



Supapixel-Guided Label Softening for Medical Image Segmentation

Hang Li¹, Dong Wei², Shilei Cao², Kai Ma², Liansheng Wang^{1(✉)},
and Yefeng Zheng²

¹ Xiamen University, Xiamen, China

hangli@stu.xmu.edu.cn, lswang@xmu.edu.cn

² Tencent Jarvis Lab, Shenzhen, China

{donwei, eliaslcao, kylekma, yefengzheng}@tencent.com

Abstract. Segmentation of objects of interest is one of the central tasks in medical image analysis, which is indispensable for quantitative analysis. When developing machine-learning based methods for automated segmentation, manual annotations are usually used as the ground truth toward which the models learn to mimic. While the bulky parts of the segmentation targets are relatively easy to label, the peripheral areas are often difficult to handle due to ambiguous boundaries and the partial volume effect, *etc.*, and are likely to be labeled with uncertainty. This uncertainty in labeling may, in turn, result in unsatisfactory performance of the trained models. In this paper, we propose superpixel-based label softening to tackle the above issue. Generated by unsupervised over-segmentation, each superpixel is expected to represent a locally homogeneous area. If a superpixel intersects with the annotation boundary, we consider a high probability of uncertain labeling within this area. Driven by this intuition, we soften labels in this area based on signed distances to the annotation boundary and assign probability values within $[0, 1]$ to them, in comparison with the original “hard”, binary labels of either 0 or 1. The softened labels are then used to train the segmentation models together with the hard labels. Experimental results on a brain MRI dataset and an optical coherence tomography dataset demonstrate that this conceptually simple and implementation-wise easy method achieves overall superior segmentation performances to baseline and comparison methods for both 3D and 2D medical images.

Keywords: Soft labeling · Superpixel · Medical image segmentation

H. Li, D. Wei and S. Cao—Contributed equally.

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-59719-1_23) contains supplementary material, which is available to authorized users.

© Springer Nature Switzerland AG 2020

A. L. Martel et al. (Eds.): MICCAI 2020, LNCS 12264, pp. 227–237, 2020.

https://doi.org/10.1007/978-3-030-59719-1_23

1 Introduction

Segmentation of objects of interest is an important task in medical image analysis. Benefiting from the development of deep neural networks and the accumulation of annotated data, fully convolutional networks (FCNs) have demonstrated remarkable performances [8, 19] in this task. In general, these models assume that the ground truth is given precisely. However, for tasks with a large number of category labels, the peripheral areas are often difficult to annotate due to ambiguous boundaries and the partial volume effect (PVE) [2], *etc.*, and are likely to be labeled with uncertainty. With a limited number of data, FCNs may have difficulties in coping with such uncertainty, which in turn affects the performance. Taking brain MRI for example, in Fig. 1, we show a slice of a multi-sequence MRI, in which the pink area shows barely or non-discernible boundaries from its surroundings, causing great difficulties in the manual annotation.




Fig. 1. Illustration of ambiguous boundaries in medical images with a slice of a multi-sequence brain MRI. The first three images show the MRI sequences, and the last shows the ground truth annotation. As we can see, the boundaries of the tissue marked in pink are barely or even not discernible from its surroundings. Best viewed in color.

To reduce the impact of imprecise boundary annotation, a potential solution is the label softening technique, and at this moment, we are only aware of few of them [5, 10, 11]. Based on the anatomical knowledge that the lesion-surrounding pixels may also include some lesion level information, Kats *et al.* [11] employed 3D morphological dilation to expand the binary mask of multiple sclerosis (MS) lesions and assigned a fixed pseudo probability to all pixels within the expanded region, such that these pixels can also contribute to the learning of MS lesions. Despite the improved Dice similarity coefficient in the experiments, the inherent contextual information of images was not utilized when determining the extent of dilation or exact value of the fixed pseudo probability. To account for uncertainties in the ground truth segmentation of atherosclerotic plaque in the carotid artery, Engelen *et al.* [5] proposed to blur the ground truth mask with a Gaussian filter for label softening. One limitation of this work was that, similar to [11], the creation of the soft labels was only based on the ground truth while ignoring the descriptive contextual information in the image. From another perspective, soft labels can also be obtained by fusing multiple manual annotations, e.g., in

[10] masks of MS lesions produced by different experts were fused using a soft version of the STAPLE algorithm [22]. However, obtaining multiple segmentation annotations for medical images can be practically difficult. An alternative to label softening is the label smoothing technique [16, 20] which assumes a uniform prior distribution over labels; yet again, this technique did not take the image context into consideration, either.




