



Meta Corrupted Pixels Mining for Medical Image Segmentation

Jixin Wang¹, Sanping Zhou¹, Chaowei Fang², Le Wang¹, and Jinjun Wang¹(✉)

¹ Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University,
Xi'an, Shaanxi, People's Republic of China
jinjun@mail.xjtu.edu.cn

² School of Artificial Intelligence, Xidian University,
Xi'an, Shaanxi, People's Republic of China

Abstract. Deep neural networks have achieved satisfactory performance in piles of medical image analysis tasks. However the training of deep neural network requires a large amount of samples with high-quality annotations. In medical image segmentation, it is very laborious and expensive to acquire precise pixel-level annotations. Aiming at training deep segmentation models on datasets with probably corrupted annotations, we propose a novel Meta Corrupted Pixels Mining (MCPM) method based on a simple meta mask network. Our method is targeted at automatically estimate a weighting map to evaluate the importance of every pixel in the learning of segmentation network. The meta mask network which regards the loss value map of the predicted segmentation results as input, is capable of identifying out corrupted layers and allocating small weights to them. An alternative algorithm is adopted to train the segmentation network and the meta mask network, simultaneously. Extensive experimental results on LIDC-IDRI and LiTS datasets show that our method outperforms state-of-the-art approaches which are devised for coping with corrupted annotations.

Keywords: Meta Corrupted Pixels Mining · Deep neural network · Medical image segmentation

1 Introduction

Recent years have witnessed the blooming of Deep Neural Networks (DNNs) in medical image analysis, including image segmentation, image registration, image reconstruction [14], and etc. Due to the powerful representation capability of DNN, significant progress has been achieved in medical image analysis. However, training a DNN usually requires a large number of high-quality labeled samples, which is hard to acquire in various applications. For example, it is very expensive to generate a precise segment of input image, because the pathological tissue needs to be marked by professional radiologists [26, 27]. As a result, a question was naturally raised: How can we train a powerful segmentation network only using a small number of high-quality labeled samples?




Fig. 1. Corrupted labels effect network’s performance

To address this situation, researchers have paid much attention to train DNNs in a semi-supervised manner. For example, Yang et al. [23] presented an active learning method for 2D biomedical image segmentation, which can improve segmentation accuracy through suggesting the most effective rather than all samples for labeling. In [25], Zhao et al. applied a modified Mask R-CNN to volumetric data for instance segmentation, and they used bounding boxes for all instances and voxel-wise labels for a small proportion of instances. Nie et al. [17] proposed an attention based semi-supervised deep networks, which adopted the adversarial learning strategy to deal with the insufficient data problem in training complex networks. In practice, the success of these semi-supervised methods depends on mining a kind of knowledge which can be used to find out more accurate labels in the training process. However, most of the existing methods use a fixed prior knowledge to guide the pseudo label estimation. Therefore, they are very unstable when dealing with training samples with complex noise distributions. As shown in Fig. 1, the segmentation network’s results are seriously affected when corrupted labels are taken as supervisory signals. This phenomenon reveals that mining corrupted labels is a critical issue in semi-supervised image segmentation.

In this paper, we design a novel Meta Corrupted Pixels Mining (MCPM) method for medical image segmentation, which can alleviate the impacts of corrupted labels in the training process. To achieve this goal, we design a simple meta mask network to protect the training of the segmentation network from the influence of pixels with incorrect labels. Specifically, the meta mask network absorbs in the loss value map of the segmentation prediction as input, and estimate a weight map indicating the importance of every pixel in the training of the segmentation network. Once the meta mask network is learned, small weights are allocated to pixels with corrupted labels. Therefore influences from these pixels are weakened when updating the segmentation network. In the training process, we update the segmentation network and meta mask network in an alternate manner, which can learn a powerful segmentation network from images with corrupted labels. The main contributions of this work can be highlighted as follows:

- We design a novel meta learning framework to mine pixels with corrupted labels during the process of training a segmentation network.

- Based on the fully convolutional structure, we build up a meta mask network which can automatically estimate pixel-wise importance factors for mitigating the influence of corrupted labels.
- Extensive experiments on both LIDC-IDRI and LiTS datasets indicate that our method achieves the state-of-the-art performance in medical image segmentation with incorrect labels.

2 Related Works

Because our method takes U-Net [19] as segmentation network and applies the meta learning regime [1] to mine pixels of corrupted labels, we briefly review a few existing works in terms of U-Net and meta learning in the following paragraphs.

