

Robust Asymmetric Loss for Multi-Label Long-Tailed Learning

Wongi Park^{1*}, Inhyuk Park^{2*}, Sungeun Kim², and Jongbin Ryu^{1,2†}

¹ Department of Software and Computer Engineering, Ajou University

² Department of Artificial Intelligence, Ajou University

{psboys, inhyuk, kimsungeun, jongbinryu}@ajou.ac.kr

Abstract

In real medical data, training samples typically show long-tailed distributions with multiple labels. Class distribution of the medical data has a long-tailed shape, in which the incidence of different diseases is quite varied, and at the same time, it is not unusual for images taken from symptomatic patients to be multi-label diseases. Therefore, in this paper, we concurrently address these two issues by putting forth a robust asymmetric loss on the polynomial function. Since our loss tackles both long-tailed and multi-label classification problems simultaneously, it leads to a complex design of the loss function with a large number of hyper-parameters. Although a model can be highly fine-tuned due to a large number of hyper-parameters, it is difficult to optimize all hyper-parameters at the same time, and there might be a risk of overfitting a model. Therefore, we regularize the loss function using the Hill loss approach, which is beneficial to be less sensitive against the numerous hyper-parameters so that it reduces the risk of overfitting the model. For this reason, the proposed loss is a generic method that can be applied to most medical image classification tasks and does not make the training process more time-consuming. We demonstrate that the proposed robust asymmetric loss performs favorably against the long-tailed with multi-label medical image classification in addition to the various long-tailed single-label datasets. Notably, our method achieves Top-5 results on the CXR-LT dataset of the ICCV CVAMD 2023 competition. We open-source our implementation of the robust asymmetric loss in the public repository: <https://github.com/kalelpark/RALoss>.

1. Introduction

Multi-label classification, which predicts more than one label from a single image, has received lots of interest in recent years. Especially in the field of medical image recogni-

*Equal contribution.

†Corresponding author.

CXR-LT : Long-tailed Distribution

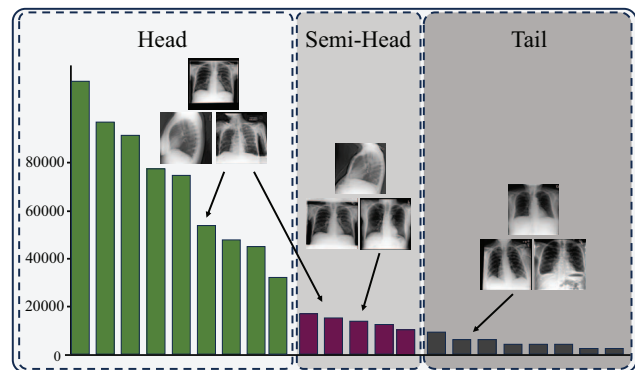


Figure 1: Label distribution of the CXR-LT Dataset[17]. Typical radiological images have a lengthy tail distribution since there are few positive samples for certain classes. Further single radiological image contains multiple classes in most cases. These long-tailed and multi-label data is common but critical issues in real-world medical image recognition task.

tion, several studies [1, 27, 43, 32] have been conducted to tackle the problem of the coexistence of multiple symptoms in a single radiology image. However, these multi-label classification studies have overlooked another critical issue: the long-tailed distribution of medical data. [7, 5, 43, 18] That is the long-tailed distribution problem. In general, multi-label data, the more classes there are used, the more long-tailed the distribution. In this case, for classes with fewer labels (tail labels), the performance of the model will drop significantly, and the model will be biased to the data of head labels with more training data. Therefore, it is challenging to generalize the learned model in practice. To solve this imbalance data distribution, there have been studies that re-sample[2, 30, 35] or re-weight[39, 19, 12] the data to make the model learn more from the tail labels, but these studies have not been addressed the long-tailed problem in a multi-label classification environment.