Fig. 2. Pipeline of our proposed method.

In this paper, we propose a new label softening method driven by the image contextual information, for improving segmentation performance especially near the boundaries of different categories. Specifically, we employ the concept of superpixels [1] for the utilization of local contextual information. Via unsupervised over-segmentation, the superpixels group original image pixels into locally homogeneous blocks, which can be considered as meaningful atomic regions of the image. Conceptually, if the scale of superpixel is appropriate, pixels within the same superpixel block can be assumed belonging to the same category. Based on this assumption, if a superpixel intersects with the annotation boundary of the ground truth, we consider a high probability of uncertain labeling within the area prescribed by this superpixel. Driven by this intuition, we soften labels in this area based on the signed distance to the annotation boundary, producing probability values spanning the full range of $[0, 1]$ —in contrast to the original “hard” binary labels of either 0 or 1. The softened labels can then be used to train the segmentation models. We evaluate the proposed approach on two

publicly available datasets: the Grand Challenge on MR Brain Segmentation at MICCAI 2018 (MRBrainS18) [7] dataset and an optical coherence tomography (OCT) image [3] dataset. The experimental results verify the effectiveness of our approach.

2 Method

The pipeline of our method is illustrated in Fig. 2. We employ the SLIC algorithm [1] to produce superpixels, meanwhile converting the ground truth annotation to multiple one-hot label maps (the “hard” labels). Soft labels are obtained by exploiting the relations between the superpixels and hard label maps (the cross symbol \otimes in Fig. 2). Then, the soft and hard labels are used jointly to supervise the training of the segmentation network.

Superpixel-Guided Region of Softening. Our purpose is to model the uncertainty near the boundaries of categories in the manual annotation for improving model performance and robustness. For this purpose, we propose to exploit the relations between superpixels and the ground truth annotation to produce soft labels. Specifically, we identify three types of relations between a superpixel and the foreground region in a one-hot ground truth label map (Fig. 3): (a) the superpixel is inside the region, (b) the superpixel is outside the region, and (c) the superpixel intersects with the region boundary. As the superpixel algorithms [1] group pixels into locally homogeneous pixel blocks, pixels within the same superpixel can be assumed to belong to the same category given that superpixels are set to a proper size. Based on this assumption, it is most likely for uncertain annotations to happen in the last case, where the ground truth annotation indicates different labels for pixels inside the same superpixel block. Therefore, our label softening works exclusively in this case.

Formally, let us denote an image by $x \in \mathbb{R}^{W \times H}$, where W and H are the width and height, respectively. (Without loss of generalization, x can also be a 3D image $x \in \mathbb{R}^{W \times H \times T}$, where T is the number of slices, and our method still applies.) Then, its corresponding ground truth annotation can be denoted by a set of one-hot label maps: $Y = \{y^c | y^c \in \mathbb{R}^{W \times H}\}_{c=1}^C$, where C is the number of categories, and y^c is the binary label map for category c , in which any pixel $y_i^c \in \{0, 1\}$, where $i \in \{1, \dots, N\}$ is the pixel index, and N is the total number of pixels; besides, we denote the foreground area in y^c by ϕ^c . We can generate superpixel blocks $S(x) = \{s^{(j)}\}_{j=1}^M$ for x using an over-segmentation algorithm, where M is the total number of superpixels. In this paper, we adopt SLIC [1] as our superpixel-generating algorithm, which is known for computational efficiency and quality of the generated superpixels. We denote the set of soft label maps to be generated by $Q_c = \{q^c | q^c \in \mathbb{R}^{W \times H}\}$; note that $q_i^c \in [0, 1]$ is a continuous value, in contrast with the binaries in y^c . As shown in Fig. 3, the relations between any ϕ^c and $s^{(j)}$ can be classified into three categories: (a) $s^{(j)}$ is inside ϕ^c ; (b) $s^{(j)}$ is outside ϕ^c ; and (c) $s^{(j)}$ intersects with boundaries of ϕ^c . For the first two cases, we use the original values of y_i^c in the corresponding locations in

q^c . Whereas as for the third case, we employ label softening strategies to assign a soft label q_i^c to each pixel i based on its distance to boundaries of ϕ^c , which is described below.