Methods Based on U-Net. This type of methods aim to design a powerful network structure, which can obtain accurate segmentation results at the output layer. In [19], Ronneberger et al. proposed a well-known U-shaped structure for 2D medical image segmentation, in which the low-level and high-level feature are recursively concatenated together from top to down, to improve segmentation results. Inspired by this idea, a number of variants have been introduced in the past few years. For example, Milletari et al. [15] extended the U-shaped structure into 3D version and built an objective function and adopted Dice coefficient maximisation to supervise the training process. In [13], Kohl et al. proposed a generative segmentation model based on a combination of a U-Net, in which a conditional variational autoencoder is designed to produce an unlimited number of plausible hypotheses. Because its superior performance in medical image segmentation, we simply choose U-Net as our segmentation network. Then, we concentrate on designing a meta learning regime which can help learn a robust segmentation network from training samples with corrupted labels.

Methods Based on Meta Learning. This kind of methods aim to learn a kind of knowledge which can be used to guide the training of the network for solving the target problem [1], which has a wide application in the few-shot learning community. For example, a number of methods, such as FWL [8]. MentorNet [11] used the concept of meta learning to learn an adaptive weighting function to make the training process more robust to noisy images. However, the meta learners used in these methods have complex forms and require complicated inputs, which are very hard to be implemented in the training process. To overcome this problem, Ren et al. [18] proposed a novel meta learning algorithm which can learn an implicit function to assign weights to training samples based on their gradient directions. In [20], Shu et al. designed a meta weight network to learn an explicit function which can impose small weights to noisy samples, therefore the noisy samples will not severely affect the training process. The difference between our proposed model and the meta weight network is that, we design a simple meta mask network to learn a knowledge which can mine the pixels of corrupted labels, so as to learn a powerful segmentation network from low-quality labeled images.




Fig. 2. The architecture and workflow of one loop

3 Meta Corrupted Pixels Mining Algorithm

We propose a novel MCPM method which can learn a powerful segmentation network from images with corrupted labels. Given a small set of images with clean labels and a large set of images with corrupted labels, our method is capable of identifying out the pixels with corrupted labels, and excluding them during the optimization procedure. As shown in Fig. 2, our network architecture is constituted by two modules: (1) a U-Net based module for segmentation; and (2) a meta mask network for mining pixels with corrupted labels. In the following paragraphs, we will introduce our method in detail.

3.1 Objective Functions

Let $\mathcal{S} = \{(\mathbf{X}^i, \mathbf{Y}^i)\}_{i=1}^N$ represent training images with probably noisy segmentation annotations, in which the width and height of training images are denoted by w and h respectively, and N indicates the number of training samples. Besides, $\mathbf{Y}^i \in \{0, 1\}^{h \times w \times c}$ denote the corrupted labels, where c is the number of classes to be segmented out. First of all, we set up a segmentation network based on U-Net [19], which yields a pixel-level prediction \mathbf{P}^i from input image \mathbf{X}^i . We define $\mathbf{P}^i = \mathcal{F}(\mathbf{X}^i; \mathbf{W})$ where \mathbf{W} represents parameters of the segmentation network. To learn \mathbf{W} , an objective function is usually adopted to calculate pixel-wise loss values as function $L_{xy}^i = \text{loss}(\mathbf{P}_{xy}^i, \mathbf{Y}_{xy}^i)$, where $x \in [1, h]$ and $y \in [1, w]$ indicate the pixel coordinates. Here, the cross entropy loss function is used as the objective function.

As mentioned above, there might exist errors in segmentation annotations. These errors will severely hamper the optimization procedure, for example, providing incorrect gradient directions in the training process. A straightforward approach to cope with this issue is ignoring these pixels with incorrect labels through reweighting loss values. Inspired from [21], we design our meta mask network in a fully convolutional structure, which can learn an accurate mask map \mathbf{R}^i for the input loss value map \mathbf{L}^i . We denote $\mathbf{R}^i = \mathcal{G}(\mathbf{L}^i; \Theta)$, where R_{xy}^i indicates the reweighting factor of the pixel at (x, y) , and Θ represents parameters of our meta mask network. Given a fixed Θ , the optimized solution to \mathbf{W}




Fig. 3. Illustration of how meta mask network works in the training process

can be found through minimizing the following objective function:

$$\mathbf{W}^*(\Theta) = \arg \min_{\mathbf{W}} \frac{1}{Nhw} \sum_{i=1}^N \sum_{x=1}^h \sum_{y=1}^w R_{xy}^i L_{xy}^i. \quad (1)$$