Another issue is that many studies[38, 15, 37] exploit ad-

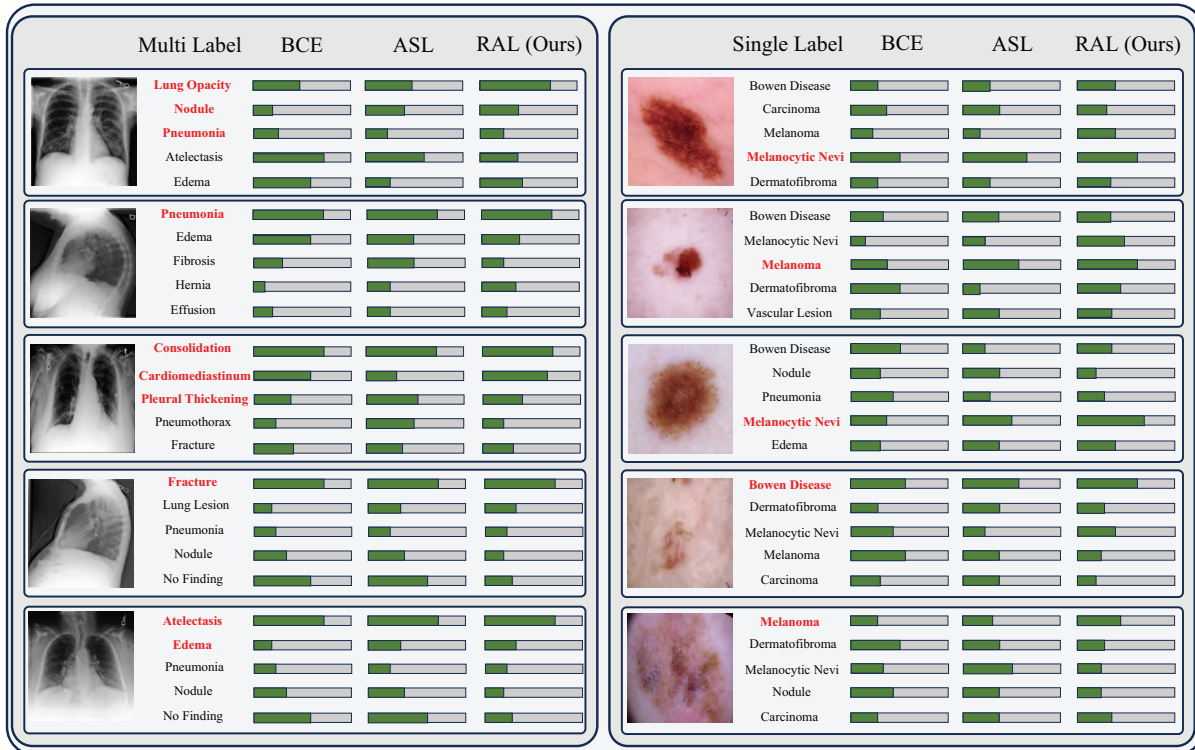


Figure 2: Examples of estimated probabilities of the BCE, ASL, and our RAL from the CXR-LT[17] and ISIC2018[10] datasets. Example result of our proposed methods, BCE and ASL[4] from CXR-LT Dataset[17] and ISIC2018 Dataset[10]. The color red denotes the positive labels. The model trained by BCE loss has a tendency to overfit the single label. On the other hand, models trained by ours exhibit higher probabilities for multiple labels, implying that ours are more reliable for the long-tailed multi-label classification task.

ditional resources to tackle the long-tailed distribution and multi-label classification problems. Using larger models or increasing computational complexity can help with these problems, but conversely, larger models can also be limited by high budgets that make them less applicable in practice. Therefore, in this paper, we propose a robust asymmetric loss that does not require additional resources to learn the multi-label long-tailed medical data. The proposed loss is based on asymmetric weighting, which ensures that the importance of the negative sample’s loss is regarded differently from that of positive samples so that even hard negative samples can be robustly learned. Compared to the existing cross-entropy, focal loss[23], asymmetric loss[4], and balanced loss[11], the proposed robust asymmetric focal loss effectively learns long-tailed multi-label data reliably while being less sensitive to hyper-parameters. Specifically, we re-weight the negative samples adopting the hill loss[42] so that ours performs favorably against the hard negative samples while being robust to the settings on the variety of hyper-parameters. To this end, we expand the asymmetric loss using a Taylor series-based approach[20] to account for the negative loss. The Taylor series ensures

that negative samples below a certain threshold are not used for training so that the stable gradient can be passed to the deep neural networks.

We evaluate the proposed method on CXR-LT, a long-tailed multi-label classification medical dataset, and demonstrate that ours improve the performance of the classification task over existing methods. Notably, our method achieved 0.351 mAP, which is within the Top-5 of the final ranking in the ICCV CVAMD 2023 competition. Furthermore, we evaluate that our robust asymmetric loss works well on long-tailed distributions, even on single-label medical datasets. For this evaluation, we utilize the ISIC2018 and APTOS2019 datasets and show that our method achieves considerably better performance compared to the existing methods.

We present the contributions of this paper as follows.

- We propose robust asymmetric loss, which is effective for long-tailed multi-label classification. The proposed loss can be finely tuned but is not sensitive to hyper-parameter settings.
- We improve the performance of long-tailed multi-label




Figure 3: Grad-CAM visualization [34] from the model trained by the proposed RAL.

classification without additional training data, model parameters, and computational budget.

- We achieved Top-5 results in the CVAMD2023 competition on the long-tailed multi-label CXR-LT dataset. In addition, we confirm that the proposed method works well on single-label medical image classification as well as the multi-label dataset.

2. Related Work

Multi-Label classification has been extensively studied to predict more than one class label[25, 16]. Recently, to understand the correlation between multiple labels, several studies have introduced network architectures that enable the model to predict the inherent relation between features and corresponding labels. Most of the studies have used Graph Convolution Network(GCN)[9, 40, 8] that learn label’s feature relation in the graph structure. Subsequently, semantic representation in images using attention mechanism has been researched in many studies[22, 36]. On the other hand, there have been training algorithms to tackle the multi-label classification, such as investigation on the weight re-weighting and class frequency[11, 29, 33]. More recently, asymmetric loss[4, 20] has been introduced to optimize the imbalanced positive and negative losses.

