



Universal Loss Reweighting to Balance Lesion Size Inequality in 3D Medical Image Segmentation

Boris Shirokikh^{1,2,3(✉)}, Alexey Shevtsov^{2,3}, Anvar Kurmukov^{2,4},
Alexandra Dalechina⁵, Egor Krivov^{2,3}, Valery Kostjuchenko⁵,
Andrey Golanov⁶, and Mikhail Belyaev¹

¹ Skolkovo Institute of Science and Technology, Moscow, Russia
`boris.shirokikh@phystech.edu`

² Kharkevich Institute for Information Transmission Problems, Moscow, Russia

³ Moscow Institute of Physics and Technology, Moscow, Russia

⁴ Higher School of Economics, Moscow, Russia

⁵ Moscow Gamma-Knife Center, Moscow, Russia

⁶ Burdenko Neurosurgery Institute, Moscow, Russia

Abstract. Target imbalance affects the performance of recent deep learning methods in many medical image segmentation tasks. It is a twofold problem: class imbalance – positive class (lesion) size compared to negative class (non-lesion) size; lesion size imbalance – large lesions overshadow small ones (in the case of multiple lesions per image). While the former was addressed in multiple works, the latter lacks investigation. We propose a loss reweighting approach to increase the ability of the network to detect small lesions. During the learning process, we assign a weight to every image voxel. The assigned weights are inversely proportional to the lesion volume, thus smaller lesions get larger weights. We report the benefit from our method for well-known loss functions, including Dice Loss, Focal Loss, and Asymmetric Similarity Loss. Additionally, we compare our results with other reweighting techniques: Weighted Cross-Entropy and Generalized Dice Loss. Our experiments show that *inverse weighting* considerably increases the detection quality, while preserves the delineation quality on a state-of-the-art level. We publish a complete experimental pipeline (https://github.com/neuro-ml/inverse_weighting) for two publicly available datasets of CT images: LiTS and LUNA16. We also show results on a private database of MR images for the task of multiple brain metastases delineation.

Keywords: Segmentation · CNN · Lung nodules · Brain metastases · CT · MRI

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-59719-1_51) contains supplementary material, which is available to authorized users.

1 Introduction

In recent years, convolutional neural networks (CNNs) have become the dominant approach to solve medical image segmentation tasks [14]. A wide variety of CNN models, training procedures and loss functions were built under the BRATS [16] and ISLES [15] competitions. The most common way to measure the performance of such a new method is to use segmentation voxel-wise metrics, e.g. Dice Score [2]. However, in the case of multiple lesions per image, clinical tasks also require analyzing algorithm in terms of the detection quality. For instance, all tumors, including the smallest ones, should be found and delineated in the brain stereotactic radiosurgery or in the lung cancer screening process. But since the Dice Score is a voxel-wise metric, it does not differentiate between missing several True Positives in a large lesion or in a small one.

Learning a model under the presence of extremely small targets is challenging. This is especially the problem for 3D medical image segmentation tasks. The total fraction of voxels with lesion is about 0.1% in the case of lung nodules and about 1% in case of multiple brain metastases. Moreover, in a series of medical image segmentation tasks we have a problem with the size imbalance. In some cases, large lesions could be up to 50 times bigger than the small ones (see typical lesion diameters distribution on Fig. 2).

Several approaches have been suggested to tackle the problem of target imbalance. The main idea is to add weight to a loss function to equally represent each class (lesion vs non-lesion or different lesion types in a multi-class problem). It is implemented, for example, in Weighted Cross-Entropy [18] and Generalized Dice Loss [19]. The shortcoming of this approach is that it pays attention only to the lesion type, but not the lesion size (see Fig. 1). Besides, most of the research focuses on the delineation quality and lacks an investigation into the detection performance. Ideal segmentation implies perfect detection, however, due to the substantial differences between large and small lesions, almost a perfect delineation could have poor detection quality. Here we address this problem by applying the idea of weighting a loss function with respect to target sizes.

Our contribution is twofold:

- We propose a loss function reweighting strategy, that balances the lesions of different sizes. We call our approach **inverse weighting**, since the generated weights are inversely proportional to the lesion size.
- We evaluate the effect of using the most popular segmentation loss functions on segmentation quality and network's ability to detect lesions of different sizes. On a series of medical image segmentation tasks, we show how our approach improves the detection quality, especially for small lesions (Fig. 3), while preserving delineation performance.