Fig. 3. Illustration of three types of relations between the foreground region in a binary ground truth label map (GT) and a superpixel block (SP). (a) SP is inside GT; (b) SP is outside GT; (c) SP intersects with boundaries of GT. We identify the region enclosed by the SP in the third case for label softening, based on the signed distances to GT boundaries.

Soft Labeling with Signed Distance Function. Assume a superpixel block s intersects with the boundaries of a foreground ϕ (for simplicity, the superscripts can be safely omitted here without confusion). For a pixel s_i in s , the absolute value of the distance d_i from s_i to ϕ is defined as the minimum distance among all the distances from s_i to all pixels on the boundaries of ϕ . We define $d_i > 0$ if s_i is inside ϕ , and $d_i \leq 0$ otherwise. As aforementioned, in the case of a superpixel block intersecting with the boundaries of ϕ , we need to assign each pixel in this block a pseudo-probability as its soft label according to its distance to ϕ . The pseudo-probability should be set to 0.5 for a pixel right on the boundary (i.e. $d_i = 0$), gradually approach 1 as d_i increases, and gradually approach 0 otherwise. Thus, we define the distance-to-probability conversion function as

$$q_i = f_{\text{dist}}(d_i) = \frac{1}{2} \left(\frac{d_i}{1 + |d_i|} + 1 \right), \quad (1)$$

where $q_i \in [0, 1]$ is the obtained soft label for pixel i .

Model Training with Soft and Hard Labels. We adopt the Kullback-Leibler (KL) divergence loss [13] to supervise model training with our soft labels:

$$\mathcal{L}_{\text{KL}} = \frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C q_i^c \log(q_i^c/p_i^c), \quad (2)$$

where p_i^c is the predicted probability of the i -th pixel belonging to the class c , and q_i^c is the corresponding soft label defined with Eq. (1). Along with \mathcal{L}_{KL} , we also adopt the commonly used Dice loss $\mathcal{L}_{\text{Dice}}$ [15] and cross-entropy (CE) loss \mathcal{L}_{CE} for medical image segmentation. Specifically, the CE loss is defined as:

$$\mathcal{L}_{\text{CE}} = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C w_c y_i^c \log(p_i^c), \quad (3)$$

where w_c is the weight for class c . When $w_c = 1$ for all classes, Eq. (3) is the standard CE loss. In addition, w_c can be set to class-specific weights to counteract the impact of class imbalance [17]: $w_c = 1/\log(1.02 + \sum_{i=1}^N y_i^c/N)$, and we refer to this version of the CE loss as weighted CE (WCE) loss. The final loss is defined as a weighted sum of the three losses: $\mathcal{L} = \mathcal{L}_{\text{CE}} + \alpha\mathcal{L}_{\text{Dice}} + \beta\mathcal{L}_{\text{KL}}$, where α and β are hyperparameters to balance the three losses. We follow the setting in nnU-Net [8] to set $\alpha = 1.0$, and explore the proper value of β in our experiments, since it controls the relative contribution of our newly proposed soft labels which are of interest.