Long-tailed distribution has been regarded as a practical problem for real-world machine learning applications. Studies that address the problem of the long-tailed distribution can be divided into two categories. First, re-sampling methods[41, 21, 14] have been introduced to under-sample or over-sample the data according to its class distribution to construct the balanced training set. Re-weighting strategy assigns different weights to the samples to adjust the long-tailed distribution[23, 11]. However, there is an ambiguity in applying the re-sampling methods to multi-label datasets.

When a single image contains both head and tail class labels, it is difficult to determine whether the sample should be over- or under-sampled. For this reason, re-sampling methods are hardly applicable to the multi-label classification so we encounter the problem that most gradients are computed from negative samples. To address this issue, the second category, studies[13, 29, 33] for dealing with the long-tailed class distribution exploiting the loss functions, are being researched. The focal loss[23] is the landmark method to tackle the long-tailed class distribution using the loss function. Focal loss adds the modulating factors(*i.e.* focusing and balance parameters) to the cross-entropy loss so that it can mitigate the long-tailed distribution by controlling such modulating factors. Then, to further tailor the loss function, asymmetric loss[20, 31] was proposed, which determines the focusing parameters of negative and positive loss separately. Recently, the loss function has been expanded to polynomial functions[20] to use only several principle terms in computing the gradients. On the other hand, the Hill loss[42] was proposed to prevent the gradient from being too large in certain samples. The Hill loss performed well for multi-label classification, but they did not apply their method to complex formulas unfolded as polynomial. Therefore, in this paper, we apply the Hill term loss to the formula extended to polynomials to learn a long-tailed multi-label classification model more accurately and robustly.

3. Method

In this section, we introduce our robust asymmetric loss function. We first describe the existing long-tailed and multi-class losses as the background of our loss function. Then, we introduce the robust asymmetric loss function by adding the Hill loss term to the polynomial function.

Label	Positive		Label	Positive	
	#Sample(K)	Portion(%)		#Sample(K)	Portion(%)
Atelectasis	67.6	10.6	Mass	5.5	0.9
Calcification	4.3	0.7	No Finding	41.8	6.6
Cardiomegaly	76.9	12.1	Nodule	7.6	1.2
Consolidation	16.0	2.5	Pleural Effusion	69.2	10.8
Edema	38.6	6.1	Pleural Other	0.6	0.1
Emphysema	4.3	0.7	Pleural Thickening	3.3	0.5
Cardiomediastinum	30.1	4.7	Pneumomediastinum	0.7	0.1
Fibrosis	1.1	0.2	Pneumonia	49.1	7.6
Fracture	11.9	1.9	Pneumoperitoneum	0.5	0.1
Hernia	4.0	0.6	Pneumothorax	14.9	2.4
Infiltration	10.2	1.6	Emphysema	2.4	0.4
Lung Lesion	2.5	0.4	Support Devices	89.1	14
Lung Opacity	79.9	12.6	Tortuous Aorta	3.4	0.6

Table 1: Specification of the CXR-LT dataset. It shows that the samples are heavily distributed in a few classes while several classes have very few samples.

3.1. Long-tailed and Multi-label Classification Loss

Traditionally, multi-label classification tasks use the Binary Cross-Entropy (BCE) Loss as:

$$\mathcal{L}_{BCE} = - \sum_{i=1}^K (y_i L_i^+ + (1 - y_i) L_i^-) \quad (1)$$

$$\begin{cases} \mathcal{L}^+ = \log(\hat{y}) \\ \mathcal{L}^- = \log(1 - \hat{y}) \end{cases}, \quad (2)$$

where L_i^+ and L_i^- are positive and negative sample losses and y and \hat{y} denote the ground-truth and estimated probability for the class labels.

However, since this BCE function computes the same weights for all class samples in training data with the long-tailed distribution, it excessively focuses on learning the head classes with a large number of training samples. This problem is addressed by the focal loss[23] \mathcal{L}_{Focal} with balancing the positive and negative losses as:

$$\begin{cases} \mathcal{L}_{Focal}^+ = \alpha_+ (1 - \hat{y})^\gamma \log(\hat{y}) \\ \mathcal{L}_{Focal}^- = \alpha_- \hat{y}^\gamma \log(1 - \hat{y}) \end{cases}, \quad (3)$$

where α_+ and α_- represent the balancing parameter and γ denotes the focusing parameter that is the key hyper-parameters of the focal loss function. Controlling the hyper-parameters, the focal loss can balance the head- and tail-class samples. However, this focal loss has a weakness in that positive and negative losses share the same focusing parameter γ . Therefore asymmetric weighting approach[31]

for the loss function alleviates this problem by assigning different focusing parameters as:

$$\begin{cases} \mathcal{L}_{ASL}^+ = (1 - \hat{y})^{\gamma^+} \log(\hat{y}) \\ \mathcal{L}_{ASL}^- = \hat{y}^{\gamma^-} \log(1 - \hat{y}) \end{cases} \quad (4)$$

$$\hat{y}_\tau = \max(\hat{y} - \tau, 0),$$

where γ^+ and γ^- are the positive and negative focusing parameters and y_τ denotes the rectified probability thresholded by τ . This ASymmetric Loss (ASL) is efficient for optimizing the training of positive and negative samples separately and able to mitigate the gradient vanishing problem due to too small a value of \hat{y} in the negative loss.

