

# Robust T-Loss for Medical Image Segmentation

Alvaro Gonzalez-Jimenez<sup>1</sup>, Simone Lionetti<sup>2</sup>, Philippe Gottfrois<sup>1</sup>,  
Fabian Gröger<sup>1,2</sup>, Marc Pouly<sup>2</sup>, and Alexander Navarini<sup>1</sup>

<sup>1</sup> University of Basel, Switzerland

<sup>2</sup> Lucerne School of Computer Science and Information Technology, Switzerland  
`alvaro.gonzalezjimenez@unibas.ch`

**Abstract.** This paper presents a new robust loss function, the T-Loss, for medical image segmentation. The proposed loss is based on the negative log-likelihood of the Student-t distribution and can effectively handle outliers in the data by controlling its sensitivity with a single parameter. This parameter is updated during the backpropagation process, eliminating the need for additional computation or prior information about the level and spread of noisy labels. Our experiments show that the T-Loss outperforms traditional loss functions in terms of dice scores on two public medical datasets for skin lesion and lung segmentation. We also demonstrate the ability of T-Loss to handle different types of simulated label noise, resembling human error. Our results provide strong evidence that the T-Loss is a promising alternative for medical image segmentation where high levels of noise or outliers in the dataset are a typical phenomenon in practice. The project website can be found at <https://robust-tloss.github.io>.

**Keywords:** robust loss · medical image segmentation · noisy labels

## 1 Introduction

Convolutional Neural Networks (CNNs) and Visual Transformers (ViTs) have become the standard in semantic segmentation, achieving state-of-the-art results in many applications [1, 24, 16]. However, supervised training of CNNs and ViTs requires large amounts of annotated data, where each pixel in the image is labeled with the category it belongs to. In the medical domain, obtaining these annotations can be costly and time-consuming as it requires expertise and domain knowledge that is often scarcely available [6]. In addition, medical image annotations can be affected by human bias and poor inter-annotator agreement [23], further complicating the process. Despite efforts to obtain labels through automated mining [31] and crowd-sourcing methods [11], the quality of datasets gathered using these methods remains challenging due to often high levels of label noise.

For instance, the Fitzpatrick 17k dataset, commonly used in dermatology research, contains non-skin images and noisy annotations. In a random sample of 504 images, 5.4% were labeled incorrectly or as other classes [10]. The dataset was scraped from online atlases, which makes it vulnerable to inaccuracies and

noise [10]. Noisy labels are and will continue to be, a problem in medical datasets. This is a concern as label noise has been shown to decrease the accuracy of supervised models [20, 22, 35], making it a key area of focus for both research and practical applications.

Previous literature has explored many methods to mitigate the problem of noisy labels in deep learning. These methods can be broadly categorized into label correction [27, 28, 32], loss function correction based on an estimated noise transition matrix [21, 33, 29], and robust loss functions [34, 30, 2, 18]. Compared to the first two approaches, which may suffer from inaccurate estimates of the noise transition matrix, robust loss functions enable joint optimization of model parameters and variables related to the noise model and have shown promising results in classification tasks [8, 34]. Despite these advances, semantic segmentation with noisy labels is relatively understudied. Existing research in this area has focused on the development of noise-resistant network architectures [15], the incorporation of domain-specific prior knowledge [29], or more recent strategies that update the noisy masks before memorization [17].

Although previous methods have shown robustness in semantic segmentation, they often have limitations, such as more hyper-parameters, modifications to the network architecture, or complex training procedures. In contrast, robust loss functions offer a much simpler solution as they could be incorporated with a simple change in a single modeling component. However, their effectiveness has not been thoroughly investigated.

In this work, we show that several traditional robust loss functions are vulnerable to memorizing noisy labels. To overcome this problem, we introduce a novel robust loss function, the T-Loss, which is inspired by the negative log-likelihood of the Student-t distribution. The T-Loss, whose simplest formulation features a single parameter, can adaptively learn an optimal tolerance level to label noise directly during backpropagation, eliminating the need for additional computations such as the Expectation Maximization (EM) steps.

