



# Superpixel-Guided Iterative Learning from Noisy Labels for Medical Image Segmentation

Shuailin Li<sup>1</sup>(✉), Zhitong Gao<sup>1</sup>, and Xuming He<sup>1,2</sup>

<sup>1</sup> ShanghaiTech University, Shanghai, China  
{lishl, gaozht, hexm}@shanghaitech.edu.cn

<sup>2</sup> Shanghai Engineering Research Center of Intelligent Vision and Imaging, Shanghai, China

**Abstract.** Learning segmentation from noisy labels is an important task for medical image analysis due to the difficulty in acquiring high-quality annotations. Most existing methods neglect the pixel correlation and structural prior in segmentation, often producing noisy predictions around object boundaries. To address this, we adopt a superpixel representation and develop a robust iterative learning strategy that combines noise-aware training of segmentation network and noisy label refinement, both guided by the superpixels. This design enables us to exploit the structural constraints in segmentation labels and effectively mitigate the impact of label noise in learning. Experiments on two benchmarks show that our method outperforms recent state-of-the-art approaches, and achieves superior robustness in a wide range of label noises. Code is available at [https://github.com/gaozhitong/SP\\_guided\\_Noisy\\_Label\\_Seg](https://github.com/gaozhitong/SP_guided_Noisy_Label_Seg).

**Keywords:** Learning with noisy labels · Semantic segmentation

## 1 Introduction

Semantic segmentation of medical images, a fundamental task in computer-aided clinical diagnoses, has recently achieved remarkable progress thanks to the effective feature learning based on deep neural networks (DNN) [3, 13, 19]. Training such DNNs for segmentation typically requires a large dataset with pixelwise annotations that accurately delineate object boundaries. In the medical domain, however, acquiring such high-quality annotations is often difficult due to lack

---

S. Li and Z. Gao—Equal contribution.

This work was supported by Shanghai Science and Technology Program 21010502700 and by the ShanghaiTech-UII Joint Lab.

---

**Electronic supplementary material** The online version of this chapter ([https://doi.org/10.1007/978-3-030-87193-2\\_50](https://doi.org/10.1007/978-3-030-87193-2_50)) contains supplementary material, which is available to authorized users.

of experienced annotators and/or visual ambiguity in object boundaries [6, 11]. Consequently, the annotated datasets often include a varying amount of label noise in practice, ranging from small boundary offsets to large region errors. Learning from those noisy annotations has been particularly challenging for deep segmentation networks due to the memorization effect [2].

There have been several attempts to tackle the problem of training segmentation networks from noisy labels, which can be largely grouped into two categories. The first type of methods view the annotation of each image as either clean or corrupted, and iteratively select or reweight image samples during training [23, 26]. In particular, Zhu et al. [26] implicitly reweight image losses by simultaneously training a label evaluation network and a segmentation network, while Xue et al. [23] explicitly select a subset of images by extending the Co-teaching [8] scheme into a tri-network framework. Such image-level weighting strategies, however, are less robust under severe noise settings as they are unable to fully utilize the pixels with clean annotations in each image. To address this limitation, the second group of training methods consider the segmentation as a pixel-wise classification task [24, 25], and perform pixel-wise sample selection or label refinement based on the state-of-the-art robust classifier learning strategies, such as the confidence learning technique [16] and the Co-teaching method with tri-networks [25]. Despite their better use of annotations, the pixel-level approaches ignore the pixel correlation and spatial prior in image segmentation, and hence tend to produce noisy prediction around object boundaries.