To learn the parameters of our meta mask network, we introduce an additional meta dataset $\hat{\mathcal{S}} = \{(\hat{\mathbf{X}}^j, \hat{\mathbf{Y}}^j)\}_{j=1}^M$ which contains images with high-quality annotations. In particular, given an input image $\hat{\mathbf{X}}^j$ and optimized parameters $\mathbf{W}^*(\Theta)$, the segmentation network will produce a pixel-wise prediction map $\hat{\mathbf{P}}^j = \mathcal{F}(\hat{\mathbf{X}}^j, \mathbf{W}^*(\Theta))$ at the output layer. Again, we can obtain a loss value map $\hat{\mathbf{L}}^j$ through comparing $\hat{\mathbf{P}}^j$ against $\hat{\mathbf{Y}}^j$ according to the cross entropy loss function. With the optimized \mathbf{W} , the optimized solution to Θ can be acquired through minimizing the following objective function:

$$\Theta^*(\mathbf{W}) = \arg \min_{\Theta} \frac{1}{Mhw} \sum_{j=1}^M \sum_{x=1}^h \sum_{y=1}^w \hat{L}_{xy}^j. \quad (2)$$

In the training process, we update \mathbf{W} and Θ in an alternation manner. As a result, the Θ can cope with the varying \mathbf{W} , which is beneficial to effectively mine more corrupted pixels from the predictions of the segmentation network.

3.2 Meta Mask Network

We take a fully convolutional structure to design our meta mask network, which can explore more local information to locate the pixels with corrupted labels. The particularities of the network are two aspects: (1) It has two convolutional layers with kernels in size of 3×3 and 5×5 , which can extract multi-scale context information from \mathbf{L}^i . (2) The resulting outputs and input are further fused through another 1×1 convolutional layer, giving rise to the final mask map \mathbf{R}^i . This simple structure can be trained under the guidance of a few high-quality labeled samples, which will in turn help train a powerful segmentation network by using a large number of low-quality labeled samples.

In Fig. 3, we visualize how our meta mask network alleviates the side effect of corrupted labels in the training process, in which: (1) shows the input image and ground truth annotations; (2) indicates the corrupted labels; (3) represents the predicted result obtained by the segmentation network; (4) denotes the mined pixels of corrupted labels. As we can observe in (5) and (6), the loss between (1)

and (3) is very small, while the loss between (2) and (3) is large. This indicates that our meta network can help train a powerful segmentation network with a large number of images accompanied with corrupted labels.

3.3 Network Optimization

We employ the iterative optimization algorithm to train our model. It is implemented with a single loop and mainly contains the following steps.

- At first, \mathbf{W} and Θ are randomly initialized as \mathbf{W}^0 and Θ^0 .
- For the t -th iteration, the parameters of the segmentation network are temporarily renovated as in Eq. (3), via one step of gradient descent in minimizing the objective function (1),

$$\mathbf{W}'^{(t)}(\Theta) = \mathbf{W}^{(t)} - \alpha \frac{1}{Nhw} \sum_{i=1}^N \sum_{x=1}^h \sum_{y=1}^w R_{xy}^{i(t)} \frac{\partial L_{xy}^i}{\partial \mathbf{W}} \Big|_{\mathbf{W}^{(t)}}, \quad (3)$$

where α is the learning rate. $R_{xy}^{i(t)}$ is computed through feeding the loss value map into the meta mask network with parameters $\Theta^{(t)}$.

- Then Θ can be updated via optimizing the objective function (2),

$$\Theta^{(t+1)} = \Theta^{(t)} - \beta \frac{1}{Mhw} \sum_{j=1}^M \sum_{x=1}^h \sum_{y=1}^w \frac{\partial \hat{L}_{xy}^j}{\partial \mathbf{W}'(\Theta)} \Big|_{\mathbf{W}'^{(t)}} \frac{\partial \mathbf{W}'(\Theta)}{\partial \Theta} \Big|_{\Theta^{(t)}}, \quad (4)$$

where β is the learning rate.

- Finally, \mathbf{W} is updated through minimizing objective function (1),

$$\mathbf{W}^{(t+1)} = \mathbf{W}^{(t)} - \alpha \frac{1}{Nhw} \sum_{i=1}^N \sum_{x=1}^h \sum_{y=1}^w R_{xy}^{i(t+1)} \frac{\partial L_{xy}^i}{\partial \mathbf{W}} \Big|_{\mathbf{W}^{(t)}}. \quad (5)$$

Here $R_{xy}^{i(t+1)}$ is computed through feeding the loss value map into the meta mask network with updated parameters $\Theta^{(t+1)}$.