3 Experiments

Datasets. To verify the effectiveness of our method on both 2D and 3D medical image segmentation, we use datasets of both types for experiments. The MRBrainS18 dataset [7] provides seven 3T multi-sequence (T1-weighted, T1-weighted inversion recovery, and T2-FLAIR) brain MRI scans with the following 11 ground truth labels: 0-background, 1-cortical gray matter, 2-basal ganglia, 3-white matter, 4-white matter lesions, 5-cerebrospinal fluid in the extracerebral space, 6-ventricles, 7-cerebellum, 8-brain stem, 9-infarction and 10-other, among which labels 9 and 10 were officially excluded from the evaluation and we follow this setting. We randomly choose five scans for training and use the rest for evaluation. For preprocessing, the scans are preprocessed by skull stripping, nonzero cropping, resampling, and data normalization. The other dataset [3] includes OCT images with diabetic macular edema (the OCT-DME dataset) for the segmentation of retinal layers and fluid regions. It contains 110 2D B-scan images from 10 patients. Eight retinal layers and fluid regions are annotated. We use the first five subjects for training and the last five subjects for evaluation (each set has 55 B-scans). Since the image quality of this dataset is poor, we firstly employ a denoising convolutional neural networks (DnCNN) [23] to reduce image noise and improve the visibility of anatomical structures. To reduce memory usage, we follow He *et al.* [6] to flatten a retinal B-scan image to the estimated Bruch’s membrane (BM) using an intensity gradient method [14] and crop the retina part out.

Experimental Setting and Implementation. For the experiments on each dataset, we first establish a baseline, which is trained without the soft labels. Then, we re-implement the Gaussian blur based label softening method [5], in which the value of σ is empirically selected, for a comparison with our proposed method. Considering the class imbalance in both datasets, we present results using the standard CE and WCE losses for all methods. We notice that the Dice loss adversely affects the performance on the OCT-DME dataset, therefore those results are not reported. We use overlap-based, volume-based, and distance-based mean metrics [21], including: Dice coefficient score, volumetric similarity (VS), 95th percentile Hausdorff distance (HD95), average surface distance (ASD), and average symmetric surface distance (ASSD) for a comprehensive evaluation of the methods. We employ a 2D U-Net [19] segmentation model

(with the Xception [4] encoder) for the OCT-DME dataset, and a 3D U-Net [8] model for the MRBrainS18 dataset (patch-based training and sliding window test tricks [9] are employed in the implementation). All experiments are conducted with the PyTorch framework [18] on a standard PC with an NVIDIA GTX 1080Ti GPU. The Adam optimizer [12] is adopted with a learning rate of 3×10^{-4} and a weight decay of 10^{-5} . The learning rate is halved if the validation performance does not improve for 20 consecutive epochs. The batch size is fixed to 2 for the MRBrainS18 dataset, and 16 for the OCT-DME dataset.

Results. The quantitative evaluation results are summarized in Table 1 and Table 2 for the MRBrainS18 and OCT-DME datasets, respectively. (Example segmentation results on both datasets are provided in the supplementary material.) As expected, the weighted CE loss produces better results than the standard CE loss for most evaluation metrics on both datasets. We note that the Gaussian blur based label softening [5] does not improve upon the baselines either with the CE or WCE loss, but only obtains results comparable to those of the baselines. The reason might be that this method indiscriminately softens all boundary-surrounding pixels with a fixed standard deviation without considering the actual image context, which may potentially harm the segmentation near originally precisely annotated boundaries. In contrast, our proposed method consistently improves all metrics when using the generated soft labels with the WCE loss. In fact, with this combination of losses, our method achieves the best performances for all evaluation metrics. It is also worth mentioning that, although our method is motivated by improving segmentation near category boundaries, it also improves the overlap-based evaluation metrics (Dice) by a noticeable extent on the OCT-DME dataset. These results verify the effectiveness of our method in improving segmentation performance, by modeling uncertainty in manual labeling with the interaction between superpixels and ground truth annotations.

Table 1. Evaluation results on the MRBrainS18 dataset [7]. The KL divergence loss is used by our method for model training with our soft labels.