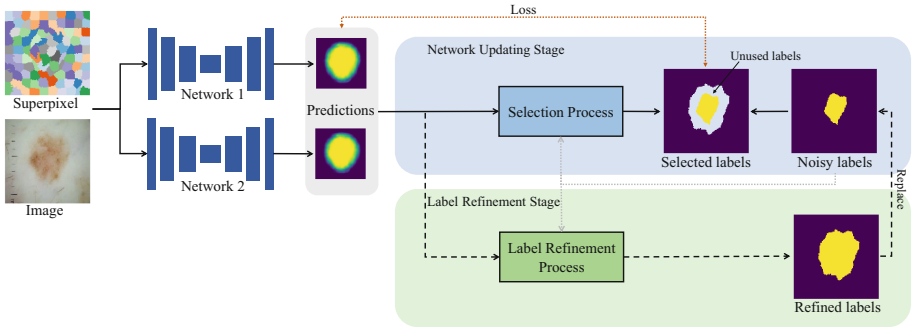
In this work, we propose a novel robust learning strategy for semantic image segmentation, aiming to exploit the structural prior of images and correlation in pixel labels. To this end, we adopt a superpixel representation and develop an iterative learning scheme that combines noise-aware training of segmentation network and noisy label refinement, both guided by the superpixels. Such integration allows us to better utilize the structural constraint in segmentation labels for model learning, which can effectively mitigate the impact of label noise. We note that while superpixel has been employed in recent work [12], they only use it to correct noisy labels and ignore the impact of noise during training.

Specifically, in each iteration, we first jointly train two deep networks using selected subsets of superpixels with small loss values, following the multi-view learning framework [8, 22]. As in the Co-teaching method, such a multi-view learning strategy regularizes the network training via the predictions of the peer networks. Here we treat each superpixel as a data sample in selection, which enables us to enforce spatial smoothness and provide better object boundary cues in network training. To avoid overfitting to label noise, we design an automatic stopping criterion for the joint learning based on the loss statistics of superpixels. After the network training, we then use the network predictions to estimate the reliability of superpixel labels and relabel a subset of most unreliable ones. Such label refinement allows us to improve the label quality for the subsequent model training. The network and label updates are repeated until no further improvement can be achieved for the label refinement.

We evaluate our method on two public benchmarks, ISIC skin lesion dataset [7] and JSRT chest x-ray dataset [5, 20], under extensive noise settings. Empirical results show that our method consistently outperforms the previous state of the art and demonstrates training robustness in a wide range of label noises.

## 2 Method

We now introduce our robust learning strategy for semantic segmentation, which aims to exploit the structural constraints in the label masks and to fully utilize reliable pixel-level labels for effective learning. To achieve this, we adopt a superpixel-based data representation, and develop an iterative learning method that jointly optimizes the network parameters and refines noisy labels.



**Fig. 1.** Overview of our robust training process. We use superpixels as our guidance in an iterative learning process which jointly updates network parameters and refines noisy labels. Each iteration selects superpixels with small losses to update two networks and relabels a set of superpixels based on network outputs.

Specifically, given a target network and noisy training data, we first compute the superpixels of the input images. Based on such pixel groupings, our iterative learning procedure alternates between a noise-aware network training stage and a label refinement stage until no improvement can be achieved. For the network training stage, we adopt the multi-view learning framework which jointly trains two instances of the segmentation network. For the label refinement stage, we use the output of two trained networks to estimate the reliability of superpixel labels and to update the unreliable labels. An overview of our training pipeline is shown in Fig. 1. Below we first present our superpixelization procedure in Sect. 2.1, followed by the two stages in the iterative learning in Sect. 2.2.

### 2.1 Superpixel Representation

To exploit the image structural prior and spatial correlation in pixel labels, we first compute a superpixel representation for the training images. Such superpixel

representation has been shown effective for different medical image modalities in literature, e.g. [17] for CT, [21] for MR and [4] for US images. Specifically, we use the off-the-shelf superpixelization method, SLIC [1], to partition each image into a set of homogeneous regions. For color images, we adopt the CIE-lab color space to represent pixel features, while for other modalities, such as X-ray images, we use both the pixel intensity and deep features from a U-net trained with a noise-aware method, e.g. [22].<sup>1</sup>

We assume the pixels share similar groundtruth labels in each superpixel, which enable us to enforce the structural constraints on the label masks and better preserve object boundaries. More importantly, we treat each superpixel as a data sample in the subsequent robust network learning as well as the label refinement. This allows us to estimate the noise level of pixel annotations in a more reliable manner by pooling the pixel information from each superpixel.

## 2.2 Iterative Model Learning

We now present our iterative learning strategy based on the superpixel representation, which aims to fully utilize the clean pixel annotations and meanwhile reduce the impact of noisy labels. To this end, we introduce an iterative optimization process for model training as below. Each iteration consists of two stages: a noise-aware network learning stage to update the network parameters and a label refinement stage to correct unreliable annotations.

