and self-training [32] approaches to obtain reliable predictions. Although these methods perform well in low-labeled data regimes, their high computation and complex training techniques might limit their applicability to broader applications in practice. For instance, the uncertainty estimation is approximated via Monte-Carlo Dropout [7] or an ensembling, which requires multiple predictions per image. Co-training methods require two or more networks to be trained simultaneously, whereas self-training-based methods rely on costly iterations. Lastly, adversarial training is challenging in terms of convergence [20].

Prior-based methods in semi-supervised segmentation typically incorporate anatomical knowledge of the target object during training the model. For instance, He *et al.* [8] encode the unlabeled images in an autoencoder and combine the learnt features as prior knowledge in the segmentation networks. Recent attempts use signed distance maps (SDM) as shape constraints during training [11, 24, 29]. For instance, Le *et al.* [11] propose an additional task of predicting SDM and enforcing consistency with an adversarial loss. Zheng *et al.* [31] exploit a probabilistic atlas in their loss function. These knowledge-based methods require an additional task to constraints shape prior, or it requires aligned images.

These limitations motivate our approach, which leverages a learnt labeling representation to approximate the uncertainty. Our main idea is to mimic a shape prior by learning a representation using segmentation masks such that each prediction is mapped into a set of plausible segmentations. In contrast to [31], our approach does not require aligned images. The mapped segmentation is subsequently used to estimate the uncertainty maps to guide the segmentation network. We hypothesize that the proposed uncertainty estimates are more robust than those derived from the entropy variance, requiring multiple inferences strategy.

Our Contributions. We propose a novel way to estimate the pixel-wise uncertainty to guide the training of a segmentation model. In particular, we integrate

mean teacher approach used in semi-supervised segmentation, which consists of a student (S) and teacher (T) model, both having the same segmentation architecture. The overall objective function is defined as follows:

$$\mathcal{L} = \min_{\theta_s} \sum_{i=1}^N \mathcal{L}_s(f(x_i; \theta_s), y_i) + \lambda_c \sum_{i=1}^{N+M} \mathcal{L}_c(f(x_i; \theta_s, \eta), f(x_i; \theta_t; \eta')),$$

where $f(\cdot)$ denotes the segmentation network, and θ_s and θ_t are the learnable weights of the student and teacher models. The supervised loss \mathcal{L}_s measures the segmentation quality on the labeled data, whereas the consistency loss \mathcal{L}_c measures the prediction consistency of student and teacher models for the same input volume x_i under different perturbations (η, η'). The balance between supervised and unsupervised loss is controlled by a ramp-up weighting co-efficient λ_c . In the mean teacher training, the student model parameters are optimized with stochastic gradient descent (SGD), whereas exponential moving average (EMA) is employed at each training step t , i.e., $\theta_t = \alpha\theta_{t-1} + (1 - \alpha)\theta_s$ to update the teacher model parameters. Note that α is the smoothing coefficient of EMA that controls the update rate.

2.2 Labeling Representation Prior

Incorporating object shape prior in deep segmentation models is not obvious. One of the reasons is that, in order to integrate such prior knowledge during training, one needs to augment the learning objective with a differentiable term, which in the case of complex shapes is not trivial. To circumvent these difficulties, a simpler solution is to resort to an autoencoder trained with pixel-wise labels, which can represent anatomical priors and be used as a global regularizer during training. This strategy has been adopted for fully supervised training in [18] and as a post-processing step in [10] to correct the segmentation predictions. Motivated by this, we represent the available labels in a non-linear latent space using a denoising autoencoder (DAE) [23], which somehow mimics a shape prior. The DAE model consists of an encoder $f_e(\cdot)$ and a decoder module $f_d(\cdot)$ with a d -dimensional latent space as shown in the Fig. 1. The DAE is trained to reconstruct the clean labels y_i from its corrupted version \tilde{y}_i , which can be achieved with a mean squared error loss: $\frac{1}{H \times W \times D} \sum_v \|f_d(f_e(\tilde{y}_{i,v})) - y_{i,v}\|^2$.