3.2. Robust Asymmetric Loss

Asymmetric loss can be expanded to the polynomial equation using the Taylor series[24] based method[20]. In a polynomial equation, using several principle low-order terms can improve the performance of multi-label classification tasks. This is because the higher-order terms in the polynomial form can be regarded as noise or redundant, so using only a few low-order terms is effective. Therefore, the asymmetric polynomial loss is formulated as follows:

$$\begin{cases} \mathcal{L}_{APL}^+ = y \sum_{m=1}^M \alpha_m (1 - \hat{y})^{m+\gamma^+} \\ \mathcal{L}_{APL}^- = (1 - y) \sum_{n=1}^N \beta_n \hat{y}_\tau^{n+\gamma^-} \end{cases}, \quad (5)$$

where M and N are parameters that determine the number of low-order terms to be used in the positive and negative losses, and α_m and β_n stand for the balance parameter of each term in the positive and negative losses. This

Asymmetric Polynomial Loss (APL)[20] has the advantage of controlling the positive and negative losses on a term-by-term basis, but it also has the significant drawback of requiring a large number of hyper-parameters to be configured by the user. Optimizing such a large number of hyper-parameters can be a time-consuming process and often leads to overfitting the models.

To be less sensitive to optimizing the numerous hyper-parameters, we introduce robust asymmetric loss. It is noticeable that, especially in multi-label data, the number of negative samples is much larger than that of positives, so making the negative loss less sensitive is the most decisive factor in the long-tailed multi-label classification task. Therefore, we adopt the Hill loss[42] so that we prevent an excessively large gradient of the negative loss in the learning process. Adding the Hill loss term to APL, we define our Robust Asymmetric Loss (RAL) as:

$$\begin{cases} \mathcal{L}_{RAL}^+ = y \sum_{m=1}^M \alpha_m (1 - \hat{y})^{m+\gamma^+} \\ \mathcal{L}_{RAL}^- = \psi(\hat{y}) \cdot (1 - y) \sum_{n=1}^N \beta_n \hat{y}_\tau^{n+\gamma^-} \end{cases} \quad (6)$$

$$\psi(\hat{y}) = \lambda - \hat{y},$$

where ψ denotes the Hill loss term and λ is set to 1.5 value. Our RAL is robust to the change of numerous hyper-parameters due to the less sensitive negative loss in the training process. In the negative loss, when \hat{y} is close to 0, that is, the estimated probability of the training data is close to the correct negative answer, the gradient value is already small. Therefore, in this case, the hyper-parameter is not sensitive. On the other hand, when the estimated probability is around 0, which is a hard negative sample, the gradient value is too large, making the network training sensitive to the hyper-parameter settings. Our RAL loss regularizes the gradient of these hard negative samples to make them less sensitive to hyper-parameters. As we expand the asymmetric loss to polynomial form, there is an unavoidable problem of setting too many hyper-parameter, so we propose RAL with Hill loss term to alleviate such a problem.

Dataset	classes	Samples	Imbalance Ratio
CXR-LT	26	377,110	142
APTOS2019	7	10,015	58
ISIC2018	5	3,662	10

Table 2: The details of long-tailed medical datasets.

4. Experiments Setup

4.1. Dataset and Metrics

CXR-LT. The 377,110 CXRs in the ICCV CVAMD 2023 Dataset(CXR-LT) dataset, which is included in the competition, have at least one label in 26 clinical findings. The

Method	Image size	mAP	mAUC	mF1
CE	224	0.301	0.808	0.218
	384	0.314	0.813	0.227
Focal loss[23]	224	0.304	0.807	0.231
	384	0.295	0.803	0.224
ASL[4]	224	0.307	0.808	0.225
	384	0.317	0.811	0.237
RAL (Ours)	224	0.314	0.815	0.225
	384	0.323	0.817	0.233

Table 3: Experimental comparison on the loss functions with ours.

class labels consist of the "No Finding" class and 12 new disease labels introduced from mimic-cxr-jpg, which cover chest X-rays(CXR). The detailed specification of the CXR-LT dataset can be found in Table 1. We randomly divide the image sets into a test and training set with a ratio of 8:2 for the CXR-LT dataset.

ISIC2018 and APTOS2019. The ISISC2018 dataset has 10,015 skin images with 7 lesion classes, and the APTOS dataset includes 3,662 diabetic retinopathy images with 5 disease classes. For these APTOS2019 and ISIC2018 datasets, we follow the same protocol of the previous study[28].

Metric. To evaluate our methods for the CXR-LT dataset, we use three metrics such as mean Average Precision(mAP), mean Area Under Curve(mAUC), and F1-Score considering the multi-label dataset. For APTOS2019 and ISIC2018, which is the single-label dataset, we use two metrics as Accuracy and F1-Score. The details of these three datasets are specified in Table 2. The imbalance ratio for measuring the significance of the long-tailed distribution is denoted as N_{max}/N_{min} , where N is the number of each class sample in each class.