**Network Update.** In the first stage, we perform a noise-aware network learning by incorporating superpixel representation into a multi-view learning framework. Specifically, following the Co-teaching strategy [8, 10, 18], we jointly train two instances of the target segmentation network using partial data with small losses [2]. To better select data samples with clean labels, we design a superpixel-wise loss that combines the loss values of two networks with an agreement-based regularization [22] on superpixels. Our loss provides a reliable guidance to sample selection thanks to the structural prior encoded in the superpixels.

Formally, given an image  $\mathbf{X}$ , we denote its annotation as  $\mathbf{Y} = \{Y_i\}_{i=1}^M, Y_i \in \{1, \dots, C\}$  where  $C$  is the number of semantic classes and  $M$  is the number of pixels. The superpixel map is represented by  $\mathbf{S} = \{S_i\}_{i=1}^M$  where  $S_i \in \{1, 2, \dots, K\}$  and  $K$  is the number of superpixels. Here  $S_j = k$  means that pixel  $j$  belongs to superpixel  $k$ .

We aim to train two deep neural networks denoted by  $f(\cdot, \theta_1)$  and  $f(\cdot, \theta_2)$ . To define a loss for each image, we first generate the predicted probability maps from two networks, denoted by  $\mathbf{P}^1, \mathbf{P}^2 \in \mathbb{R}^{C \times M}$ , where  $\mathbf{P}^i = f(\mathbf{X}, \theta_i), i = 1, 2$ . We then compute the superpixel-wise probabilities  $\mathbf{P}_s^i \in \mathbb{R}^{C \times K}, i = 1, 2$  and the corresponding soft labels  $\mathbf{Y}_s \in [0, 1]^{C \times K}$  by averaging over each superpixel:

---

<sup>1</sup> On the chest x-ray dataset [5, 20], our superpixels achieve undersegmentation errors lower than 0.32 for 800 superpixels per image, which is comparable to the natural image setting[1].

$$\mathbf{P}_s^i(c, k) = \frac{1}{N(k)} \sum_{j: S_j=k} \mathbf{P}^i(c, j), \quad \mathbf{Y}_s(c, k) = \frac{1}{N(k)} \sum_{j: S_j=k} \mathbf{1}(Y_j = c) \quad (1)$$

where  $N(k) = |\{j : S_j = k\}|$  is the size of the superpixel. Inspired by [22], we define our superpixel-wise loss function  $\ell^{sp}$  by considering both classification losses and prediction agreement on each superpixel:

$$\ell^{sp} = (1 - \lambda) * (\ell_{ce}(\mathbf{P}_s^1, \mathbf{Y}_s) + \ell_{ce}(\mathbf{P}_s^2, \mathbf{Y}_s)) + \lambda * \ell_{kl}(\mathbf{P}_s^1, \mathbf{P}_s^2) \quad (2)$$

where  $\ell_{ce}$  is the cross-entropy loss with soft labels,  $\ell_{kl}$  is the symmetric Kullback-Leibler(KL) Divergence, and  $\lambda$  is a balance factor. By considering both two terms in the small loss criterion, we aim to select and update on training data with low label noise while maximizing networks' agreement. Denote  $R$  as the ratio of pixels being selected, we perform the small-loss selection by choosing the superpixel set  $\hat{\mathcal{D}}_s$  as follows:

$$\hat{\mathcal{D}}_s = \arg \min_{\mathcal{D}_s: N(\mathcal{D}_s) \geq R \cdot M} \sum_{k \in \mathcal{D}_s} \ell_k^{sp} \quad (3)$$

where  $N(\mathcal{D}_s) = \sum_{k \in \mathcal{D}_s} N(k)$  is the total number of pixels in the superpixel set. Given the small-loss selection, we train the network based on the average loss:

$$\mathcal{L} = \frac{1}{N(\hat{\mathcal{D}}_s)} \sum_{S_i \in \hat{\mathcal{D}}_s} \ell_i \quad (4)$$

where  $\ell$  has the same form as Eq. 2 except it is defined on the pixel level. Here we skip the superpixel-level pooling as in Eq. 1 for more efficient back propagation.