2.3 Uncertainty from a Labeling Representation

The role of the uncertainty is to gradually update the student model with reliable target regions from the teacher predictions. Our proposed method estimates the uncertainty directly from the labeling representation network $f_d(f_e(\cdot))$, requiring only one inference step. First, we map the prediction from the teacher model p_i^t with a DAE model to produce a plausible segmentation \hat{p}_i^t . We subsequently estimate the uncertainty as the pixel-wise difference between the DAE output and the prediction, i.e., $U_i = \|\hat{p}_i^t - p_i^t\|^2$. Then, the reliable target for the consistency loss

is obtained as $e^{-\gamma U_i}$, similarly to [15], where γ is an uncertainty weighting factor empirically set to 1. Finally, our consistency loss is defined as:

$$\mathcal{L}_c(p_i^s, p_i^t) = \frac{\sum_v e^{-\gamma U_{i,v}} \|p_{i,v}^s - p_{i,v}^t\|^2}{\sum_v e^{-\gamma U_{i,v}}}$$

where v is a voxel. We jointly optimize the consistency loss \mathcal{L}_c and supervised loss \mathcal{L}_s as learning objectives, where \mathcal{L}_s uses the cross-entropy and dice losses.

3 Results

Our proposed method is compared with the state-of-the-art semi-supervised segmentation methods [11, 14, 15, 30]. We group the uncertainty-based methods to assess the effectiveness of our uncertainty estimation for segmentation. For a fair comparison, all experiments are run three times with a fixed set of seeds on the same machines, and their average results are reported.

Dataset and Evaluation Metrics. Our method is evaluated on the Left Atrium (LA) dataset from the 2018 Atrial Segmentation Challenge [28]. The dataset consists of 100 3D MR volumes of LA with an isotropic resolution of 0.625 mm^3 and corresponding segmentation masks. In our experiments, we use a 80/20 training/testing split and apply the same preprocessing as in [11, 14, 30]. The training set is partitioned into N/M labeled/unlabeled splits, fixed across all methods for each setting. We employ Dice Score Coefficient (DSC) and 95% Hausdorff Distance (HD) metrics to assess quantitative segmentation performance.

Implementation and Training details. Following [11, 14, 30], we use V-net [16] as backbone architecture for teacher, student and DAE models. The skip connections are removed, and a dense layer is added at the bottleneck layer for the DAE model. The student model is trained by a SGD optimizer with an initial learning rate (lr) of 0.1 and momentum 0.9 for 6000 iterations with a cosine annealing [13] decaying. The teacher weights are updated by an EMA with an update rate of $\alpha = 0.99$ as in [22]. The consistency weight is updated with Gaussian warming up function $\lambda_c = \beta * e^{-5(1-t/t_{max})^2}$, where t and t_{max} denotes current and maximum training iterations, and β is set to 0.1, as in [30]. The DAE model is also trained with SGD with $lr = 0.1$, momentum of 0.9 and decaying the lr by 2 for every 5000 iterations. Input to both segmentation and DAE networks are random cropped to $112 \times 112 \times 80$ size and employ online standard data augmentation techniques such as random flipping and rotation. In addition, input labels to the DAE model are corrupted with a random swapping of pixels around class boundaries, morphological operations (erosion and dilation), resizing and adding/removing shapes. The batch size is set to 4 in both networks. Input batch for segmentation network uses two labeled and unlabeled data. For testing, generating segmentation predictions uses the sliding window strategy, and the method is evaluated at the last iteration as in [30]. Our experiments were run on an NVIDIA RTX A6000 GPU with PyTorch 1.8.0.

Table 1. Segmentation results on the LA test set for 10% labeled data experiments averaged over three runs. Uncertainty methods with K inferences are grouped at the bottom, while $K = -$, indicates non-uncertainty methods.

Methods	#K	N/M	DSC (%)	HD (mm)
Upper bound	-	80/0	91.23 ± 0.44	6.08 ± 1.84
Lower bound	-	8/0	76.07 ± 5.02	28.75 ± 0.72
MT [22]	-	8/72	78.22 ± 6.89	16.74 ± 4.80
SASSnet [11]	-	8/72	83.70 ± 1.48	16.90 ± 1.35
DCT [14]	-	8/72	83.10 ± 0.26	12.62 ± 1.44
UAMT [30]	8	8/72	85.09 ± 1.42	18.34 ± 2.80
URPC [15]	1	8/72	84.47 ± 0.31	17.11 ± 0.60
Ours	1	8/72	86.58 ± 1.03	11.82 ± 1.42

Comparison with the State-of-the-Art. We now compare our method with relevant semi-supervised segmentation approaches under the 10% and 20% labeled data settings and report their results in Tables 1 and 2. Non-uncertainty-based method such as MT [22], DCT [14], and SASSnet [11] are grouped in the middle of the table, while uncertainty-based methods UAMT [30], URPC [15]¹ and our methods are grouped at the bottom of each table. The upper and lower bound from the backbone architecture V-net [16] are reported in the top.