Method	Losses	Dice (%) \uparrow	VS (%) \uparrow	HD95 (mm) \downarrow	ASD (mm) \downarrow	ASSD (mm) \downarrow
Baseline	CE+Dice	85.47	96.53	3.5287	0.9290	0.7722
	WCE+Dice	85.56	96.55	3.5116	0.8554	0.7371
Engelen <i>et al.</i> [5]	CE+Dice	85.16	95.44	3.6396	0.9551	0.8004
	WCE+Dice	85.47	95.58	3.5720	0.8432	0.7502
Ours	KL+CE+Dice	85.26	93.95	3.4684	0.9460	0.8061
	KL+WCE+Dice	85.63	96.60	3.1194	0.8146	0.7153

Ablation Study on Number of Superpixels. The proper scale of the superpixels is crucial for our proposed method, as superpixels of different sizes may

Table 2. Evaluation results on the OCT-DME dataset [3]. The KL divergence loss is used by our method for model training with our soft labels.

Method	Losses	Dice (%) \uparrow	VS (%) \uparrow	HD95 (mm) \downarrow	ASD (mm) \downarrow	ASSD (mm) \downarrow
Baseline	CE	82.50	96.52	3.22	1.075	1.075
	WCE	82.82	96.44	3.27	1.082	1.087
Engelen <i>et al.</i> [5]	CE	82.69	96.38	3.24	1.092	1.092
	WCE	82.94	96.36	3.27	1.080	1.094
Ours	KL+CE	82.72	96.41	3.25	1.081	1.087
	KL+WCE	83.94	96.73	3.17	1.058	1.066

describe different levels of image characteristics, and thus may interact differently with the ground truth annotation. Since in the SLIC [1] algorithm, the size of superpixels is controlled by the total number of generated superpixel blocks, we conduct experiments to study how the number of superpixels influences the performance on the MRBrainS18 dataset. In Fig. 4, we show performances of our method with different numbers of superpixels ranged from 500 to 3500 with a sampling interval of 500. As we can see, as the number of superpixels increases, the performance first increases due to the more image details incorporated, and then decreases after reaching the peak. This is in line with our intuition, since the assumption that pixels within the same superpixel belong to the same category can hold only if the scale of superpixels is appropriate. Large superpixels can produce flawed soft labels. In contrast, as the number of superpixels grows and their sizes shrink, soft labels will degenerate into hard labels, which does not provide additional information.

**Fig. 4.** Performances of our method with different numbers of superpixels on the MRBrainS18 dataset [7]. The HD95, ASD and ASSD are in mm. Best viewed in color.

Ablation Study on Weight of Soft Label Loss. The weight β controls the contribution of the soft labels in training. To explore the influence of the soft label loss, we conduct a study on the MRBrainS18 dataset to compare the performance of our method with different values of β . We set β to $1/4$, $1/2$, 1 , 2 , 4 , and 8 . The mean Dice, HD95, ASD, and ASSD of our proposed method with




Fig. 5. Performances of our method using different values of β on the MRBrainS18 dataset [7]. The HD95, ASD and ASSD are in mm. Best viewed in color.

these values of β are shown in Fig. 5. Note that the x -axis uses a log scale since values of β differ by orders of magnitude. Improvements in performance can be observed when β increases from $1/4$ to 1 . When β continues to increase, however, the segmentation performances start to drop. This indicates that the soft labels are helpful to segmentation, although giving too much emphasis to them may decrease the generalization ability of the segmentation model.

4 Conclusion

In this paper, we presented a new label softening method that was simple yet effective in improving segmentation performance, especially near the boundaries of different categories. The proposed method first employed an over-segmentation algorithm to group image pixels into locally homogeneous blocks called superpixels. Then, the superpixel blocks intersecting with the category boundaries in the ground truth were identified for label softening, and a signed distance function was employed to convert the pixel-to-boundary distances to soft labels within $[0, 1]$ for pixels inside these blocks. The soft labels were subsequently used to train a segmentation network. Experimental results on both 2D and 3D medical images demonstrated the effectiveness of this simple approach in improving segmentation performance.

Acknowledgments. This work was supported by the National Natural Science Foundation of China (Grant No. 61671399), Fundamental Research Funds for the Central Universities (Grant No. 20720190012), Key Area Research and Development Program of Guangdong Province, China (No. 2018B010111001), National Key Research and Development Project (2018YFC2000702), and Science and Technology Program of Shenzhen, China (No. ZDSYS201802021814180).