To evaluate the effectiveness of the T-Loss as a robust loss function for medical semantic segmentation, we conducted experiments on two widely-used benchmark datasets in the field: one for skin lesion segmentation and the other for lung segmentation. We injected different levels of noise into these datasets that simulate typical human labeling errors and trained deep learning models using various robust loss functions. Our experiments demonstrate that the T-Loss outperforms these robust state-of-the-art loss functions in terms of segmentation accuracy and robustness, particularly under conditions of high noise contamination. We also observed that the T-Loss could adaptively learn the optimal tolerance level to label noise which significantly reduces the risk of memorizing noisy labels.

This research is divided as follows: Section 2 introduces the motivation behind our T-Loss and provides its mathematical derivation. Section 3 covers the datasets used in our experiments, the implementation and training details of T-Loss, and the metrics used for comparison. Section 4 presents the main findings of our study, including the results of the T-Loss and the baselines on both

datasets and an ablation study on the parameter of T-Loss. Finally, in section 5, we summarize our contributions and the significance of our study for the field.

## 2 Methodology

Let  $\mathbf{x}_i \in \mathbb{R}^{c \times w \times h}$  be an input image and  $\mathbf{y}_i \in \{0, 1\}^{w \times h}$  be its noisy annotated binary segmentation mask, where  $c$  represents the number of channels,  $w$  the image's width, and  $h$  its height. Given a set of images  $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$  and corresponding masks  $\{\mathbf{y}_1, \dots, \mathbf{y}_N\}$ , our goal is to train a model  $f_{\mathbf{w}}$  with parameters  $\mathbf{w}$  such that  $f_{\mathbf{w}}(\mathbf{x})$  approximates the accurate binary segmentation mask for any given image  $\mathbf{x}$ .

To this end we note that, heuristically, assuming error terms to follow a Student-t distribution (as suggested e.g. in [19]) allows for significantly larger noise tolerance with respect to the usual gaussian form. Recall that the Student-t distribution for a  $D$ -dimensional variable  $\mathbf{y}$  is defined by the Probability Density Function (PDF)

$$p(\mathbf{y}|\boldsymbol{\mu}, \boldsymbol{\Sigma}; \nu) = \frac{\Gamma(\frac{\nu+D}{2})}{\Gamma(\frac{\nu}{2})} \frac{|\boldsymbol{\Sigma}|^{-1/2}}{(\pi\nu)^{D/2}} \left[ 1 + \frac{(\mathbf{y} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \boldsymbol{\mu})}{\nu} \right]^{-\frac{\nu+D}{2}}, \quad (1)$$

where  $\boldsymbol{\mu}$  and  $\boldsymbol{\Sigma}$  are respectively the mean and the covariance matrix of the associated multivariate normal distribution,  $\nu$  is the number of degrees of freedom, and  $|\cdot|$  indicates the determinant (see e.g. [3]). From this expression, we see that the tails of the Student-t distribution follow a power law that is indeed heavier compared to the usual negative quadratic exponential. For this reason, it is well known to be robust to outliers [26, 7].

Since the common Mean Squared Error (MSE) loss is derived by minimizing the negative log-likelihood of the normal distribution, we choose to apply the same transformation and get

$$-\log p(\mathbf{y}|\boldsymbol{\mu}, \boldsymbol{\Sigma}; \nu) = -\log \Gamma\left(\frac{\nu+D}{2}\right) + \log \Gamma\left(\frac{\nu}{2}\right) + \frac{1}{2} \log |\boldsymbol{\Sigma}| + \frac{D}{2} \log(\pi\nu) + \frac{\nu+D}{2} \log \left[ 1 + \frac{(\mathbf{y} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \boldsymbol{\mu})}{\nu} \right]. \quad (2)$$

The functional form of our loss function for one image is then obtained with the identification  $\mathbf{y} = \mathbf{y}_i$  and the approximation  $\boldsymbol{\mu} = f_{\mathbf{w}}(\mathbf{x}_i)$ , and aggregated with

$$\mathcal{L}_T = \frac{1}{N} \sum_{i=1}^N -\log p(\mathbf{y}_i | f_{\mathbf{w}}(\mathbf{x}_i), \boldsymbol{\Sigma}; \nu). \quad (3)$$