In the first setting, 10% of labeled data is used, and the remaining images are used as unlabeled data. From Table 1, we can observe that leveraging unlabeled data improves the lower bound in all baselines. The uncertainty-based baselines seem to improve the segmentation performance by 1% in Dice score compared to non-uncertainty-based baselines. However, their performance drops in terms of HD up to 5 mm. Among baseline methods, UAMT and DCT achieve the best Dice and HD scores, respectively. Compared to these best performing baselines, our method brings 1.5% and 0.8 mm improvements in Dice and HD scores. Moreover, uncertainty estimation in our method requires a single inference from a labeling representation, whereas UAMT uses $K = 8$ inferences per training step to obtain an uncertainty map.

Furthermore, we also validate our method on the 20% of labeled data scenario, whose results are reported in Tables 2. Results demonstrate a similar trend as compared to the 10% experiments. The uncertainty-based baselines improve 1% in terms of Dice and drop up to 1 mm in HD, compared to non-uncertainty-based methods. Our method improves the best performing baseline in both Dice and HD scores. Particularly, our method improves the HD score by 2.5 mm compared to the best performing baseline (SASSnet).

Visual results of different segmentation results are depicted in Fig. 2. In the top row of the figure, the segmentation of SASSnet produces holes in segmentation, and their method employs a post-processing tool to improve the segmenta-

¹ Note that URPC [15] use multi-scale 3D U-Net [5] architecture.

Table 2. Segmentation results on the LA test set for 20% labeled data experiments averaged over three runs. Uncertainty methods with K inferences are grouped at the bottom, while $K = -$, indicates non-uncertainty methods.

Methods	#K	N/M	DSC (%)	HD (mm)
Upper bound	-	80/0	91.23 ± 0.44	6.08 ± 1.84
Lower bound	-	16/0	81.46 ± 2.96	23.61 ± 4.94
MT [22]	-	16/64	86.06 ± 0.81	11.63 ± 3.4
SASSnet [11]	-	16/64	87.81 ± 1.45	10.18 ± 0.55
DCT [14]	-	16/64	87.35 ± 1.26	10.25 ± 2.49
UAMT [30]	8	16/64	87.78 ± 1.03	11.1 ± 1.91
URPC [15]	1	16/64	88.58 ± 0.10	13.1 ± 0.60
Ours	1	16/64	88.60 ± 0.82	7.61 ± 0.78

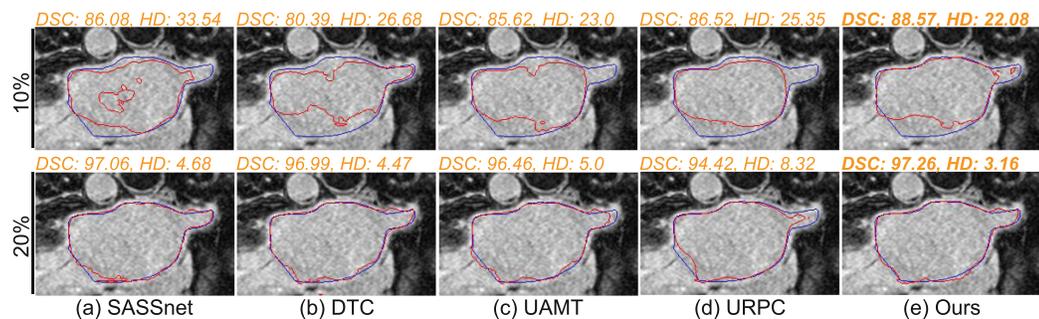


Fig. 2. Qualitative comparison under 10% and 20% annotation setting. DSC (%) and HD (mm) scores are mentioned at the top of each image. Coloring is the prediction (Red) and ground truth (Blue). (Color figure online)

tion, which is avoided for a fair comparison. DTC captures the challenging top right side region in segmentation; however, the prediction is under-segmented and noisy. The uncertainty-based methods improve the segmentation in UAMT and produce smooth segmentation boundaries in URPC. Our method improves the segmentation region further compared to URPC. In the case of 20% labeled data experiments, all methods improve the segmentation due to having access to more labels during training, while the boundary regions are either under or over-segmented. Our method produces better and smoother segmentation, which can be due to the knowledge derived from the labeling representation.