**Stopping Criterion.** While the selection strategy enables the network training with mostly clean-labeled data, some noisy labels are inevitably selected and gradually affect model performance. To tackle the problem, we propose a criterion to stop network training before such overfitting. Our criterion is defined based on the loss gap  $G_l$  between the selected data and the rest of the training set as follows:

$$G_l = \frac{1}{K - |\hat{\mathcal{D}}_s|} \sum_{k \notin \hat{\mathcal{D}}_s} \ell_k^{sp} - \frac{1}{|\hat{\mathcal{D}}_s|} \sum_{k \in \hat{\mathcal{D}}_s} \ell_k^{sp} \quad (5)$$

Intuitively, the model tends to first learn the relatively simple patterns in clean data, then starts to overfit to label noise [2]. Consequently,  $G_l$  first gradually increases, then starts to decrease. Based on this observation, we stop the model training when  $G_l$  reached the maximum before decreasing.

After training, we observe that the outputs of two peer networks are very similar to each other. Consequently, we arbitrarily choose one network to make predictions during test/deployment.

**Label Refinement.** In this stage, we use the trained networks to estimate the reliability of the superpixel annotations and relabel a subset of unreliable ones. Specifically, we choose superpixels with large losses, which indicates strong inconsistency between model predictions and their labels. We then relabel them according to predicted class labels. Formally, we define the unreliable superpixel set  $\hat{\mathcal{D}}_u$  based on the superpixel losses and compute the predicted superpixel labels  $\hat{\mathbf{Y}} = \{\hat{Y}_i\}_{i=1}^K, \hat{Y}_i \in \{1, \dots, C\}$  as below,

$$\hat{\mathcal{D}}_u = \arg \max_{\mathcal{D}_u: N(\mathcal{D}_u) \leq (1-R) \cdot M} \sum_{k \in \mathcal{D}_u} \ell_k^{sp} \quad (6)$$

$$\hat{Y}_k = \arg \max_c \frac{1}{2} (\mathbf{P}_s^1(c, k) + \mathbf{P}_s^2(c, k)) \quad (7)$$

where  $\ell_k^{sp}$  is defined in Eq. 2,  $R$  is the selection ratio mentioned before. Finally, we update the pixel-wise label map  $\mathbf{Y}' = \{Y'_i\}_{i=1}^M, Y'_i \in \{1, \dots, C\}$  as

$$Y'_i = \mathbb{1}(S_i \in \hat{\mathcal{D}}_u) \hat{Y}_{S_i} + \mathbb{1}(S_i \notin \hat{\mathcal{D}}_u) Y_i \quad (8)$$

After the label refinement, we replace  $\mathbf{Y}$  by  $\mathbf{Y}'$ , increase  $R$  by a fixed ratio  $\gamma$  and start the next iteration.

### 3 Experiment

We validate our method on two public datasets, ISIC [7] and JSRT [5, 20], which consists of images from two different modalities. We follow the literature and use simulated label noises as no public benchmark with real label noises is available.

#### 3.1 Dataset

**ISIC Dataset.** ISIC 2017 dataset [7] is a public large-scale dataset of dermoscopy images, acquired from a variety of devices used at multiple sites. This dataset contains 2000 training and 600 test images with corresponding segmentation masks. We resize all images to  $128 \times 128$  in resolution.

**JSRT Dataset.** JSRT dataset [5, 20] is a public chest x-ray dataset containing three classes of annotations: lung, heart and clavicle. There are 247 chest radiographs in total, with unified resolution  $2048 \times 2048$ . We split them into a training set of 197 images and a test set of 50 images, and resize them into  $256 \times 256$ <sup>2</sup>.

**Noise Patterns.** To simulate manual noisy annotations, we randomly select a ratio  $\alpha$  of samples from the training data to apply morphological [23–26] or affine transformation with noise level controlled by  $\beta$ . For affine transformation, we use a combination of rotation and translation to imitate other real-world noise patterns. Unlike prior works, we use the relative size w.r.t the target object region

<sup>2</sup> Clavicles are particularly small in chest x-ray images. To facilitate fine-grained segmentation and reduce consuming time, we crop their region of interest by statistics on the training set.

