

## Bridging the Gap between Model Explanations in Partially Annotated Multi-label Classification

Youngwook Kim<sup>1</sup>  
 Cordelia Schmid<sup>4</sup>

Jae Myung Kim<sup>2</sup>  
 Zeynep Akata<sup>2,5</sup>

Jieun Jeong<sup>1,3</sup>  
 Jungwoo Lee<sup>1,3\*</sup>

<sup>1</sup>Seoul National University <sup>2</sup>University of Tübingen <sup>3</sup>HodooAI Lab

<sup>4</sup>Inria, Ecole normale supérieure, CNRS, PSL Research University <sup>5</sup>MPI for Intelligent Systems

### Abstract

Due to the expensive costs of collecting labels in multi-label classification datasets, partially annotated multi-label classification has become an emerging field in computer vision. One baseline approach to this task is to assume unobserved labels as negative labels, but this assumption induces label noise as a form of false negative. To understand the negative impact caused by false negative labels, we study how these labels affect the model’s explanation. We observe that the explanation of two models, trained with full and partial labels each, highlights similar regions but with different scaling, where the latter tends to have lower attribution scores. Based on these findings, we propose to boost the attribution scores of the model trained with partial labels to make its explanation resemble that of the model trained with full labels. Even with the conceptually simple approach, the multi-label classification performance improves by a large margin in three different datasets on a single positive label setting and one on a large-scale partial label setting. Code is available at <https://github.com/youngwk/BridgeGapExplanationPAMC>.

### 1. Introduction

Multi-label image classification is the task of predicting all labels corresponding to a given image. Since web-crawled images often contain multiple objects/concepts [3, 32, 35, 44], the importance of this task is rising. However, it faces a significant issue of huge annotation costs. We need C binary labels for each training image to provide exhaustive annotation for a model that classifies images into C categories. It acts as a severe obstacle to scaling multi-label classification datasets.

For this reason, partially annotated multi-label classification [2, 11, 13, 17, 21, 24] has recently become an actively

\*Corresponding author.




Figure 1. **CAM Observation.** We compare the class activation map (CAM) output from two multi-label classification models: one trained with full labels ( $CAM_{full}$ ) and the other trained with partial labels and AN assumption ( $CAM_{partial}$ ). We observe that the overall structure of  $CAM_{partial}$  is not much affected by the noisy false negative labels during training. This observation motivates us to make  $CAM_{partial}$  similar to  $CAM_{full}$  by boosting its relatively large attribution scores. Best viewed in color.

studied topic. In this setting, instead of exhaustive annotation, only a few categories are labeled for each training image. We can effectively reduce the burden of annotation by adopting partial annotation strategies.

One baseline approach for solving a partially annotated multi-label classification task is assuming unobserved labels as negative labels (Assume Negative, AN) [4, 6, 36, 40]. It is a reasonable assumption since most labels are negative labels in the multi-label scenario [33]. However, this assumption causes label noise in a form of false negatives since the actually positive but unannotated labels are incor-

rectly assumed to be negative. Since this label noise perturbs the learning process of the model [1, 7, 18, 45], recent studies on a partially annotated multi-label classification focus on suppressing the influence of label noise by ignoring or correcting the loss of samples that are likely to be false negatives [2, 21].

Aside from recent research directions, we delve into “how” false negative labels influence a multi-label classification model. We conduct control experiments with two models. One is the model trained with partial labels and AN assumption where false negative labels exist. The other is the model trained with full annotations and thus trained without false negatives. We compare the class activation map (CAM) [49] output between the two models to see the difference in how each model understands the input image and makes a prediction result.

Figure 1 shows that a model trained with false negatives still highlights similar regions to one trained with full annotation. However, the attribution scores in the highlighted areas are much small. This observation leads us to think that if we scale up the damaged score of the highlighted region in the model trained with false negatives, the explanation of this model will become similar to that of the model trained with full annotation.