4.2. Implementation details.

We use the ConvNeXT-B[26] as the backbone for the proposed loss. We resize the input images as 384×384 and exploit the data augmentation schemes following the previous[3, 9]. We train our networks using the Adam optimizer with 0.9 momentum and 0.001 weight decay. The batch size is 256, and the initial learning rate is set to $1e-4$. Our networks are trained on PyTorch version 1.11.0 with RTX A5000 GPUs.

Label	BCE	APL	RAL (Ours)	Label	BCE	APL	RAL (Ours)
Atelectasis	0.578	0.599	0.610	Mass	0.159	0.200	0.222
Calcification	0.120	0.137	0.151	No Finding	0.445	0.477	0.479
Cardiomegaly	0.626	0.633	0.648	Nodule	0.148	0.205	0.234
Consolidation	0.203	0.209	0.224	Pleural Effusion	0.801	0.813	0.821
Edema	0.527	0.552	0.562	Pleural Other	0.015	0.048	0.059
Emphysema	0.258	0.313	0.334	Pleural Thickening	0.065	0.094	0.111
Cardiomediastinum	0.155	0.166	0.173	Pneumomediastinum	0.103	0.138	0.203
Fibrosis	0.099	0.121	0.131	Pneumonia	0.289	0.306	0.167
Fracture	0.175	0.223	0.270	Pneumoperitoneum	0.134	0.143	0.524
Hernia	0.483	0.509	0.560	Pneumothorax	0.394	0.478	0.483
Infiltration	0.058	0.061	0.075	Emphysema	0.391	0.459	0.544
Lung Lesion	0.054	0.059	0.079	Support Devices	0.892	0.906	0.913
Lung Opacity	0.579	0.601	0.613	Tortuous Aorta	0.055	0.052	0.056

Table 4: Experimental results of the development phase of the CVAMD 2023 competition. Our RAL works well on most cases compared to the other loss functions.

5. Experimental results

In this section, we show the experimental results to validate the effectiveness of RAL. We first compare the proposed RAL with previous state-of-the-art loss functions such as focal loss, LDAM, and ASL. We then dissect the proposed loss function into its component level to demonstrate its robustness. In this experiment, RAL performs well consistently for variations of numerous hyper-parameters. We also show that the proposed RAL works favorably on both multi- and single-label long-tailed medical image classification tasks. Further, we validate that ours is robust to several noisy conditions.

Methods	ICIS2018		APTOS2019	
	Accuracy	F1-score	Accuracy	F1-score
CE	0.850	0.716	0.812	0.608
Focal loss[23]	0.861	0.735	0.815	0.629
LDAM[6]	0.849	0.728	0.813	0.620
ASL [†] [4]	0.854	0.734	0.820	0.660
RAL (Ours)	0.852	0.740	0.826	0.673

Table 5: Experimental results of the ICIS2018 and APTOS2019 datasets. [†] denotes the result from our implementation with the official code: <https://github.com/Alibaba-MIIL/ASL>.

5.1. Comparison on Loss Functions

In this subsection, we compare our RAL with other loss functions on three datasets such as ISIC2018, APTOS2019, and CXR-LT. Table 3 and 4 show the result submitted to the CVAMD 2023 competition site¹ using our RAL at the development phase. In this result, our RAL achieves competitive performance compared to others. Ours performs well on most classes consistently in Table 4. Further, our proposed RAL works well on single-label long-tailed datasets, such as ISIC2018 and APTOS2019. It outperforms the other methods in such datasets in Table 5. These findings of the experimental results highlight the competitiveness of our RAL in diverse long-tailed medical image classification datasets.

5.2. Ablation Study

For a more in-depth analysis of the proposed method, we broke RAL into three components in our ablation study. The three components are focal, asymmetric, and Hill loss where we apply the polynomial expansion to the Hill loss. All results of this ablation study are taken using the ConvNeXt-B model with 384×384 image size in Table 6. Through this ablation study, we demonstrate that each component of our RAL is effective for the long-tailed multi-label classification task.

¹https://codalab.lisn.upsaclay.fr/competitions/12599#learn_the_details




Figure 4: Experimental result on the evaluation of polynomial coefficient. The Y-axis F1 Diff shows the difference from the best F1 score for each method (*i.e.* APL and our RAL). Therefore, a value less than 0 indicates a larger difference from the best result and is sensitive to the polynomial coefficient hyper-parameter. In this result, RAL is less sensitive against the APL considerably.

Focal Loss	Asymmetric	Hill	mAP	mAUC
✓	✗	✗	0.295	0.803
✓	✓	✗	0.307	0.815
✓	✓	✓	0.323	0.817

Table 6: Experimental result of the ablation study of the proposed RAL.




Figure 5: Examples of noisy images such as Gaussian Blur, Speckle, and Salt-Pepper.