Equation (2) has  $D(D+1)/2$  free parameters in the covariance matrix, which should be estimated from the data. In the case of images, this can easily be in the order of  $10^4$  or larger, which makes a general computation highly non-trivial and may deteriorate the generalization capabilities of the model. For these reasons,

we take  $\Sigma$  to be the identity matrix  $\mathbf{I}_D$ , despite knowing that pixel annotations in an image are not independent. The loss term for one image simplifies to

$$-\log p(\mathbf{y}|\boldsymbol{\mu}, \mathbf{I}_D; \nu) = -\log \Gamma\left(\frac{\nu + D}{2}\right) + \log \Gamma\left(\frac{\nu}{2}\right) + \frac{D}{2} \log(\pi\nu) + \frac{\nu + D}{2} \log \left[1 + \frac{(\mathbf{y} - \boldsymbol{\mu})^2}{\nu}\right]. \quad (4)$$

To clarify the relation with known loss functions, let  $\delta = |\mathbf{y}_i - f_{\mathbf{w}}(\mathbf{x}_i)|$ , and fix the value of  $\nu$ . For  $\delta \rightarrow 0$ , the functional dependence from  $\delta$  reduces to a linear function of  $\delta^2$ , i.e. MSE. For large values of  $\delta$ , though, eq. (4) is equivalent to  $\log \delta$ , thus penalizing large deviations even less than the much-advocated robust Mean Absolute Error (MAE). The scale of this transition, the sensitivity to outliers, is regulated by the parameter  $\nu$ .

We optimize the parameter  $\nu$  jointly with  $\mathbf{w}$  using gradient descent. To this end, we reparametrize  $\nu = e^{\tilde{\nu}} + \epsilon$  where  $\epsilon$  is a safeguard for numerical stability. Loss functions with similar dynamic tolerance parameters were also studied in [2] in the context of regression, where using the Student-t distribution is only mentioned in passing.

### 3 Experiments

In this section, we demonstrate the robustness of the T-Loss for segmentation tasks on two public image collections from different medical modalities, namely ISIC [5] and Shenzhen [4, 13, 25]. In line with the literature, we use simulated label noise in our tests, as no public benchmark with real label noise exists [15].

#### 3.1 Datasets

The **ISIC** 2017 dataset [5] is a well-known public benchmark of dermoscopy images for skin cancer detection. It contains 2000 training and 600 test images with corresponding lesion boundary masks. The images are annotated with lesion type, diagnosis, and anatomical location metadata. The dataset also includes a list of lesion attributes, such as size, shape, and color. We resized the images to  $256 \times 256$  pixels for our experiments.

**Shenzhen** [4, 13, 25] is a public dataset containing 566 frontal chest radiographs with corresponding lung segmentation masks for tuberculosis detection. Since there is not a predefined split for Shenzhen as in ISIC, to ensure representative training and testing sets, we stratified the images by their tuberculosis and normal lung labels, with 70% of the data for training and the remaining 30% for testing. Resulting in 296 training images and 170 test images. All images were resized to  $256 \times 256$  pixels.

Without a public benchmark with real noisy and clean segmentation masks, we artificially inject additional mask noise in these two datasets to test the model’s robustness to low annotation quality. This simulates the real risk of

errors due to factors like annotator fatigue and difficulty in annotating certain images. In particular, we follow [15], randomly sample a portion of the training data with probability  $\alpha \in \{0.3, 0.5, 0.7\}$ , and apply morphological transformations with noise levels controlled by  $\beta \in \{0.5, 0.7\}$ . The morphological transformations included erosion, dilation, and affine transformations, which respectively reduced, enlarged, and displaced the annotated area.

### 3.2 Setup

We train a nnU-Net [12] as a segmentation network from scratch. To increase variations in the training data, we augment them with random mirroring, flipping, and gamma transformations. The T-loss was initialized with  $\tilde{\nu} = 0$  and  $\epsilon = 10^{-8}$ . The nnU-Net was trained for 100 epochs using the Adam optimizer with a learning rate of  $10^{-3}$  and a batch size of 16 for the ISIC dataset and 8 for the Shenzhen dataset. The network was trained on a single NVIDIA Tesla V100 with 32 GB of memory.