To this end, we introduce a simple piece-wise linear function, named BoostLU, that bridges the gap between the explanation of two models trained with false negatives and with full annotation each. Concretely, we use the modified CNN model to get CAM during the forward pass directly [47], and the logit in the modified CNN model is the mean of attribution scores of CAM. The BoostLU function is applied element-wisely to the CAM output of the modified CNN to boost the scores of the highlighted regions, thereby compensating for the decrease of attribution scores in CAM caused by false negatives. It increases the logit value for positive labels and thus makes a better prediction. Furthermore, when we combine BoostLU with the recently proposed methods [21] that explicitly detect and modify false negatives during training, it helps to detect false negatives better, thus leading to better performance. As a result, we achieve state-of-the-art performance on PASCAL VOC [14], MS COCO [28], NUSWIDE [10], and Openimages V3 [23] datasets in a partial label setting.

We summarize the contributions of this paper as follows.

1. We analyze how the false negative labels affect the explanation of the model in a partially annotated multi-label classification scenario.
2. We propose a simple but effective function, named BoostLU, that compensates for the damage of false negatives in a multi-label classification model with little extra computational cost.
3. When applied during inference, BoostLU boosts the

baseline method (AN)’s test performance without additional training.

4. Combined with recent methods of detecting and modifying false negatives during training, BoostLU boosts the state-of-the-art performance on single positive and large-scale partial label settings.

## 2. Related Works

**Partially annotated multi-label classification.** One primary stream to solve the partially annotated multi-label classification problem is to view unobserved labels as *missing labels*. Earlier works tackled this problem by solving matrix completion [5, 15, 43] or employing the Bayesian model [19, 37]. However, these works require loading all data into memory at once, thus making it infeasible to train deep neural networks. Curriculum labeling [13] proposed a bootstrapping strategy using model prediction. IMCL [17], SE [24], and SST [8] exploited label correlation and image similarity to generate regularization losses or pseudo-labels for missing labels. SARB [31] performed a category-wise mixup on feature space between labeled and unlabeled images to propagate information into missing labels. Zhou et al. [50] proposed entropy maximization loss that suppresses gradients from missing labels to promote learning from observed labels.

Since a significant part of labels is negative in a multi-label setting [33], there is another stream to treat unobserved labels as negatives and try to lessen the harmful impact of false negatives. In other words, it views unobserved labels as *noisy labels*. ROLE [11] proposed to estimate unobserved labels while simultaneously regularizing the estimation with an average number of positive labels online. Kim et al. [21] observed the memorization effect [1] in a noisy multi-label classification setting that the model learns from clean labels first. Thus false negative labels are likely to show large loss values during training. Then they suggested three methods, LL-R, LL-Ct, and LL-Cp, that prevent false negatives from being memorized by rejecting, temporally correcting, and permanently correcting samples with large losses, respectively. P-ASL [2] assigned different scaling rates between annotated negatives and assumed negatives. It also ignored losses from categories with high prediction scores or label prior values. In this work, we look at false negatives differently and study their effect on model explanation.

**Class activation mapping.** Class activation mapping (CAM) [49] provides information about where the classification model is attending to generate prediction scores. There are several follow-up works, including Grad-CAM [34], which generates model-agnostic attention maps, and CALM [20], which strengthens the interpretability of attention maps.

Since CAM provides localization ability to classification models, it has been widely used for various vision tasks, such as weakly supervised object localization [9, 12, 41, 42] and weakly supervised semantic segmentation [25, 26, 29, 42, 46]. Recently, Zhang et al. [48] utilized CAM in facial expression recognition in the presence of noisy labels. They found that the model trained with noisy labels highlights only part of the features and suggested a random masking strategy to prevent memorizing partial features. Although there is a similarity in that they inspected the CAM output of the model in noisy label situations, our work is different since we focus on the noisy multi-label classification setting with another type of noise.

### 3. Preliminary

This section introduces the formal definition of a partially annotated multi-label classification in §3.1. Next, we briefly summarize the class activation map (CAM) in §3.2.

#### 3.1. Problem Definition

We aim to train a multi-label classification model with dataset  $\mathcal{D}$  consisting of pairs of input image  $\mathbf{x}$  and partially annotated label  $\mathbf{y}$ . Each category can have three kinds of labels: 0, 1, and  $\phi$ . In other words,  $\mathbf{y} \in \mathcal{Y} = \{0, 1, \phi\}^C$  where  $\phi$  indicates the absence of annotation and  $C$  is the number of total categories. Denote the index set of positive labels, negative labels, and unannotated labels as  $\mathcal{I}^p$ ,  $\mathcal{I}^n$ , and  $\mathcal{I}^\phi$ , respectively. We study the setting where labels are sparsely annotated, i.e.,  $|\mathcal{I}^p| + |\mathcal{I}^n| \ll |\mathcal{I}^\phi|$ .