5.3. Robustness Analysis

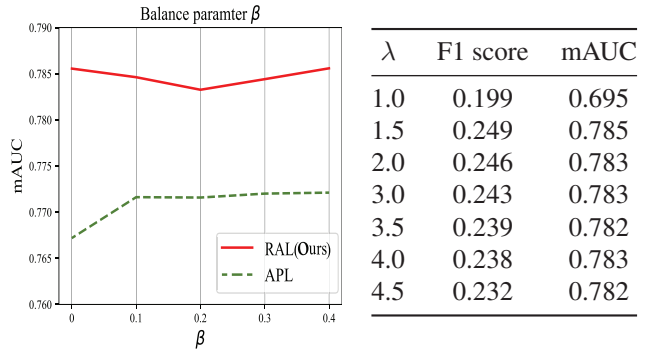
We carry out further experiments to investigate the robustness of RAL. In our experiment, we introduce Gaussian Blur, Salt-Pepper, and Speckle noise to the original images of the CXR-LT dataset, as shown in Fig.5. To validate that our method performs well even with noisy conditions, we compare our method’s mAUC to that of Binary Cross-Entropy(BCE) Loss and ASymmetric Loss(ASL). Under the noisy condition, our method outperforms the others by about 1 3% as shown in Table 8. Therefore, our RAL has been empirically demonstrated to be more robust than existing methods when images are impacted by noisy conditions.

5.4. Performance Analysis on Hyper-parameters

We conduct more experiments for the performance analysis on hyper-parameter settings used in RAL. We compare

ours with Asymmetric Polynomial Loss (APL)[20] in this experiment. We employ the same hyper-parameters for ours and APL.

First, we adjust the polynomial coefficient value of α_m in Eq.6 for our RAL and APL. Figure 4 shows that, across all three datasets, our RAL is generally less sensitive to changes in the polynomial coefficient than APL, leading to less variance in the performance.



(a) Effect of parameter β .

(b) Effect of parameter λ .

Figure 6: Experimental result of ours according to the hyper-parameters β and λ . In (a), we conduct an evaluation to compare our RAL with APL with regard to β in Eq.6, the weight that regularizes the negative loss. In (b), we evaluate F1 and AUC scores in relation to λ in Eq.6, which is utilized to avoid significant gradients in the negative loss.

Furthermore, we evaluate the balance parameter β , which governs the negative loss. Fig.6(a) shows that for all the β values, our RAL outperforms APL in terms of AUC score, highlighting the resilience to the hyper-parameter settings of our RAL. In addition, we experimented with the balance parameter β , which controls the negative loss. Fig.6(b) shows the F1 and AUC score for different λ val-

Image size	1024	1512	2048	Image size	1024	1512	2048
Label	# 1	# 2	# 3	Label	# 1	# 2	# 3
Atelectasis	0.607	0.576	0.574	Mass	0.206	0.127	0.128
Calcification	0.142	0.111	0.115	No Finding	0.478	0.452	0.452
Cardiomegaly	0.648	0.610	0.610	Nodule	0.200	0.184	0.182
Consolidation	0.218	0.182	0.182	Pleural Effusion	0.831	0.814	0.815
Edema	0.555	0.518	0.517	Pleural Other	0.039	0.032	0.033
Emphysema	0.193	0.174	0.176	Pleural Thickening	0.109	0.082	0.083
Cardiomediastinum	0.184	0.164	0.164	Pneumomediastinum	0.338	0.305	0.313
Fibrosis	0.153	0.119	0.122	Pneumonia	0.309	0.284	0.286
Fracture	0.289	0.234	0.237	Pneumoperitoneum	0.282	0.239	0.230
Hernia	0.550	0.397	0.402	Pneumothorax	0.552	0.507	0.507
Infiltration	0.060	0.056	0.055	Emphysema	0.560	0.545	0.544
Lung Lesion	0.038	0.028	0.029	Support Devices	0.913	0.896	0.896
Lung Opacity	0.596	0.556	0.556	Tortuous Aorta	0.060	0.049	0.049

Table 7: Experimental result of test phase of the CVAMD 2023 competition. It shows that image size of 1024×1024 achieves the best result. We assume that this result is because the resolution we used for training is 1024×1024 , which is not very large.

Methods	Img	Blur	Speckle	SaltPepper
BCE	0.789	0.718	0.502	0.501
ASL	0.791	0.734	0.512	0.513
RAL (Ours)	0.796	0.547	0.534	0.745

Table 8: Experimental result on the noisy conditions. Our RAL shows better performance compared to the others consistently.

ues; it can be found that the best performance is obtained at $\lambda = 1.5$ consistent with the previous study[42].

Submit	mAP	mAUC	mF1
# 1	0.351	0.837	0.256
# 2	0.317	0.814	0.061
# 3	0.318	0.814	0.143

Table 9: Experimental result of three submissions of test phase of the CVAMD 2023 competition.

6. Result on CVAMD2023 Competition

Our method results in the Top-5 of the ICCV CVAMD 2023 competition’s final rankings. To get this result, we scale the input image to 1024×1024 and use ConvNeXT-B models from [26]. We increase the input image size in the test phase using the checkpoint file saved during the development phase. We configure hyper-parameters in the same

values as Sec.4.2. In the final score of the competition, ours recorded 0.351 mAP, 0.837 mAUC, and 0.256 mF1 scores, which are included in the Top-5 ranking. With the efficient loss function design, we show improved performance on the multi-label long-tailed classification of the CVAMD 2023 challenge, which does not use additional model parameters or inference complexity. Table 9 shows the test phase results of our three submissions.