The model is trained using noisy masks. However, by using the ground truth for the corresponding noisy mask, we can evaluate the robustness of the model and measure noisy-label memorization. This is done by analyzing the dice score of the model’s prediction compared to the actual ground truth.

In addition to the T-Loss, we train several other losses for comparison. Our analysis includes some traditional robust losses, such as Mean Absolute Error (MAE), Reverse Cross Entropy (RCE), Normalized Cross Entropy (NCE), and Normalized Generalized Cross Entropy (NGCE), as well as more recent methods, such as Generalized Cross Entropy (GCE) [34], Symmetrical Cross Entropy (SCE) [30], and Active-Passive Loss (APL) [18]. For APL, in particular, we consider three combinations: 1) NCE+RCE, 2) NGCE+MAE, and 3) NGCE+RCE. We consider the mean of the predictions for the last 10 epochs with a fixed number of epochs and report its mean and standard deviation over 3 different random seeds.

Finally, we complete our evaluation with statistical significance tests. We use the ANOVA test [9] to compare the differences between the means of the dice scores and obtain a *p-value*. In addition, if the difference is significant, we perform the Tukey *post-hoc* test [14] to determine which means are different. We assume statistical significance for *p-values* of less than  $p = 0.05$  and denote this with a <sup>\*</sup>.

## 4 Results

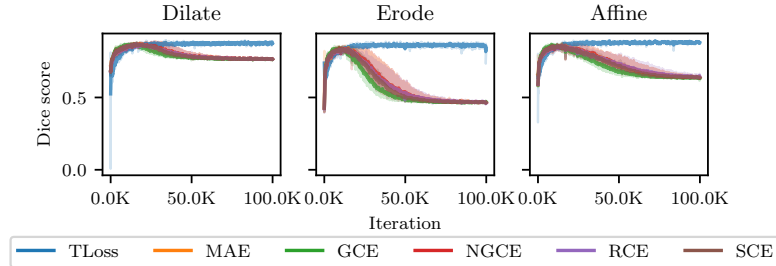
### 4.1 Results on the ISIC dataset

We present experimental results for the skin lesion segmentation task on the ISIC dataset in table 1. Our results show that conventional losses perform well with no noise or under low noise levels, but their performance decreases significantly with increasing noise levels due to the memorization of noisy labels. This can

be observed from the training dice scores in fig. 1, where traditional robust losses overfit data in later stages of learning while metrics for the T-Loss do not deteriorate. Our method achieves a dice score of  $0.788 \pm 0.007$  even for the most extreme noise scenario under exam. Examples of the obtained masks can be seen in the supplementary material.

**Table 1.** Dice score on the ISIC dataset with different noise ratios. The values refer to the mean and standard deviation over 3 different random seeds for the mean score over the last 10 epochs.

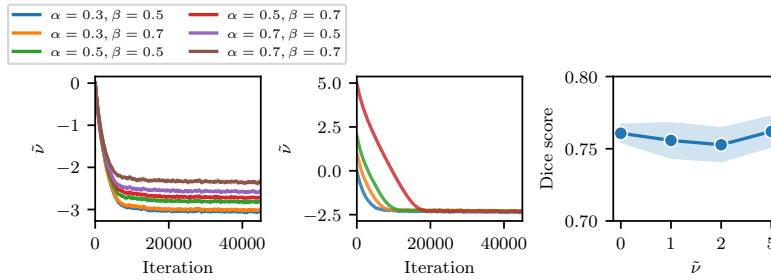
Loss	$\alpha = 0.0$	$\alpha = 0.3$		$\alpha = 0.5$		$\alpha = 0.7$	
		$\beta = 0.5$	$\beta = 0.7$	$\beta = 0.5$	$\beta = 0.7$	$\beta = 0.5$	$\beta = 0.7$
GCE	0.828(7)	0.805(9)	0.785(14)	0.772(15)	0.736(17)	0.743(12)	0.691(22)
MAE	0.826(6)	0.803(7)	0.786(12)	0.771(10)	0.742(14)	0.751(09)	0.698(20)
RCE	0.827(5)	0.802(6)	0.791(11)	0.779(11)	0.745(17)	0.752(10)	0.695(16)
SCE	0.828(6)	0.806(9)	0.793(11)	0.774(12)	0.738(14)	0.756(10)	0.691(20)
NGCE	0.825(6)	0.803(7)	0.788(08)	0.773(11)	0.745(16)	0.745(11)	0.688(19)
NCE+RCE	0.829(6)	0.799(8)	0.792(12)	0.777(11)	0.751(18)	0.746(11)	0.696(13)
NGCE+MAE	0.828(5)	0.802(8)	0.788(13)	0.774(10)	0.741(18)	0.748(11)	0.693(15)
NGCE+RCE	0.827(7)	0.807(7)	0.790(11)	0.776(13)	0.736(17)	0.748(10)	0.689(17)
T-Loss (Ours)	0.825(5)	0.809(6)	0.804(5)*	0.800(11)*	0.790(5)*	0.788(7)*	0.761(6)*