2 Related Work

A large number of neural network architectures, improved training procedures, and loss functions have been proposed in recent years. We extensively investigate

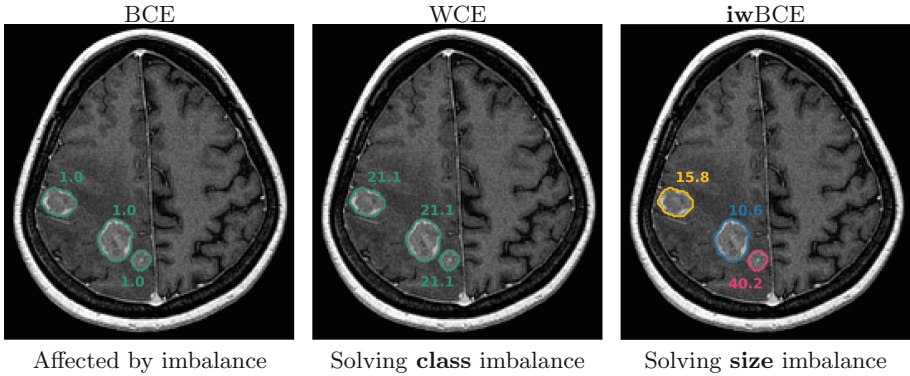


Fig. 1. The effect of inverse weighting. No reweighting applied (left), class balancing via weighted-cross entropy (center), **inverse weighting** (right). Weights for every tumor are calculated using formulas in Table 1 and placed near the tumors.

the behavior of loss functions keeping the rest of the deep learning pipeline on the state-of-the-art level without diving into details.

The *Binary Cross-Entropy* (BCE) is the standard loss function commonly used for segmentation tasks. It does not handle the problem of class imbalance and differently sized objects thus often yielding poor results. Authors of [13] suggested using *Focal Loss* as an extension of BCE in highly class imbalanced detection tasks and it is widely used in segmentation tasks as well [8]. Focal Loss does not apply any type of reweighting but automatically focuses the network attention on difficult examples. *Dice Loss* [17] has recently become one of the state-of-the-art losses for medical image segmentation tasks. The authors claim that Dice Loss establishes the right balance between classes without assigning any weights. But for the tasks with multiple targets, a large object overshadows the small one, hence the network tends to miss small lesions. Recent work [8] proposed *Asymmetric Similarity Loss* (ASL) based on F_β score. ASL extends Dice Loss (the special case with $\beta = 1$) and allows training a network with a better balance between precision and recall. But it shares the same drawback with Dice Loss: differently sized overshadowing objects. Authors of [5] proposed Sensitivity-Specificity loss which we left without consideration. It performs worse than Dice Loss on a 3D medical image segmentation task in [19] and utilizes a similar idea with ASL.

Several approaches reweight BCE and Dice Loss to improve network performance in medical image segmentation tasks. In [18] authors use *Weighted Cross-Entropy* (WCE) loss and [19] suggest Generalized Dice Loss (GDL) to tackle the problem of class imbalance. Both approaches utilize the same idea of reweighting the corresponding losses with weights inverse to the sizes of classes (see Table 1). Our approach simultaneously solves class imbalance problem and imbalance between differently sized objects. A deeper modification of Cross-Entropy loss to handle class imbalance is evaluated in [11], but the goal is quite

different – overfitting on small datasets. In [21] authors suggest, a highly depend-able on hyperparameters, a combination of Cross-Entropy and logarithmic Dice Loss to solve multiclass (19 classes) segmentation problem. In our work, we show an improvement for both of these losses independently.

We focus our attention on the most relevant loss functions and their explicitly reweighted modifications. Below we detail how our method is applied to state-of-the-art losses and compare it with WCE and GDL.

3 Method

We find out that all models tend to miss small targets when training with BCE or Focal Loss. We assume poor performance comes from the inability of these losses to equally represent differently sized targets. Dice Loss and ASL have the same drawback: large targets overshadow the small ones. Moreover, already developed losses handle only the imbalance between classes, not between lesion sizes. We aim to close the gap and propose a simple methodology to reweight loss functions in the way that all targets contribute equally, e.g. small targets have greater weights.