3.4 Discussion

Under the guidance of a small meta set with clean annotations, the meta mask network is learned in a gradient descent by gradient descent manner as shown in (3) and (4). The update of parameters in the meta mask network is dependent on the gradients of losses calculated on pixels from both meta and training images. After putting (3) into (4), it can be easily observed that the ascending direction of the weight coefficient of every pixel relies on the inner product (it can also be interpreted as a similarity metric) between the gradient of the pixel (formulated as $\frac{\partial L_{xy}^i}{\partial \mathbf{W}} \Big|_{\mathbf{W}^{(t)}}$) and the average gradient of pixels of meta images (formulated as $\frac{1}{Mhw} \sum_{j=1}^M \sum_{x=1}^h \sum_{y=1}^w \frac{\partial \hat{L}_{xy}^j}{\partial \mathbf{W}'(\Theta)} \Big|_{\mathbf{W}'^{(t)}}$). A positive inner product pushes the

parameters of the meta mask network towards a direction which can give rise to a larger weighting coefficient for the corresponding pixel; a negative inner product pushes the network towards the opposite direction. This is the reason why our method can effectively identify corrupted pixels.

4 Experiments

4.1 Datasets and Metrics

Two datasets are exploited to validate the superiority of our method in medical image segmentation with noisy annotations, including LIDC-IDRI (Lung Image Database Consortium and Image Database Resource Initiative) [2, 3, 7] and LiTS (Liver Tumor Segmentation) [9]. 64×64 patches covering lesions are cropped out as training or testing samples.

- (1) The LIDC-IDRI dataset contains 1018 lung CT scans from 1010 patients with lesion masks annotated by four experts. 3591 patches are cropped out. They are split into a training set of 1906 images and a testing set of 1385 images. The remain 300 images are used as the meta set.
- (2) The LiTS dataset includes 130 abdomen CT scans accompanied with annotations of liver tumors. 2214 samples are sampled from this dataset. 1471, 300 and 443 images are used for training, meta weight learning, and testing respectively.

Three metrics, including IOU (also referred as the Jaccard Index), Dice coefficient and Hausdorff distance, are employed for quantitatively measuring performances of segmentation algorithms.

Synthesizing Noisy Annotations. In practice, it is difficult to localize the boundary of the target region during the annotating procedure. Considering this phenomenon, we synthesize noisy annotations through creating masks which loosely encompasses target lesions. We use 2 operators to simulate corrupted annotations. 1) The dilation morphology operator is employed to extend the foreground region by several pixels (randomly drawn from $[0, 6]$). 2) The toolkit of deformation provided in ElasticDeform [6, 19, 22], which includes more complicate operations such as rotation, translation, deformation and morphology dilation, is used to contaminate ground-truths of training images. In our experiment, only a part of samples are contaminated with the above strategies. We denote the percent of images which are selected out to generate noisy labels as r .

4.2 Implementation Details

Adam and SGD is used to optimize to network parameters on LIDC-IDRI and LiTS, respectively. The learning rates α and β are initialized as 10^{-4} and 10^{-3} respectively, and decayed by 0.1 in 20^{th} epoch and 40^{th} epoch. The batch size is set as 128. All models are trained with 120 epochs.

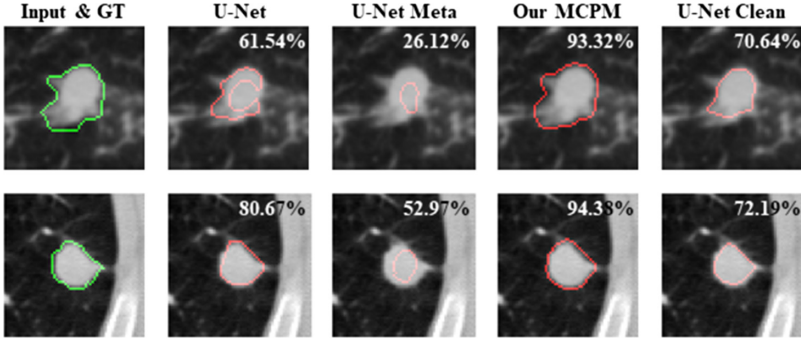


Fig. 4. Visualization of segmentation results. Green and red contours indicate the ground-truths and segmentation results, respectively. The Dice value is shown at the top-right corner, and our method produces much better results than other methods. (Color figure online)