**Fig. 1.** The dice score of training set predictions compared to ground truth annotations during the training process on the ISIC 2017 dataset for each type of noisy mask with  $\alpha = 0.7, \beta = 0.7$ . The model memorizes the noisy labels after the first  $\sim 20$ K iterations, thus negatively affecting the dice score for all losses except the T-Loss.

## 4.2 Results on the Shenzhen dataset

The results of lung segmentation for the Shenzhen test set are reported in table 2. Similar to the ISIC dataset, all considered robust losses perform well at low noise



**Fig. 2.** The behavior of  $\tilde{\nu}$  in the skin lesion segmentation task. Left: convergence of  $\tilde{\nu}$  with different levels of label noise. Center: sensitivity of  $\tilde{\nu}$  to initialization for  $\alpha = 0.7$ ,  $\beta = 0.7$ . Right: sensitivity of the dice score to the initialization of  $\tilde{\nu}$  with the same settings.

levels. However, as the noise level increases, their dice scores deteriorate. On the other hand, the T-Loss stands out by consistently achieving the highest dice score, even in the most challenging scenarios. The statistical test results also support this claim, with the T-Loss being significantly superior to the other methods.

**Table 2.** Dice score on the Shenzhen dataset with different noise ratios. The values refer to the mean and standard deviation over 3 different random seeds for the mean score over the last 10 epochs.

Loss	$\alpha = 0.0$	$\alpha = 0.3$		$\alpha = 0.5$		$\alpha = 0.7$	
		$\beta = 0.5$	$\beta = 0.7$	$\beta = 0.5$	$\beta = 0.7$	$\beta = 0.5$	$\beta = 0.7$
GCE	0.948(1)	0.933(6)	0.930(8)	0.909(12)	0.880(19)	0.856(14)	0.807(23)
MAE	0.949(2)	0.937(6)	0.931(8)	0.910(11)	0.880(16)	0.864(20)	0.823(23)
RCE	0.949(2)	0.938(4)	0.931(8)	0.910(10)	0.886(20)	0.863(15)	0.818(26)
SCE	0.949(3)	0.938(4)	0.930(6)	0.908(09)	0.881(18)	0.865(15)	0.821(24)
NGCE	0.949(2)	0.936(7)	0.933(9)	0.906(12)	0.875(23)	0.865(17)	0.822(24)
NCE+RCE	0.949(2)	0.936(7)	0.928(9)	0.906(12)	0.879(18)	0.863(18)	0.818(23)
NGCE+MAE	0.949(1)	0.938(5)	0.934(6)	0.906(11)	0.877(19)	0.865(16)	0.824(21)
NGCE+RCE	0.949(2)	0.936(5)	0.930(10)	0.909(10)	0.884(17)	0.862(12)	0.821(26)
T-Loss (Ours)	0.949(1)	0.948(1)*	0.939(1)*	0.914(5)*	0.904(8)*	0.896(7)*	0.870(31)*

### 4.3 Dynamic tolerance to noise

The value of  $\tilde{\nu}$  is crucial for the model’s performance, as it controls the sensitivity to label noise. To shed light on this mechanism, we study the behavior of  $\tilde{\nu}$  during training for different label noise levels and initializations on the ISIC dataset.