During the training stage, we generate a tensor of weights for every incoming patch. To form such a tensor we split the corresponding ground-truth patch into $K + 1$ connected components L_0, \dots, L_K , where L_0 is the non-lesion component (background) and K is the number of lesions in the current patch. Next, we assign the weight to every component which is inverse to the component’s volume:

$$w_j = \frac{\sum_{k=0}^K |L_k|}{(K + 1) \cdot |L_j|}, \quad (1)$$

here w_j is the weight, assigned to every voxel inside the corresponding component L_j . The constant in the denominator ensures that the sum of our weights is equal to the sum of the unit tensor of the same size (see derivation details in Supplementary Materials). We call this method **inverse weighting (iw)**. Note, how our approach assigns greater weights to the smaller tumors (Fig. 1). At this point, we can modify any of the discussed loss functions with our reweighting. Since WCE and GDL explicitly reweight state-of-the-art losses, we do not apply reweighting twice. Corresponding modifications for BCE, Focal Loss, Dice Loss, and ASL are shown in the Table 1.

4 Experiments

4.1 Data

We report our results on three datasets. Two publicly available datasets that include 3D CT images: LUNA16 [10] with lung cancerous nodules and LiTS [4] with liver tumors; and one private dataset with MR images of multiple brain metastases.

Table 1. Loss functions and their modifications. Here y_i denotes the i^{th} element of the ground truth binary mask, p_i is the corresponding predicted probability, and \mathbf{w}_i is the proposed inverse weight.

Loss	Original expression	Proposed modification (\mathbf{iw})
BCE	$-(y_i \log p_i + (1 - y_i) \log(1 - p_i))$	$-\mathbf{w}_i (y_i \log p_i + (1 - y_i) \log(1 - p_i))$
Focal Loss $_{\gamma, \alpha}$	$-\alpha(1 - p_i)^\gamma y_i \log p_i$ $+ (1 - \alpha)p_i^\gamma (1 - y_i) \log(1 - p_i)$	$-\mathbf{w}_i (\alpha(1 - p_i)^\gamma y_i \log p_i$ $+ (1 - \alpha)p_i^\gamma (1 - y_i) \log(1 - p_i))$
WCE	$-w y_i \log p_i - (1 - y_i) \log(1 - p_i), w = \frac{n - \sum_i p_i}{\sum_i p_i}$	—
Dice Loss	$1 - \frac{2 \sum_i p_i y_i}{\sum_i (p_i^2 + y_i^2)}$	$1 - \frac{2 \sum_i \mathbf{w}_i p_i y_i}{\sum_i \mathbf{w}_i (p_i^2 + y_i^2)}$
ASL $_{\beta}$	$1 - \frac{(1 + \beta^2) \sum_i p_i y_i}{\sum_i (\beta^2 y_i + p_i)}$	$1 - \frac{(1 + \beta^2) \sum_i \mathbf{w}_i p_i y_i}{\sum_i \mathbf{w}_i (\beta^2 y_i + p_i)}$
GDL	$1 - \frac{2 \sum_{c=1}^2 w_c^2 \sum_i p_i y_i}{\sum_{c=1}^2 w_c^2 \sum_i (p_i^2 + y_i^2)}, w_c = \frac{1}{\sum_i y_i}$	—

LUNA16 includes 816 (we have excluded 72 cases with nodules located outside of lung masks) annotated chest scans from LIDC/IDRI database [1]. For every image, we clip intensities between -1000 and 300 Hounsfield units (HU), and then set the voxels outside the given binary lung masks to -1000 . Ground truth mask was formed by averaging 4 given annotations.

Metastases (private dataset) includes 1952 unique patients with the T1-weighted MRI of the head. We apply no preprocessing steps to these images.

LITS includes 131 annotated CT abdomen scans. For every image, we clip intensities between -300 and 300 HU and then apply a given binary mask of liver the same way we did it with LUNA16 data.

Before passing through the network, we scale images to have voxel’s intensities between 0 and 1.