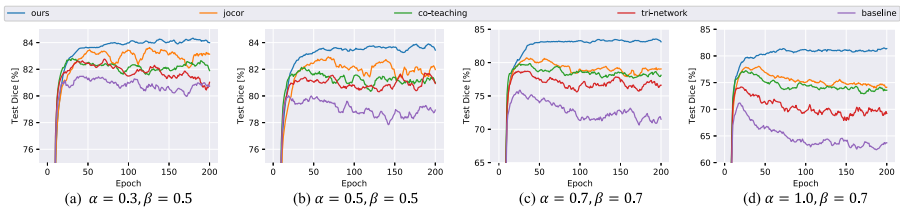
when controlling the noise level  $\beta$ , as people usually annotate target object in a favorable field of view by zooming in or out images. We investigate our algorithm in several noisy settings with  $\alpha$  being  $\{0.3, 0.5, 0.7, 1.0\}$  and  $\beta$  being  $\{0.5, 0.7\}$ . Some noisy examples are shown in the supplementary.

### 3.2 Experiment Setup

**Comparisons.** We compare our method with several state-of-the-art approaches, including Co-teaching [8], Tri-network [25] and JoCoR [22], which employ the robust learning at the pixel level. We do not include methods such as [15, 24] as they rely on a clean validation set. For fair comparison, we re-implement these methods with the same network backbone and training policy.

**Table 1.** Quantitative comparisons of noisy-labeled segmentation methods on ISIC dataset, where the metric is Dice[%] over the last 10 epochs.  $\alpha$  and  $\beta$  control the noise ratio and noise level, respectively.

	Baseline	Co-teaching[8]	Tri-network[25]	JoCoR [22]	Ours
Original data	82.49	82.72	82.96	83.64	<b>84.26</b>
$\alpha = 0.3, \beta = 0.5$	80.75	81.44	81.50	82.65	<b>84.00</b>
$\alpha = 0.3, \beta = 0.7$	79.46	81.47	80.73	81.58	<b>83.34</b>
$\alpha = 0.5, \beta = 0.5$	78.95	81.22	80.94	82.41	<b>83.90</b>
$\alpha = 0.5, \beta = 0.7$	75.44	80.06	80.24	81.06	<b>83.19</b>
$\alpha = 0.7, \beta = 0.5$	76.61	79.61	79.55	80.55	<b>83.83</b>
$\alpha = 0.7, \beta = 0.7$	71.51	78.50	76.61	79.05	<b>83.12</b>
$\alpha = 1.0, \beta = 0.5$	71.13	76.69	75.61	78.43	<b>82.23</b>
$\alpha = 1.0, \beta = 0.7$	63.71	73.68	70.01	74.30	<b>81.39</b>



**Fig. 2.** Curves of test dice vs. epoch on four different noise settings.

**Implementation Details.** We adopt nnU-Net [9] as the segmentation network. Following [14], we take two networks sharing the same architecture yet with different initializations. Following [8], the noise rate is assumed to be known, and we set initial selection ratio  $\mathcal{R}$  as  $(1 - \text{noise rate})$  and  $\gamma$  as 1.1. The balance

factor  $\lambda$  is 0.65. We train our model by a SGD optimizer, with a constant learning rate 0.005. The batch sizes are 32 for ISIC dataset and 8 for JSRT dataset. We implement the code framework with PyTorch on TITAN Xp GPU.

**Evaluation Metric.** During testing, we use the standard metric Dice coefficient (Dice) to evaluate the quality of predicted masks. We stop iterative learning when label refinement cannot bring any benefit, i.e.,  $G_t$  no longer shows the rising trend for training. To make fair comparisons, we train all methods for maximum 200 epochs and report average Dice over the last 10 epochs.