As seen in fig. 2,  $\tilde{\nu}$  dynamically adjusts annotation noise tolerance in the early stages of training, independently of its initial value. The plots demonstrate that  $\tilde{\nu}$  clearly converges to a stable solution during training, with initializations far from this solution only mildly prolonging the time needed for convergence and having no significant effect on the final dice score.

## 5 Conclusions

In this contribution, we introduced the T-Loss, a loss function based on the negative log-likelihood of the Student-t distribution. The T-Loss offers the great advantage of controlling sensitivity to outliers through a single parameter that is dynamically optimized. Our evaluation on public medical datasets for skin lesion and lung segmentation demonstrates that the T-Loss outperforms other robust losses by a statistically significant margin. While other robust losses are vulnerable to noise memorization for high noise levels, the T-Loss can reabsorb this form of overfitting into the tolerance level  $\nu$ . Our loss function also features remarkable independence to different noise types and levels.

It should be noted that other methods, such as [15] offer better performance for segmentation on the ISIC dataset with the same synthetic noisy labels, while the T-Loss offers a simple alternative. The trade-off in terms of performance, computational cost, and ease of adaption to different scenarios remains to be investigated. Similarly, combinations of the T-Loss with superpixels and/or iterative label refinement procedures are still to be explored.

The T-Loss provides a robust solution for binary segmentation of medical images in the presence of high levels of annotation noise, as frequently met in practice e.g. due to annotator fatigue or inter-annotator disagreements. This may be a key feature in achieving good generalization in many medical image segmentation applications, such as clinical decision support systems. Our evaluations and analyses provide evidence that the T-Loss is a reliable and valuable tool in the field of medical image analysis, with the potential for broader application in other domains.

## 6 Data use declaration and acknowledgment

We declare that we have used the ISIC dataset [5] under the Apache License 2.0, [publicly available](#), and the Shenzhen dataset [4, 13, 25] [public available](#) under the CC BY-NC-SA 4.0 License.

## References

- [1] Ludovic Amruthalingam et al. “Objective Hand Eczema Severity Assessment with Automated Lesion Anatomical Stratification”. In: *Experimental Dermatology* (Jan. 10, 2023), exd.14744. DOI: [10.1111/exd.14744](https://doi.org/10.1111/exd.14744).