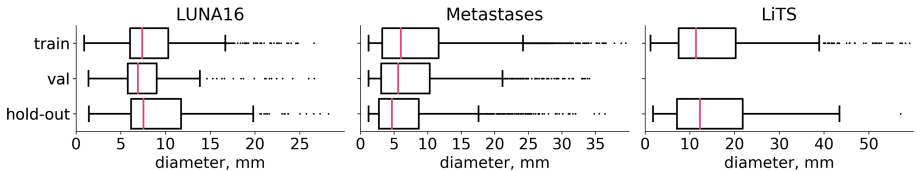


Fig. 2. Lesion diameters distribution. Metastases under 5 mm, lung nodules under 10 mm and liver tumors under 12 mm are considered *small*, according to the clinical recommendations [3, 12]

We use *train-validation* setup to compare different architectures and hyper-parameters for loss functions. Then the merged combination of *training* and *validation* data is used to train the chosen methods and we report final results on previously unseen *hold-out* set. LUNA16 is presented as 10, approximately

equal, subsets [10] thus we use the first 6 for *training* (534 images), next 2 for *validation* (178 images) and the last pair as *hold-out* (174 images). We divide Metastases into *training* (1250 images), *validation* (402 images) and *hold-out* (300 images). LiTS is also presented as 2 subsets, so we use the first for *training* (104 images) and the second as *hold-out* (27 images). We do not shrink the validation part of the LiTS, since this dataset is used only once for the final results reporting.

4.2 Architecture and Training

For all our experiments we consistently use a single CNN model – slightly modified 3D U-Net [6]. Implemented architecture within PyTorch framework is available in our repository along with a schematic image. Following the suggestion of [9], we do not focus our attention on fine-tuning the CNN model.

In all scenarios we train the model for 100 epochs, starting with learning rate of 10^{-2} , and reducing it to 10^{-3} at the epoch 80. Each epoch consists of 100 iterations of stochastic gradient descent with Nesterov momentum (0.9). At every iteration we sample patches of size $128 \times 128 \times 128$ and batch size of two. With the probability of 0.5 we sample the patch so that it contains at least one voxel with lesion, otherwise we sample it uniformly. The training takes about 26 hours on a 24 GB NVIDIA Tesla M40 GPU.

Note, that only two of the considered loss functions have hyperparameters: ASL (β) and Focal Loss (γ, α). We use ASL with $\beta = 1.5$ originally recommended in [8]. For Focal Loss we also use $\gamma = 2$ originally recommended in [13], but change α to be 0.75 chosen on validation.

4.3 Metric

Dice Score has a particular drawback measuring the delineation quality in the tasks with multiple lesions per image: big lesion overshadows small ones. We use **object Dice Score** – the average Dice Score over *unique found lesions*. Therefore it does not shift towards larger lesions. Note that we exclude missed lesions from this analysis, hence the delineation quality is independent from detection quality.

To measure the detection quality we suggest using a Free-response Receiver Operating Characteristic (FROC) curve analysis. It is extremely efficient operating with multiple targets and False Positive (FP) responses per case [7]. A FROC curve measures the sensitivity to detected objects instead of voxel-wise sensitivity, therefore does not have the same drawback of overshadowed lesions. A FROC curve summarizes the model’s efficiency with the trade-off between the fraction of lesions detected (Recall) and the average number of FPs per image. But it gives us only visual representation of experimental results. To compare the performance of different methods we extract a single value from the curves. Authors of [20] suggested using the **average object-wise Recall** over the pre-defined FP values (1/8, 1/4, 1/2, 1, 2, 4, 8) which is also the main metric of

LUNA16 challenge [10]. This metric gives us the average fraction of detected lesions per case which is highly interpretable in terms of detection quality.

To calculate the confidence intervals for FROC curves and for average Recall we use bootstrapping. We sample 80% of test patients and build a curve on every of the 100 iterations. Average recall is calculated for every bootstrapped curve and we report the mean value along with the standard deviation.

4.4 Results and Discussion

We visualize our main contribution with the considerable improvement of the average object-wise Recall for all four chosen loss functions on all three datasets (Fig. 3). We also report our metrics separately for three groups of lesion sizes and show a solid contribution into the small lesion detection quality which satisfies our method’s motivation. However, a comparison with WCE is worth a more detailed discussion.