### 3.3 Experiments on ISIC Dataset

Table 1 reports a summary of quantitative results of ISIC dataset. At the mild noise setting ( $\alpha = 0.3, \beta = 0.5$ ), we achieve 84.00% Dice and outperform recent methods more than 1.35% Dice. As the noise increases, the performance of baseline decreases sharply, indicating the significant impact of label noise. Other methods mitigate this impact to some extent, but their performance still drop notably. By contrast, our method consistently outperforms them and maintains high performance, validating its robustness to different noise settings. Remarkably, in the extreme noise setting ( $\alpha = 1.0, \beta = 0.7$ ), our method achieves 81.39% Dice and outperforms JoCoR (7.09% Dice), Co-teaching (7.71% Dice) and Tri-Network (11.38% Dice)<sup>3</sup>.

**Table 2.** Ablation study on our model components.

Method	Supersixel	Selection	Label Refinement	Dice [%]
Ours	✓	✓	✓	<b>83.12</b>
	×	✓	✓	81.15
	✓	×	✓	79.32
	✓	✓	×	80.56

In Fig. 2, we show curves of test dice vs. epochs. Most methods first reach a high performance then gradually decrease, indicating that their training is affected by noisy labels. In contrast, our method demonstrates a consistent high performance, which verifies the robustness of our training method. We also show some qualitative comparisons for visualization in the supplementary.

We also observe that our method is robust against the inaccuracy in supersixelization. For the ISIC dataset, we use 100 supersixels per image with relative high undersegmentation error (1.0), and our supersixel selection can potentially discard inaccurate supersixels in the noise-aware learning.

<sup>3</sup> We also observe that our method outperforms the baseline with 84.26% Dice on the original dataset, likely due to the noise in manual annotations [7].



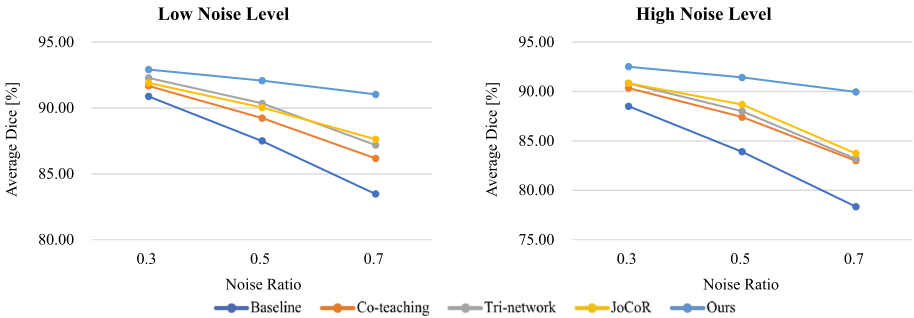
### 3.4 Ablation Study

We first verify the effect of superpixel representation based on a set of experiments on ISIC dataset under the noise setting ( $\alpha = 0.7, \beta = 0.7$ ), whose results are shown in Table 2. Row #1 is our method which achieves 83.12% Dice. Changing superpixel to pixel representation brings a performance drop of 1.97% Dice in row #2. This demonstrates the advantage of superpixel representation in learning with noisy labels.

In addition, to analyze the effect of selection module and label refinement module in iterative learning, we take a drop-one-out manner at the same setting. Ablating selection module from the model leads to a decrease of 3.80% Dice in row #3, meanwhile, removing label refinement module makes the performance drop 2.56% Dice in row #4. It is evident that both modules are essential for our robust iterative learning strategy. We also validate the effectiveness of adaptive stopping criterion and report the quality of refined labels in the supplementary.

### 3.5 Experiments on JSRT Dataset

To explore the generalization capability of our method, we also conduct experiments on JSRT dataset. Figure 3 presents the average results of three classes, and the table in supplementary reports detailed values for each class. Our method outperforms other methods consistently on all three classes.



**Fig. 3.** Average results on JSRT dataset with different noise settings: low noise level  $\beta = 0.5$  and high noise level  $\beta = 0.7$ , respectively.

## 4 Conclusion

In this paper, we propose a robust learning strategy for medical image segmentation. Unlike previous methods, we exploit structural prior and pixel correlation for segmentation model learning, which significantly mitigate the impact of label noise. We develop an iterative learning scheme based on superpixel representation. In each iteration, we jointly train two deep networks using selected subsets

of superpixels, and also relabel a subset of unreliable superpixels. Evaluation on two benchmarks with simulated noises demonstrates that our learning strategy achieves the state-of-the-art performance and robustness in extensive noise settings. We note that learning with realistic label errors is an important future research topic and building a benchmark with such label noises is a crucial step.