- [2] Jonathan T Barron. “A general and adaptive robust loss function”. In: *CVPR*. 2019, pp. 4331–4339.
- [3] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Information Science and Statistics. 2006. 738 pp. ISBN: 978-0-387-31073-2.
- [4] Sema Candemir et al. “Lung Segmentation in Chest Radiographs Using Anatomical Atlases With Nonrigid Registration”. In: *TMI* 33.2 (Feb. 2014), pp. 577–590. DOI: [10.1109/tmi.2013.2290491](https://doi.org/10.1109/tmi.2013.2290491).
- [5] Noel C. F. Codella 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)”. In: *ISBI*. Apr. 2018, pp. 168–172. DOI: [10.1109/ISBI.2018.8363547](https://doi.org/10.1109/ISBI.2018.8363547).
- [6] Ncoza Dlova et al. “Prevalence of Skin Diseases Treated at Public Referral Hospitals in KwaZulu-Natal, South Africa”. In: *The British journal of dermatology* 178 (Mar. 31, 2017). DOI: [10.1111/bjd.15534](https://doi.org/10.1111/bjd.15534).
- [7] Florence Forbes and Darren Wraith. “A New Family of Multivariate Heavy-Tailed Distributions with Variable Marginal Amounts of Tailweight: Application to Robust Clustering”. In: *Stat Comput* 24.6 (Nov. 2014), pp. 971–984. DOI: [10.1007/s11222-013-9414-4](https://doi.org/10.1007/s11222-013-9414-4).
- [8] Aritra Ghosh, Himanshu Kumar, and P. S. Sastry. “Robust Loss Functions under Label Noise for Deep Neural Networks”. In: *AAAI*. Feb. 4, 2017, pp. 1919–1925.
- [9] Ellen R. Girden. *ANOVA: Repeated Measures*. Sage University Papers. Quantitative Applications in the Social Sciences no. 07-084. 1992. 77 pp. ISBN: 978-0-8039-4257-8.
- [10] Matthew Groh et al. “Evaluating Deep Neural Networks Trained on Clinical Images in Dermatology with the Fitzpatrick 17k Dataset”. In: *CVPRW*. June 2021, pp. 1820–1828. DOI: [10.1109/CVPRW53098.2021.00201](https://doi.org/10.1109/CVPRW53098.2021.00201).
- [11] Danna Gurari et al. “How to Collect Segmentations for Biomedical Images? A Benchmark Evaluating the Performance of Experts, Crowdsourced Non-experts, and Algorithms”. In: *WACV*. Jan. 2015, pp. 1169–1176. DOI: [10.1109/WACV.2015.160](https://doi.org/10.1109/WACV.2015.160).
- [12] Fabian Isensee et al. “nnU-Net: A Self-Configuring Method for Deep Learning-Based Biomedical Image Segmentation”. In: *Nat Methods* 18.2 (2 Feb. 2021), pp. 203–211. DOI: [10.1038/s41592-020-01008-z](https://doi.org/10.1038/s41592-020-01008-z).
- [13] Stefan Jaeger et al. “Automatic Tuberculosis Screening Using Chest Radiographs”. In: *TMI* 33.2 (), pp. 233–245. DOI: [10.1109/tmi.2013.2284099](https://doi.org/10.1109/tmi.2013.2284099).
- [14] H. J. Keselman and Joanne C. Rogan. “The Tukey Multiple Comparison Test: 1953–1976.” In: *Psychological Bulletin* 84.5 (Sept. 1977), pp. 1050–1056. DOI: [10.1037/0033-2909.84.5.1050](https://doi.org/10.1037/0033-2909.84.5.1050).
- [15] Shuailin Li, Zhitong Gao, and Xuming He. “Superpixel-Guided Iterative Learning from Noisy Labels for Medical Image Segmentation”. In: *MIC-CAI*. Vol. 12901. 2021, pp. 525–535. DOI: [10.1007/978-3-030-87193-2\\_50](https://doi.org/10.1007/978-3-030-87193-2_50).
- [16] Geert Litjens et al. “A Survey on Deep Learning in Medical Image Analysis”. In: *Medical Image Analysis* 42 (Dec. 2017), pp. 60–88. DOI: [10.1016/j.media.2017.07.005](https://doi.org/10.1016/j.media.2017.07.005).