A straightforward approach to train the model given partial labels is to treat unannotated labels by assuming negative (AN) and use binary cross-entropy as a loss function:

$$\mathcal{L}_{AN} = \frac{1}{C} \left[ \sum_{i \in \mathcal{I}^p} \mathcal{L}_+ + \sum_{i \in \mathcal{I}^n \cup \mathcal{I}^\phi} \mathcal{L}_- \right] \quad (1)$$

where  $\mathcal{L}_+ = -\log(\sigma(g_i))$ ,  $\mathcal{L}_- = -\log(1-\sigma(g_i))$  and  $g_i$  is a logit for  $i$ -th category. However, labels whose true label is positive but unannotated are incorrectly assumed to be negative and become false negatives. Denote the index set of true negative and false negative labels as  $\mathcal{I}^{tn}$  and  $\mathcal{I}^{fn}$ , then  $\mathcal{I}^n \cup \mathcal{I}^\phi = \mathcal{I}^{tn} \cup \mathcal{I}^{fn}$ . We set the approach of training the model with Equation (1) as the baseline method and investigate the influence of false negatives on the multi-label classification model.

#### 3.2. Recap CAM

Most CNN architectures consist of several convolution layers (Convs), followed by a Global Average Pooling (GAP) layer [27] and a fully connected (FC) layer. Let the last convolutional feature map be  $\mathbf{F} \in \mathbb{R}^{H \times W \times D}$ , and a weight matrix of the FC layer be  $\mathbf{W} \in \mathbb{R}^{C \times D}$  where

$(H, W)$  and  $D$  are the spatial size and channel size of the feature map, respectively. We can obtain the class activation map (CAM) [49] for class  $c$  ( $M_c$ ) by

$$M_c = \sum_{d=1}^D \mathbf{W}_{cd} \mathbf{F}_d, \quad (2)$$

where  $\mathbf{F}_d$  denotes  $d$ -th channel of  $\mathbf{F}$ .  $M_c$  explains the model's prediction by attributing scores on each pixel.

Instead of performing post-processing to get CAM as in Equation (2), we can directly get CAM during the forward pass by reordering the last two layers from Convs-GAP-FC to Convs-1x1Conv-GAP where 1x1Conv is the one-by-one convolutional layer with the weight  $\mathbf{W}$  [47]. The output feature maps of 1x1Conv become the same as  $M$ , and the logit  $g_c$  becomes

$$g_c = \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W (M_c)_{ij}. \quad (3)$$

Thus, we can interpret each element  $(M_c)_{ij}$  as an *attribution score* at spatial location  $(i, j)$  contributing to the logit for class  $c$ . For the following sections, we utilize this modified architecture to facilitate the application of our method.

### 4. Impact of False Negatives on CAM

It is well known that neural networks can memorize wrong labels due to their large model capacities [45]. Likewise, if we train a multi-label classification model with AN loss (Equation (1)) when given partial labels, the model is damaged by memorizing false negative labels [21]. It results in poor performance compared to the model trained with full labels, which false negatives have not influenced.

To better understand why the model trained with partial labels performs less than that with full labels, we analyze the behavioral difference between these two models. Concretely, we use a class activation map (CAM) [49] to explain each model's prediction and compare the explanation results. We train two multi-label classification models on a COCO dataset [28] with the same CNN architecture ResNet-50 [16]: one model with full labels using binary cross entropy loss and the other with partial labels using AN loss (Equation (1)). We denote the CAM output from each model as  $CAM_{full}$  and  $CAM_{partial}$ , respectively.

To analyze the explanation of these two models, we first compute the Spearman correlation between  $CAM_{full}$  and  $CAM_{partial}$  on positive labels. We show the distribution of the correlation values on the test set in Figure 2a. For comparison, we consider a 2D Gaussian image centered at the midpoint and calculate the Spearman correlation coefficient between this Gaussian image and  $CAM_{full}$ . We observe that there is mainly a positive correlation between  $CAM_{full}$  and




Figure 2. **CAM Analysis on COCO test set.** (a): Distribution of Spearman correlation coefficients between  $CAM_{full}$  and  $CAM_{partial}$  from the same image. Overall positive correlation implies that  $CAM_{partial}$  has a structure similar to  $CAM_{full}$ . (b), (c): Boxplot of the average of top/bottom 5% of attribution scores, respectively. The damage of false negative labels to the model mainly lowers the upper attribution scores for positive labels while maintaining its overall structure in CAM.