7. Conclusion

In this paper, we introduce the Robust Asymmetric Loss (RAL) for long-tailed multi-label classification tasks on medical images. Our proposed RAL trains the model more robustly against the various hyper-parameters without additional resources. RAL shows competitive results on the long-tailed single- and multi-label datasets compared to previous state-of-the-art loss functions. We especially achieve a Top-5 ranking in the CVAMD 2023 competition using our method. We think that future research can benefit from our findings and incorporate ours into their work.

Acknowledgement. This research was supported in part by the MSIT(Ministry of Science and ICT), Korea, under the National Program for Excellence in SW(2022-0-01077) supervised by the IITP(Institute of Information & communications Technology Planning & Evaluation) in 2023 and by a grant of ‘Korea Government Grant Program for Education and Research in Medical AI’ through the Korea Health Industry Development Institute(KHIDI), funded by the Korea government(MOE, MOHW).

References

- [1] Nkechinyere N Agu, Joy T Wu, Hanqing Chao, Ismini Lourentzou, Arjun Sharma, Mehdi Moradi, Pingkun Yan, and James Hendler. Anaxnet: anatomy aware multi-label finding classification in chest x-ray. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part V 24*, pages 804–813. Springer, 2021.
- [2] Shin Ando and Chun Yuan Huang. Deep over-sampling framework for classifying imbalanced data. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2017, Skopje, Macedonia, September 18–22, 2017, Proceedings, Part I 10*, pages 770–785. Springer, 2017.
- [3] Shekoofeh Azizi, Basil Mustafa, Fiona Ryan, Zachary Beaver, Jan Freyberg, Jonathan Deaton, Aaron Loh, Alan Karthikesalingam, Simon Kornblith, Ting Chen, et al. Big self-supervised models advance medical image classification. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3478–3488, 2021.
- [4] Emanuel Ben-Baruch, Tal Ridnik, Nadav Zamir, Asaf Noy, Itamar Friedman, Matan Protter, and Lih Zelnik-Manor. Asymmetric loss for multi-label classification. *arXiv preprint arXiv:2009.14119*, 2020.
- [5] Dipkamal Bhusal and Dr Sanjeeb Prasad Panday. Multi-label classification of thoracic diseases using dense convolutional network on chest radiographs. *arXiv preprint arXiv:2202.03583*, 2022.
- [6] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. *Advances in neural information processing systems*, 32, 2019.
- [7] Haomin Chen, Shun Miao, Daguang Xu, Gregory D Hager, and Adam P Harrison. Deep hierarchical multi-label classification of chest x-ray images. In *International conference on medical imaging with deep learning*, pages 109–120. PMLR, 2019.
- [8] Zhaomin Chen, Xiu-Shen Wei, Peng Wang, and Yanwen Guo. Learning graph convolutional networks for multi-label recognition and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [9] Zhao-Min Chen, Xiu-Shen Wei, Peng Wang, and Yanwen Guo. Multi-label image recognition with graph convolutional networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5177–5186, 2019.
- [10] Noel Codella, Veronica Rotemberg, Philipp Tschandl, M Emre Celebi, Stephen Dusza, David Gutman, Brian Helba, Aadi Kalloo, Konstantinos Liopyris, Michael Marchetti, et al. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). *arXiv preprint arXiv:1902.03368*, 2019.
- [11] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9268–9277, 2019.
- [12] Qi Dong, Shaogang Gong, and Xiatian Zhu. Imbalanced deep learning by minority class incremental rectification. *IEEE transactions on pattern analysis and machine intelligence*, 41(6):1367–1381, 2018.
- [13] Jie Du, Xiaoci Zhang, Peng Liu, Chi-Man Vong, and Tianfu Wang. An adaptive deep metric learning loss function for class-imbalance learning via intraclass diversity and inter-class distillation. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [14] Adrian Galdran, Gustavo Carneiro, and Miguel A González Ballester. Balanced-mixup for highly imbalanced medical image classification. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part V 24*, pages 323–333. Springer, 2021.
- [15] Zongyuan Ge, Dwarikanath Mahapatra, Xiaojun Chang, Zetao Chen, Lianhua Chi, and Huimin Lu. Improving multi-label chest x-ray disease diagnosis by exploiting disease and health labels dependencies. *Multimedia Tools and Applications*, 79:14889–14902, 2020.
- [16] Meng Han, Hongxin Wu, Zhiqiang Chen, Muhang Li, and Xilong Zhang. A survey of multi-label classification based on supervised and semi-supervised learning. *International Journal of Machine Learning and Cybernetics*, 14(3):697–724, 2023.
- [17] Gregory Holste, Song Wang, Ajay Jaiswal, Yuzhe Yang, Mingquan Lin, Yifan Peng, and Atlas Wang. Cxr-It: Multi-label long-tailed classification on chest x-rays.
- [18] Gregory Holste, Song Wang, Ziyu Jiang, Thomas C Shen, George Shih, Ronald M Summers, Yifan Peng, and Zhangyang Wang. Long-tailed classification of thorax diseases on chest x-ray: A new benchmark study. In *MICCAI Workshop on Data Augmentation, Labelling, and Imperfections*, pages 22–32. Springer, 2022.
- [19] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Deep imbalanced learning for face recognition and attribute prediction. *IEEE transactions on pattern analysis and machine intelligence*, 42(11):2781–2794, 2019.
- [20] Yusheng Huang, Jiexing Qi, Xinbing Wang, and Zhouhan Lin. Asymmetric polynomial loss for multi-label classification. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023.
- [21] Lie Ju, Yicheng Wu, Lin Wang, Zhen Yu, Xin Zhao, Xin Wang, Paul Bonnington, and Zongyuan Ge. Flexible sampling for long-tailed skin lesion classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 462–471. Springer, 2022.
- [22] Elizaveta Kovtun, Galina Boeva, Artem Zabolotnyi, Evgeny Burnaev, Martin Spindler, and Alexey Zaytsev. Label attention network for sequential multi-label classification: you were looking at a wrong self-attention. *CoRR*, 2023.
- [23] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Pro-*