- [17] Sheng Liu et al. “Adaptive Early-Learning Correction for Segmentation from Noisy Annotations”. In: *CVPR*. June 2022, pp. 2596–2606. DOI: [10.1109/CVPR52688.2022.00263](https://doi.org/10.1109/CVPR52688.2022.00263).
- [18] Xingjun Ma et al. “Normalized Loss Functions for Deep Learning with Noisy Labels”. In: *ICML*. Vol. 119. 2020, pp. 6543–6553.
- [19] Kevin P. Murphy. *Machine Learning: A Probabilistic Perspective*. Illustrated edition. Cambridge, MA, Aug. 24, 2012. 1104 pp. ISBN: 978-0-262-01802-9.
- [20] David F. Nettleton, Albert Orriols-Puig, and Albert Fornells. “A Study of the Effect of Different Types of Noise on the Precision of Supervised Learning Techniques”. In: *Artif Intell Rev* 33.4 (Apr. 1, 2010), pp. 275–306. DOI: [10.1007/s10462-010-9156-z](https://doi.org/10.1007/s10462-010-9156-z).
- [21] Giorgio Patrini et al. “Making Deep Neural Networks Robust to Label Noise: A Loss Correction Approach”. In: *CVPR*. July 2017, pp. 2233–2241. DOI: [10.1109/CVPR.2017.240](https://doi.org/10.1109/CVPR.2017.240).
- [22] M. Pechenizkiy et al. “Class Noise and Supervised Learning in Medical Domains: The Effect of Feature Extraction”. In: *CBMS*. June 2006, pp. 708–713. DOI: [10.1109/CBMS.2006.65](https://doi.org/10.1109/CBMS.2006.65).
- [23] Vinícius Ribeiro, Sandra Avila, and Eduardo Valle. “Handling Inter Annotator Agreement for Automated Skin Lesion Segmentation”. In: (2019).
- [24] Dinggang Shen, Guorong Wu, and Heung-Il Suk. “Deep Learning in Medical Image Analysis”. In: *Annu. Rev. Biomed. Eng.* 19.1 (June 21, 2017), pp. 221–248. DOI: [10.1146/annurev-bioeng-071516-044442](https://doi.org/10.1146/annurev-bioeng-071516-044442).
- [25] Sergii Stirenko et al. “Chest X-Ray Analysis of Tuberculosis by Deep Learning with Segmentation and Augmentation”. In: *ELNANO*. Apr. 2018. DOI: [10.1109/elnano.2018.8477564](https://doi.org/10.1109/elnano.2018.8477564).
- [26] Jianyong Sun, Ata Kabán, and Jonathan M. Garibaldi. “Robust Mixture Clustering Using Pearson Type VII Distribution”. In: *Pattern Recognition Letters* 31.16 (Dec. 2010), pp. 2447–2454. DOI: [10.1016/j.patrec.2010.07.015](https://doi.org/10.1016/j.patrec.2010.07.015).
- [27] Tong Xiao et al. “Learning from Massive Noisy Labeled Data for Image Classification”. In: *CVPR*. June 2015, pp. 2691–2699. DOI: [10.1109/CVPR.2015.7298885](https://doi.org/10.1109/CVPR.2015.7298885).
- [28] Andreas Veit et al. “Learning from Noisy Large-Scale Datasets with Minimal Supervision”. In: *CVPR*. July 2017, pp. 6575–6583. DOI: [10.1109/CVPR.2017.696](https://doi.org/10.1109/CVPR.2017.696).
- [29] Guotai Wang et al. “A Noise-Robust Framework for Automatic Segmentation of COVID-19 Pneumonia Lesions From CT Images”. In: *TMI* 39.8 (Aug. 2020), pp. 2653–2663. DOI: [10.1109/TMI.2020.3000314](https://doi.org/10.1109/TMI.2020.3000314). PMID: [32730215](https://pubmed.ncbi.nlm.nih.gov/32730215/).
- [30] Yisen Wang et al. “Symmetric Cross Entropy for Robust Learning With Noisy Labels”. In: *ICCV*. Oct. 2019, pp. 322–330. DOI: [10.1109/ICCV.2019.00041](https://doi.org/10.1109/ICCV.2019.00041).

- [31] Ke Yan et al. “DeepLesion: Automated Mining of Large-Scale Lesion Annotations and Universal Lesion Detection with Deep Learning”. In: *J. Med. Imag.* 5.03 (July 20, 2018), p. 1. DOI: [10.1117/1.JMI.5.3.036501](https://doi.org/10.1117/1.JMI.5.3.036501).
- [32] Shuo Yang et al. “Estimating Instance-dependent Bayes-label Transition Matrix Using a Deep Neural Network”. In: *ICML*. June 28, 2022, pp. 25302–25312.
- [33] Yu Yao et al. “Dual T: Reducing Estimation Error for Transition Matrix in Label-noise Learning”. In: *NeurIPS*. Vol. 33. 2020, pp. 7260–7271.
- [34] Zhilu Zhang and Mert Sabuncu. “Generalized Cross Entropy Loss for Training Deep Neural Networks with Noisy Labels”. In: *NeurIPS*. Vol. 31. 2018.
- [35] Xingquan Zhu and Xindong Wu. “Class Noise vs. Attribute Noise: A Quantitative Study”. In: *Artif Intell Rev* 22.3 (Nov. 2004), pp. 177–210. DOI: [10.1007/s10462-004-0751-8](https://doi.org/10.1007/s10462-004-0751-8).