## 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.* (2012)
2. Arpit, D., et al.: A closer look at memorization in deep networks. In: *ICML (2017)*
3. Ching, T., Himmelstein, D.S., et al.: Opportunities and obstacles for deep learning in biology and medicine. *J. Royal Soc. Interface* **15**, 20170387 (2018)
4. Daoud, M.I., Atallah, A.A., Awwad, F., Al-Najjar, M., Alazrai, R.: Automatic superpixel-based segmentation method for breast ultrasound images. *Expert Systems with Applications* (2019)
5. Ginneken, B., Stegmann, M.B., Loog, M.: Segmentation of anatomical structures in chest radiographs using supervised methods: a comparative study on a public database. *Medical image analysis* (2006)
6. Gurari, D., Theriault, D., et al.: How to collect segmentations for biomedical images? a benchmark evaluating the performance of experts, crowdsourced non-experts, and algorithms. In: *2015 IEEE Winter Conference on Applications of Computer Vision* (2015)
7. Gutman, D., et al.: Skin lesion analysis toward melanoma detection: a challenge at the 2017 International Symposium on Biomedical Imaging (isbi), Hosted by the International Skin Imaging Collaboration (isic). *ISBI* (2018)
8. Han, B., et al.: Co-teaching: robust training of deep neural networks with extremely noisy labels. In: *NeurIPS* (2018)
9. Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods* (2020)
10. 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. *ArXiv* (2018)
11. Kohli, M., Summers, R., Geis, J.: Medical image data and datasets in the era of machine learning—whitepaper from the 2016 c-mimi meeting dataset session. *J. Digital Imaging* **30**, 392–399 (2017)
12. Li, Y., Jia, L., Wang, Z., Qian, Y., Qiao, H.: Un-supervised and semi-supervised hand segmentation in egocentric images with noisy label learning. *Neurocomputing* (2019)
13. Litjens, G., et al.: A survey on deep learning in medical image analysis. *Medical Image Analysis* (2017)
14. Malach, E., Shalev-Shwartz, S.: Decoupling “when to update” from “how to update”. In: *NIPS* (2017)
15. Mirikharaji, Z., Yan, Y., Hamarneh, G.: Learning to segment skin lesions from noisy annotations. *ArXiv* (2019)
16. Northcutt, C.G., Jiang, L., Chuang, I.: Confident learning: Estimating uncertainty in dataset labels. *ArXiv* (2019)

17. Qin, W., et al.: Supervoxel-based and boundary-sensitive convolutional neural network for automated liver segmentation. *Physics in Medicine & Biology* (2018)
18. Ren, M., Zeng, W., Yang, B., Urtasun, R.: Learning to reweight examples for robust deep learning. In: *ICML* (2018)
19. Shen, D., Wu, G., Suk, H.: Deep learning in medical image analysis. *Annual review of biomedical engineering* (2017)
20. Shiraishi, J., et al.: Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists' detection of pulmonary nodules. *AJR, American journal of roentgenology* (2000)
21. Tian, Z., Liu, L., Zhang, Z., Fei, B.: Supervoxel-based segmentation for 3d prostate mr images. *IEEE Trans. Med. Imaging* **35**, 791–801 (2015)
22. Wei, H., Feng, L., Chen, X., An, B.: Combating noisy labels by agreement: a joint training method with co-regularization. *CVPR* (2020)
23. Xue, C., Deng, Q., Li, X., Dou, Q., Heng, P.: Cascaded robust learning at imperfect labels for chest x-ray segmentation. In: *MICCAI* (2020)
24. Zhang, M., et al.: Characterizing label errors: confident learning for noisy-labeled image segmentation. In: *MICCAI* (2020)
25. Zhang, T., Yu, L., Hu, N., Lv, S., Gu, S.: Robust medical image segmentation from non-expert annotations with tri-network. In: *MICCAI* (2020)
26. Zhu, H., Shi, J., Wu, J.: Pick-and-learn: automatic quality evaluation for noisy-labeled image segmentation. In: *MICCAI* (2019)