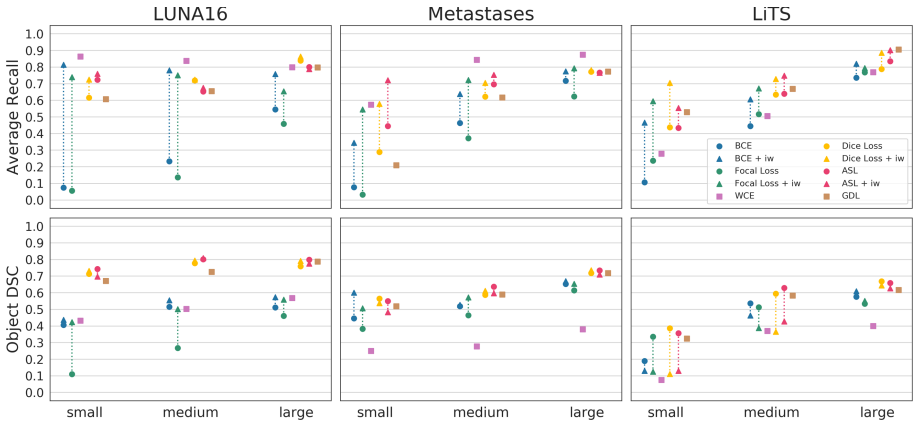


Fig. 3. The impact of inversely weighted loss functions in terms of average Recall and object-wise Dice Score. We show performance on three approximately equal subsets (1/3 each) of lesions divided by their size. Small and medium groups correspond to the clinical recommendations of small lesions (see Fig. 2).

Images from LUNA16 contain 1.3 nodules per scan on average, while Metastases and LiTS have 4.8 and 6.9 tumors per scan respectively. The latter means that LUNA16 is hardly an appropriate dataset to benefit from our method, since the majority of training patches contain only one lesion. One lesion per patch is clearly the *class imbalance* problem, and WCE outperforms the other methods in terms of average Recall. But nevertheless we show inverse weighting solving also the class imbalance task on the competitive level solidly improving BCE and Focal Loss performance. Finally, even the slight improvement in the detection quality of WCE comes with the dramatic delineation quality loss on the other two datasets, which is crucial for clinical tasks.

GDL failed to surpass inversely weighted loss function almost in all scenarios. But overall we find ASL and Dice Loss along with GDL and their inversely weighted modifications to be highly stable during the training. Respectively, Dice-like loss function sufficiently outperform BCE-like losses both in terms of the detection and the delineation qualities. We believe such a behaviour comes from two properties of Dice Loss. Firstly, it is designed to optimize the Dice Score metric, and one could clearly see the dominance of Dice-like losses in terms of object Dice Score (Fig. 3 and Table 2). Secondly, it partially solves the class imbalance problem, but only in the cases with exactly one object per patch. The latter is again perfectly demonstrated on LUNA16, as we put this dataset to be more about class imbalance problem in the previous paragraph. One could see the already high object Dice Scores and average Recall values of ASL and Dice Loss on LUNA16 along with minor changes of their reweighting.

However, modified with inverse weighting loss functions have a noticeable decrease in delineation quality on LiTS data. We consider this to be a side effect of highly increased object-wise Recall: *modified losses find more difficult cases, hence joint object Dice Score could decrease.*

Besides the separate performance on lesion sizes we also include more detailed results for all lesions in hold-out sets (Table 2). We give the visual representation of experimental results in terms of detection quality via FROC analysis (see Supplementary Materials).

Table 2. Results for all considered loss functions along with the proposed method – inverse weighting (“+” with iw, “–” without iw). The numbers in brackets are standard deviation.

	iw	LUNA16		Metastases		LiTS	
		avg Recall	obj DSC	avg Recall	obj DSC	avg Recall	obj DSC
BCE	–	.42 (.02)	.57 (.28)	.47 (.01)	.67 (.25)	.47 (.03)	.61 (.29)
	+	.67 (.01)	.56 (.20)	.52 (.01)	.66 (.23)	.59 (.03)	.53 (.27)
Focal Loss	–	.35 (.02)	.51 (.28)	.40 (.01)	.64 (.25)	.50 (.03)	.58 (.27)
	+	.55 (.01)	.54 (.20)	.52 (.01)	.63 (.21)	.62 (.03)	.48 (.27)
WCE	–	.74 (.01)	.50 (.17)	.54 (.01)	.39 (.22)	.52 (.04)	.41 (.29)
Dice Loss	–	.71 (.02)	.76 (.20)	.55 (.01)	.69 (.23)	.62 (.03)	.63 (.25)
	+	.73 (.02)	.77 (.16)	.57 (.01)	.68 (.21)	.72 (.03)	.49 (.30)
ASL	–	.68 (.02)	.77 (.16)	.55 (.01)	.71 (.20)	.66 (.03)	.63 (.24)
	+	.70 (.02)	.76 (.18)	.59 (.02)	.66 (.20)	.73 (.03)	.53 (.28)
GDL	–	.69 (.02)	.73 (.20)	.53 (.01)	.70 (.22)	.69 (.03)	.60 (.27)

5 Conclusion

We propose a universal approach to loss functions reweighting. It could be used with almost any state-of-the-art loss function. Our experiment demonstrates an improvement of network’s ability to detect lesions for Cross-Entropy, Focal Loss, Dice Loss and Asymmetric Similarity Loss on three medical tasks with multiple

targets per case. Moreover, we believe the method can also improve quality with other complex multi-stage pipelines or with any other CNN architecture which is the goal for our future research.