Ablation Study. To validate the effectiveness of our uncertainty estimation on segmentation performance, two experiments are conducted by adopting threshold strategy and entropy scheme from UAMT. Particularly, a threshold strategy is used in consistency, whereas entropy is used to estimate the uncertainty, and their results are reported in Table 3. Compared to UAMT, our threshold and entropy experiments significantly improve the segmentation performance in HD

and Dice scores, while our proposed method (L2-based exponential uncertainty) achieves the best performance. These results show the merit of our labeling representation for uncertainty estimation. Furthermore, we report the ablation on uncertainty weight γ and consistency weight β , in Table 4. Results demonstrate that $\gamma = 1$ is best for our method, while for $\beta = 1$ our method further improves Dice and HD scores; however, we report on $\beta = 0.1$ in all experiments for a fair comparison. Overall, for most of the γ and β values, our method is consistently better than UAMT baselines, demonstrating the robustness of our approach.

Table 3. Effectiveness of our proposed uncertainty estimation on segmentation results using different strategies.

Methods	N/M	DSC (%)	HD (mm)
UAMT [30]	8/72	85.09 \pm 1.42	18.34 \pm 2.80
Ours (threshold)	8/72	85.39 \pm 0.91	12.96 \pm 3.05
Ours (entropy)	8/72	85.92 \pm 1.52	11.16 \pm 0.82
Ours	8/72	86.58 \pm 1.03	11.82 \pm 1.42

Table 4. Evaluating the γ and β values under 10% annotation setting.

$\gamma, \beta = 0.1$	DSC (%)	HD (mm)	$\beta, \gamma = 1$	DSC (%)	HD (mm)
0.1	85.30 \pm 1.17	13.51 \pm 2.66	0.01	84.89 \pm 0.92	11.84 \pm 2.79
0.5	85.28 \pm 0.60	14.01 \pm 4.44	0.05	85.88 \pm 1.44	10.98 \pm 1.85
1	86.58 \pm 1.03	11.82 \pm 1.42	0.1	86.58 \pm 1.03	11.82 \pm 1.42
2	85.84 \pm 1.39	12.13 \pm 3.43	0.5	86.54 \pm 0.74	12.42 \pm 1.31
5	84.87 \pm 0.85	15.28 \pm 1.76	1	86.89 \pm 0.6	9.85 \pm 0.82

4 Conclusion

We presented a novel labeling representation-based uncertainty estimation for the semi-supervised segmentation. Our method produces an uncertainty map from a labeling representation network, which guides the reliable regions of prediction for the segmentation network, thereby achieving better segmentation results. Results demonstrate that the proposed method achieves the best performance compared to state-of-the-art baselines on left atrium segmentation from 3D MR volumes in two different settings. The ablation studies demonstrate the effectiveness and robustness of our uncertainty estimation compared to the entropy-based method. Our proposed uncertainty estimation from the labeling representation approach can be adapted to a broader range of applications where it is crucial to obtain a reliable prediction.

Acknowledgments. This research work was partly funded by the Canada Research Chair on Shape Analysis in Medical Imaging, the Natural Sciences and Engineering Research Council of Canada (NSERC), and the Fonds de Recherche du Quebec (FQRNT).

References

1. Bai, W., et al.: Semi-supervised learning for network-based cardiac MR image segmentation. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10434, pp. 253–260. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66185-8_29
2. Bortsova, G., Dubost, F., Hogeweg, L., Katramados, I., de Bruijne, M.: Semi-supervised medical image segmentation via learning consistency under transformations. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11769, pp. 810–818. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32226-7_90
3. Chaitanya, K., Karani, N., Baumgartner, C.F., Becker, A., Donati, O., Konukoglu, E.: Semi-supervised and task-driven data augmentation. In: Chung, A.C.S., Gee, J.C., Yushkevich, P.A., Bao, S. (eds.) IPMI 2019. LNCS, vol. 11492, pp. 29–41. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-20351-1_3
4. Cheplygina, V., de Bruijne, M., Pluim, J.P.: Not-so-supervised: a survey of semi-supervised, multi-instance, and transfer learning in medical image analysis. *Media* **54**, 280–296 (2019)
5. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_49
6. Cui, W., et al.: Semi-supervised brain lesion segmentation with an adapted mean teacher model. In: Chung, A.C.S., Gee, J.C., Yushkevich, P.A., Bao, S. (eds.) IPMI 2019. LNCS, vol. 11492, pp. 554–565. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-20351-1_43
7. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: ICML, pp. 1050–1059. PMLR (2016)
8. He, Y., et al.: Dense biased networks with deep priori anatomy and hard region adaptation: semi-supervised learning for fine renal artery segmentation. *Media* **63**, 101722 (2020)
9. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. arXiv preprint [arXiv:1610.02242](https://arxiv.org/abs/1610.02242) (2016)
10. Larrazabal, A.J., Martínez, C., Glocker, B., Ferrante, E.: Post-DAE: anatomically plausible segmentation via post-processing with denoising autoencoders. *IEEE TMI* **39**(12), 3813–3820 (2020)
11. Li, S., Zhang, C., He, X.: Shape-aware semi-supervised 3D semantic segmentation for medical images. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12261, pp. 552–561. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_54
12. Li, X., Yu, L., Chen, H., Fu, C.W., Xing, L., Heng, P.A.: Transformation-consistent self-ensembling model for semisupervised medical image segmentation. *IEEE Trans. Neural Netw. Learn. Syst.* **32**(2), 523–534 (2020)
13. Loshchilov, I., Hutter, F.: SGDR: stochastic gradient descent with warm restarts. arXiv preprint [arXiv:1608.03983](https://arxiv.org/abs/1608.03983) (2016)