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 Trans. Pattern Anal. Mach. Intell.* **34**(11), 2274–2282 (2012)
2. Ballester, M.A.G., Zisserman, A.P., Brady, M.: Estimation of the partial volume effect in MRI. *Med. Image Anal.* **6**(4), 389–405 (2002)

3. Chiu, S.J., Allingham, M.J., Mettu, P.S., Cousins, S.W., Izatt, J.A., Farsiu, S.: Kernel regression based segmentation of optical coherence tomography images with diabetic macular edema. *Biomed. Opt. Express* **6**(4), 1172–1194 (2015)
4. Chollet, F.: Xception: deep learning with depthwise separable convolutions. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1251–1258 (2017)
5. van Engelen, A., et al.: Supervised in-vivo plaque characterization incorporating class label uncertainty. In: *IEEE International Symposium on Biomedical Imaging*, pp. 246–249. IEEE (2012)
6. He, Y., et al.: Fully convolutional boundary regression for retina OCT segmentation. In: Shen, D., et al. (eds.) *MICCAI 2019. LNCS*, vol. 11764, pp. 120–128. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32239-7_14
7. Kuijf, H.J.: Grand challenge on MR brain segmentation at MICCAI 2018 (2018). <http://mrbrains18.isi.uu.nl>
8. Isensee, F., et al.: Abstract: nnU-Net: self-adapting framework for U-Net-based medical image segmentation. *Bildverarbeitung für die Medizin 2019. I*, pp. 22–22. Springer, Wiesbaden (2019). https://doi.org/10.1007/978-3-658-25326-4_7
9. Jin, D., et al.: Accurate esophageal gross tumor volume segmentation in PET/CT using two-stream chained 3D deep network fusion. In: Shen, D., et al. (eds.) *MICCAI 2019. LNCS*, vol. 11765, pp. 182–191. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_21
10. Kats, E., Goldberger, J., Greenspan, H.: A soft STAPLE algorithm combined with anatomical knowledge. In: Shen, D., et al. (eds.) *MICCAI 2019. LNCS*, vol. 11766, pp. 510–517. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32248-9_57
11. Kats, E., Goldberger, J., Greenspan, H.: Soft labeling by distilling anatomical knowledge for improved MS lesion segmentation. In: *IEEE International Symposium on Biomedical Imaging*, pp. 1563–1566. IEEE (2019)
12. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. In: *International Conference on Learning Representations* (2015)
13. Kullback, S., Leibler, R.A.: On information and sufficiency. *Ann. Math. Stat.* **22**(1), 79–86 (1951)
14. Lang, A., et al.: Retinal layer segmentation of macular OCT images using boundary classification. *Biomed. Opt. Express* **4**(7), 1133–1152 (2013)
15. Milletari, F., Navab, N., Ahmadi, S.A.: V-Net: fully convolutional neural networks for volumetric medical image segmentation. In: *International Conference on 3D Vision*, pp. 565–571. IEEE (2016)
16. Ouyang, X., et al.: Weakly supervised segmentation framework with uncertainty: a study on pneumothorax segmentation in chest x-ray. In: Shen, D., et al. (eds.) *MICCAI 2019. LNCS*, vol. 11769, pp. 613–621. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32226-7_68
17. Paszke, A., Chaurasia, A., Kim, S., Culurciello, E.: ENet: a deep neural network architecture for real-time semantic segmentation. *arXiv preprint arXiv:1606.02147* (2016)
18. Paszke, A., et al.: PyTorch: an imperative style, high-performance deep learning library. In: *Advances in Neural Information Processing Systems*, pp. 8024–8035 (2019)
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. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the Inception architecture for computer vision. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2818–2826 (2016)
21. Taha, A.A., Hanbury, A.: Metrics for evaluating 3D medical image segmentation: analysis, selection, and tool. *BMC Med. Imaging* **15**(1), 29 (2015)
22. Warfield, S.K., Zou, K.H., Wells, W.M.: Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. *IEEE Trans. Med. Imaging* **23**(7), 903–921 (2004)
23. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a Gaussian denoiser: residual learning of deep CNN for image denoising. *IEEE Trans. Image Process.* **26**(7), 3142–3155 (2017)