Acknowledgements. The results of the paper are based on the scientific research supported by the Russian Science Foundation under grant 17-11-01390. The authors also acknowledge the National Cancer Institute and the Foundation for the National Institutes of Health, and their critical role in the creation of the free publicly available LIDC/IDRI Database used in this study.

References

1. 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)
2. Bakas, S., et al.: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. arXiv preprint [arXiv:1811.02629](https://arxiv.org/abs/1811.02629) (2018)
3. Bankier, A.A., MacMahon, H., Goo, J.M., Rubin, G.D., Schaefer-Prokop, C.M., Naidich, D.P.: Recommendations for measuring pulmonary nodules at CT: a statement from the fleischner society. *Radiology* **285**(2), 584–600 (2017)
4. Bilic, P., et al.: The liver tumor segmentation benchmark (LiTS). arXiv preprint [arXiv:1901.04056](https://arxiv.org/abs/1901.04056) (2019)
5. Brosch, T., Yoo, Y., Tang, L.Y.W., Li, D.K.B., Traboulsee, A., Tam, R.: Deep convolutional encoder networks for multiple sclerosis lesion segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 3–11. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_1
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. DeLuca, P., Wambersie, A., Whitmore, G.: Extensions to conventional ROC methodology: LROC, FROC, and AFROC. *J ICRU* **8**(1), 31–5 (2008)
8. Hashemi, S.R., Salehi, S.S.M., Erdogmus, D., Prabhu, S.P., Warfield, S.K., Gholipour, A.: Asymmetric loss functions and deep densely-connected networks for highly-imbalanced medical image segmentation: application to multiple sclerosis lesion detection. *IEEE Access* **7**, 1721–1735 (2018)
9. Isensee, F., Kickingeder, P., Wick, W., Bendszus, M., Maier-Hein, K.H.: No new-net. In: Crimi, A., Bakas, S., Kuijf, H., Keyvan, F., Reyes, M., van Walsum, T. (eds.) BrainLes 2018. LNCS, vol. 11384, pp. 234–244. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-11726-9_21
10. Jacobs, C., Setio, A.A.A., Traverso, A., van Ginneken, B.: Lung nodule analysis 2016 (2016). <https://luna16.grand-challenge.org>
11. Li, Z., Kamnitsas, K., Glocker, B.: Overfitting of neural nets under class imbalance: analysis and improvements for segmentation. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11766, pp. 402–410. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32248-9_45

12. Lin, N.U., et al.: Response assessment criteria for brain metastases: proposal from the rano group. *Lancet Oncol.* **16**(6), e270–e278 (2015)
13. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2980–2988 (2017)
14. Litjens, G., et al.: A survey on deep learning in medical image analysis. *Med. Image Anal.* **42**, 60–88 (2017)
15. Maier, O., et al.: Isles 2015—a public evaluation benchmark for ischemic stroke lesion segmentation from multispectral MRID. *Med. Image Anal.* **35**, 250–269 (2017)
16. Menze, B.H., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Trans. Med. Imaging* **34**(10), 1993–2024 (2014)
17. 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)
18. 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
19. Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Jorge Cardoso, M.: Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In: Cardoso, M.J., et al. (eds.) *DLMIA/ML-CDS -2017*. LNCS, vol. 10553, pp. 240–248. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67558-9_28
20. Van Ginneken, B., et al.: Comparing and combining algorithms for computer-aided detection of pulmonary nodules in computed tomography scans: the ANODE09 study. *Med. Image Anal.* **14**(6), 707–722 (2010)
21. Wong, K.C.L., Moradi, M., Tang, H., Syeda-Mahmood, T.: 3D segmentation with exponential logarithmic loss for highly unbalanced object sizes. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) *MICCAI 2018*. LNCS, vol. 11072, pp. 612–619. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00931-1_70