4.3 Experimental Results

Comparison with Existing Methods. Without specification, r is set to 40% in experiments of this section which means annotations of 40% training images are contaminated. We compare our method against 7 existing segmentation models which are proposed to deal with ambiguous, low-quality or insufficient annotations on the LIDC-IDRI dataset: **Prob U-Net** [13], **Phi-Seg** [5], **UA-MT** [24] modified for 2D segmentation, **Curriculum** [12], **Few-Shot GAN** [16], **Quality Control** [4], **U^2 Net** [10], and **MWNet** [20] which is integrated with U-Net. All above models and the baseline U-Net are trained with mixed images of the training set and the meta set. We also implement another variant of U-Net which is trained merely using images from the meta set (indicated by ‘U-Net Meta’). We also report the result of U-Net trained using images with clean labels (indicated by ‘U-Net Clean’). As shown in Table 1. Our method performs significantly better than other methods. For example, the Dice value of our method surpasses that of the second best method MWNet by 3.4%. Additionally, our method outperforms baseline U-Net models by a large margin. It even achieves promising performance which is comparable to the result of ‘U-Net Clean’. This indicates that the impact of incorrect annotations fabricated as in Sect. 4.1 is almost eliminated. Visualization examples are shown in Fig. 4.

Results with Various r -s. In this section, we vary the percent of noisy images r from 0 to 0.8. The segmentation results of four methods on LIDC-IDRI and LiTS are presented in Table 2. On the LIDC-IDRI dataset, our method performs better than other methods when noises are introduced into the training set. On the LiTS dataset, our method exceeds other methods consistently under all cases.

Table 1. Results of segmentation models on LIDC-IDRI.

Noisy	Dilation			ElasticDeform		
Model name	mIOU	Dice	Hausdorff	mIOU	Dice	Hausdorff
U-Net	62.53	75.56	1.9910	65.01	76.17	1.9169
U-Net Meta	60.91	72.21	2.0047	60.91	72.21	2.0047
Prob U-Net	66.42	78.39	1.8817	68.43	79.50	1.8757
Phi-Seg	67.01	79.06	1.8658	68.55	81.76	1.8429
UA-MT	68.18	80.98	1.8574	68.84	82.47	1.8523
Curriculum	67.78	79.54	1.8977	68.18	81.30	1.8691
Few-Shot GAN	67.74	78.11	1.9137	67.93	77.83	1.9223
Quality Control	65.00	76.50	1.9501	68.07	77.68	1.9370
U^2 Net	65.92	76.01	1.9666	67.20	77.05	1.9541
MWNet	71.56	81.17	1.7762	71.89	81.04	1.7680
Our MCPM	74.69	84.64	1.7198	75.79	84.99	1.7053
U-Net Clean	<u>75.73</u>	<u>83.91</u>	<u>1.7051</u>	<u>75.73</u>	<u>83.91</u>	<u>1.7051</u>

Table 2. Results (mIOU) of segmentation models using various r -s.

Dataset name	LIDC-IDRI					LiTS				
r	0.8	0.6	0.4	0.2	0	0.8	0.6	0.4	0.2	0
U-Net	42.64	51.23	62.53	69.88	75.73	37.18	43.55	46.41	51.20	61.07
Prob U-Net	52.13	60.81	66.42	71.03	76.37	40.16	45.90	49.22	53.97	60.60
MWNet	61.28	67.33	71.56	72.07	74.40	43.14	44.97	51.96	58.65	59.18
Our MCPM	67.60	68.97	74.69	74.87	75.26	45.09	48.76	55.17	62.04	62.68

5 Conclusion

We proposed a novel Meta Corrupted Pixels Mining method to alleviate the side effect of corrupted label in medical image segmentation. Given a small number of high-quality labeled images, the deduced learning regime make our meta mask network able to locate the pixels having corrupted labels, which can be used to help train a powerful segmentation network from a large number of low-quality labeled images. Extensive experiments on two datasets, LIDC-IDRI and LiTS, show that the proposed method can achieve the state-of-the-art performance in medical image segmentation.

Acknowledgments. This work is jointly supported by the National Key Research and Development Program of China under Grant No. 2017YFA0700800, the National Natural Science Foundation of China Grant No. 61629301, 61976171, and the Key Research and Development Program of Shaanxi Province of China under Grant No. 2020GXLH-Y-008.