14. Luo, X., Chen, J., Song, T., Wang, G.: Semi-supervised medical image segmentation through dual-task consistency. In: AAAI, vol. 35, pp. 8801–8809 (2021)
15. Luo, X., et al.: Efficient semi-supervised gross target volume of nasopharyngeal carcinoma segmentation via uncertainty rectified pyramid consistency. In: de Bruijne, M., et al. (eds.) MICCAI 2021. LNCS, vol. 12902, pp. 318–329. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87196-3_30
16. Milletari, F., Navab, N., Ahmadi, S.A.: V-Net: fully convolutional neural networks for volumetric medical image segmentation. In: 3DV, pp. 565–571. IEEE (2016)
17. Nie, D., Gao, Y., Wang, L., Shen, D.: ASDNet: attention based semi-supervised deep networks for medical image segmentation. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11073, pp. 370–378. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00937-3_43
18. Oktay, O., et al.: Anatomically constrained neural networks (ACNNs): application to cardiac image enhancement and segmentation. *IEEE TMI* **37**(2), 384–395 (2017)
19. Peng, J., Estrada, G., Pedersoli, M., Desrosiers, C.: Deep co-training for semi-supervised image segmentation. *Pattern Recogn.* **107**, 107269 (2020)
20. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training GANs. In: *NeurIPS*, vol. 29 (2016)
21. Sedai, S., et al.: Uncertainty guided semi-supervised segmentation of retinal layers in OCT images. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11764, pp. 282–290. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32239-7_32
22. Tarvainen, A., Valpola, H.: Mean teachers are better role models: weight-averaged consistency targets improve semi-supervised deep learning results. In: *NeurIPS*, vol. 30 (2017)
23. Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., Manzagol, P.A., Bottou, L.: Stacked denoising autoencoders: learning useful representations in a deep network with a local denoising criterion. *JMLR* **11**(12), 3371–3408 (2010)
24. Wang, K., et al.: Triple-uncertainty guided mean teacher model for semi-supervised medical image segmentation. In: de Bruijne, M., et al. (eds.) MICCAI 2021. LNCS, vol. 12902, pp. 450–460. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87196-3_42
25. Wang, P., Peng, J., Pedersoli, M., Zhou, Y., Zhang, C., Desrosiers, C.: Self-paced and self-consistent co-training for semi-supervised image segmentation. *Media* **73**, 102146 (2021)
26. Wang, Y., et al.: Double-uncertainty weighted method for semi-supervised learning. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12261, pp. 542–551. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_53
27. Xia, Y., et al.: 3D semi-supervised learning with uncertainty-aware multi-view co-training. In: *IEEE/CVF WCCV*, pp. 3646–3655 (2020)
28. Xiong, Z., et al.: A global benchmark of algorithms for segmenting the left atrium from late gadolinium-enhanced cardiac magnetic resonance imaging. *Media* **67**, 101832 (2021)
29. Xue, Y., et al.: Shape-aware organ segmentation by predicting signed distance maps. In: AAAI, vol. 34, pp. 12565–12572 (2020)
30. Yu, L., Wang, S., Li, X., Fu, C.-W., Heng, P.-A.: Uncertainty-aware self-ensembling model for semi-supervised 3D left atrium segmentation. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11765, pp. 605–613. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_67

31. Zheng, H., et al.: Semi-supervised segmentation of liver using adversarial learning with deep atlas prior. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11769, pp. 148–156. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32226-7_17
32. Zheng, H., et al.: Cartilage segmentation in high-resolution 3D micro-CT images via uncertainty-guided self-training with very sparse annotation. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12261, pp. 802–812. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_78