- ceedings of the IEEE international conference on computer vision, pages 2980–2988, 2017.
- [24] Seppo Linnainmaa. Taylor expansion of the accumulated rounding error. *BIT Numerical Mathematics*, 16(2):146–160, 1976.
- [25] Weiwei Liu, Haobo Wang, Xiaobo Shen, and Ivor W Tsang. The emerging trends of multi-label learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(11):7955–7974, 2021.
- [26] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022.
- [27] Eman A Abdel Maksoud, Sherif Barakat, and Mohammed Elmogy. Medical images analysis based on multilabel classification. In *Machine Learning in Bio-Signal Analysis and Diagnostic Imaging*, pages 209–245. Elsevier, 2019.
- [28] Yassine Marrakchi, Osama Makansi, and Thomas Brox. Fighting class imbalance with contrastive learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 466–476. Springer, 2021.
- [29] Seulki Park, Jongin Lim, Younghan Jeon, and Jin Young Choi. Influence-balanced loss for imbalanced visual classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 735–744, 2021.
- [30] Samira Pouyanfar, Yudong Tao, Anup Mohan, Haiman Tian, Ahmed S Kaseb, Kent Gauen, Ryan Dailey, Sarah Aghajanzadeh, Yung-Hsiang Lu, Shu-Ching Chen, et al. Dynamic sampling in convolutional neural networks for imbalanced data classification. In *2018 IEEE conference on multimedia information processing and retrieval (MIPR)*, pages 112–117. IEEE, 2018.
- [31] Tal Ridnik, Emanuel Ben-Baruch, Nadav Zamir, Asaf Noy, Itamar Friedman, Matan Protter, and Lihi Zelnik-Manor. Asymmetric loss for multi-label classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 82–91, 2021.
- [32] Ario Sadafi, Niklas Koehler, Asya Makhro, Anna Bogdanova, Nassir Navab, Carsten Marr, and Tingying Peng. Multiclass deep active learning for detecting red blood cell subtypes in brightfield microscopy. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part I 22*, pages 685–693. Springer, 2019.
- [33] Dvir Samuel and Gal Chechik. Distributional robustness loss for long-tail learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9495–9504, 2021.
- [34] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [35] Li Shen, Zhouchen Lin, and Qingming Huang. Relay back-propagation for effective learning of deep convolutional neural networks. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VII 14*, pages 467–482. Springer, 2016.
- [36] Yangtao Wang, Yanzhao Xie, Jiangfeng Zeng, Hanpin Wang, Lisheng Fan, and Yufan Song. Cross-modal fusion for multi-label image classification with attention mechanism. *Computers and Electrical Engineering*, 101:108002, 2022.
- [37] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020.
- [38] Junfei Xiao, Yutong Bai, Alan Yuille, and Zongwei Zhou. Delving into masked autoencoders for multi-label thorax disease classification. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3588–3600, 2023.
- [39] Yuzhe Yang and Zhi Xu. Rethinking the value of labels for improving class-imbalanced learning. *Advances in neural information processing systems*, 33:19290–19301, 2020.
- [40] Jin Ye, Junjun He, Xiaojiang Peng, Wenhao Wu, and Yu Qiao. Attention-driven dynamic graph convolutional network for multi-label image recognition. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 649–665. Springer, 2020.
- [41] Yuhang Zang, Chen Huang, and Chen Change Loy. Fasa: Feature augmentation and sampling adaptation for long-tailed instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3457–3466, 2021.
- [42] Youcai Zhang, Yuhao Cheng, Xinyu Huang, Fei Wen, Rui Feng, Yaqian Li, and Yandong Guo. Simple and robust loss design for multi-label learning with missing labels. *arXiv preprint arXiv:2112.07368*, 2021.
- [43] Qian Zhou, Hua Zou, and Zhongyuan Wang. Long-tailed multi-label retinal diseases recognition via relational learning and knowledge distillation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 709–718. Springer, 2022.