References

1. Andrychowicz, M., et al.: Learning to learn by gradient descent by gradient descent. In: *NeurIPS*, pp. 3981–3989 (2016)
2. Armato III, S.G., et al.: Data from LIDC-IDRI. The cancer imaging archive, vol. 9, no. 7 (2015). <https://doi.org/10.7937/K9/TCIA.2015.LO9QL9SX>
3. Armato III, S.G., et al.: The lung image database consortium (LIDC) and image database resource initiative (IDRI): a completed reference database of lung nodules on CT scans. *Med. Phys.* **38**(2), 915–931 (2011)
4. Audelan, B., Delingette, H.: Unsupervised quality control of image segmentation based on Bayesian learning. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11765, pp. 21–29. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_3
5. Baumgartner, C.F., et al.: PHiSeg: capturing uncertainty in medical image segmentation. arXiv preprint [arXiv:1906.04045](https://arxiv.org/abs/1906.04045) (2019)
6. 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
7. Clark, K., et al.: The cancer imaging archive (TCIA): maintaining and operating a public information repository. *J. Digit. Imaging* **26**(6), 1045–1057 (2013)
8. Dehghani, M., Mehrjou, A., Gouws, S., Kamps, J., Schölkopf, B.: Fidelity-weighted learning. arXiv preprint [arXiv:1711.02799](https://arxiv.org/abs/1711.02799) (2017)
9. Han, X.: Automatic liver lesion segmentation using a deep convolutional neural network method. arXiv preprint [arXiv:1704.07239](https://arxiv.org/abs/1704.07239) (2017)
10. Huang, C., Han, H., Yao, Q., Zhu, S., Zhou, S.K.: 3D U²-net: a 3D universal U-net for multi-domain medical image segmentation. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11765, pp. 291–299. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_33
11. Jiang, L., Zhou, Z., Leung, T., Li, L.J., Fei-Fei, L.: Mentornet: learning data-driven curriculum for very deep neural networks on corrupted labels. In: *ICML*, pp. 2304–2313 (2018)
12. Kervadec, H., Dolz, J., Granger, É., Ben Ayed, I.: Curriculum semi-supervised segmentation. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11765, pp. 568–576. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_63
13. Kohl, S., et al.: A probabilistic u-net for segmentation of ambiguous images. In: *NeurIPS*, pp. 6965–6975 (2018)
14. Liu, H., Xu, J., Wu, Y., Guo, Q., Ibragimov, B., Xing, L.: Learning deconvolutional deep neural network for high resolution medical image reconstruction. *Inf. Sci.* **468**, 142–154 (2018)
15. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: fully convolutional neural networks for volumetric medical image segmentation. In: *2016 Fourth International Conference on 3D Vision (3DV)*, pp. 565–571. IEEE (2016)
16. Mondal, A.K., Dolz, J., Desrosiers, C.: Few-shot 3D multi-modal medical image segmentation using generative adversarial learning. arXiv preprint [arXiv:1810.12241](https://arxiv.org/abs/1810.12241) (2018)

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. Ren, M., Zeng, W., Yang, B., Urtasun, R.: Learning to reweight examples for robust deep learning. arXiv preprint [arXiv:1803.09050](https://arxiv.org/abs/1803.09050) (2018)
19. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
20. Shu, J., et al.: Meta-weight-net: learning an explicit mapping for sample weighting. arXiv preprint [arXiv:1902.07379](https://arxiv.org/abs/1902.07379) (2019)
21. Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A.A.: Inception-v4, inception-resnet and the impact of residual connections on learning. In: AAAI (2017)
22. van Tulder, G.: Package elsticdeform. <https://github.com/gvtulder/elasticdeform>
23. Yang, L., Zhang, Y., Chen, J., Zhang, S., Chen, D.Z.: Suggestive annotation: a deep active learning framework for biomedical image segmentation. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10435, pp. 399–407. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_46
24. 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
25. Zhao, Z., Yang, L., Zheng, H., Guldner, I.H., Zhang, S., Chen, D.Z.: Deep learning based instance segmentation in 3D biomedical images using weak annotation. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11073, pp. 352–360. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00937-3_41
26. Zhou, S., Wang, J., Zhang, M., Cai, Q., Gong, Y.: Correntropy-based level set method for medical image segmentation and bias correction. *Neurocomputing* **234**, 216–229 (2017)
27. Zhou, S., Wang, J., Zhang, S., Liang, Y., Gong, Y.: Active contour model based on local and global intensity information for medical image segmentation. *Neurocomputing* **186**, 107–118 (2016)