$CAM_{partial}$ , while the correlation of the control group is distributed widely but mostly around zero. It implies that the overall structure (i.e., the attribution ranking among pixels) of  $CAM_{partial}$  is preserved despite the influence of false negative labels, therefore having a high Spearman correlation with  $CAM_{full}$ . We can also visually inspect the similar structure between  $CAM_{partial}$  and  $CAM_{full}$  in Figure 1, where both CAMs highlight similar regions.

Since we know that the overall structure is similar between  $CAM_{full}$  and  $CAM_{partial}$ , we then compare the range of attribution scores between  $CAM_{full}$  and  $CAM_{partial}$ . Concretely, we compute the mean of the highest 5% of attribution scores and the lowest 5%, respectively, for each CAM and summarize the distribution of these values on the test set in Figure 2b and 2c. Note that we take an average of 5% of scores to reduce the effect of outliers. We observe that top-ranking attribution scores of  $CAM_{partial}$  from positive labels drop sharply compared to  $CAM_{full}$ , while these scores from negative labels remain similar. Also, there is little difference in bottom-ranking attribution scores between  $CAM_{full}$  and  $CAM_{partial}$ , both on positive and negative labels. It implies that false negatives mainly affect the model’s understanding in regions with relatively high attribution scores, especially for positive labels. Consequently, the decrease of attribution scores at specific regions in CAM leads to a decrease in the logit value (since logit is the average of attribution scores on CAM as in Equation (3)), making the model predicts a lower score for the positive category. The change of gradient during training can explain the occurrence of this phenomenon.

**Gradient analysis.** In Equation (1), recall that the BCE loss is  $\mathcal{L}_+$  with a positive target and  $\mathcal{L}_-$  with a negative one. Their gradients with respect to the logit  $g$  are

$$\frac{\partial \mathcal{L}_+}{\partial g} = \sigma(g) - 1, \quad \frac{\partial \mathcal{L}_-}{\partial g} = \sigma(g). \quad (4)$$

For a training image  $\mathbf{x}$ , the gradient difference on the logit

$g$  between partial label (with AN assumption) and full label case is given by

$$\begin{aligned} & \frac{1}{C} \left[ \sum_{i \in \mathcal{I}^p} \frac{\partial \mathcal{L}_+}{\partial g_i} + \sum_{i \in \mathcal{I}^{fn}} \frac{\partial \mathcal{L}_-}{\partial g_i} + \sum_{i \in \mathcal{I}^{tn}} \frac{\partial \mathcal{L}_-}{\partial g_i} \right] \\ & - \frac{1}{C} \left[ \sum_{i \in \mathcal{I}^p} \frac{\partial \mathcal{L}_+}{\partial g_i} + \sum_{i \in \mathcal{I}^{fn}} \frac{\partial \mathcal{L}_+}{\partial g_i} + \sum_{i \in \mathcal{I}^{tn}} \frac{\partial \mathcal{L}_-}{\partial g_i} \right] \quad (5) \\ & = \frac{1}{C} \left[ \sum_{i \in \mathcal{I}^{fn}} \left( \frac{\partial \mathcal{L}_-}{\partial g_i} - \frac{\partial \mathcal{L}_+}{\partial g_i} \right) \right] = \frac{|\mathcal{I}^{fn}|}{C}. \end{aligned}$$

Equation (5) shows that the logit receives more gradients proportional to the number of false negative labels on a partial label setting. Therefore, as training progresses, the additional gradients from false negatives are gradually accumulated in the logit, making the logit smaller than the model trained on full labels. Since the logit is equal to the average of CAM, the attribution scores of  $CAM_{partial}$  become lower than that of  $CAM_{full}$ .

## 5. Proposed Method

In this section, we propose a conceptually simple but effective method to make the model trained with partial labels resemble the model trained with full labels by mimicking the explanation. We propose a function BoostLU devised to compensate for the damaged attribution score of the explanation due to false negatives in §5.1. We then introduce three scenarios that utilize our function through §5.2 ~ §5.4.

### 5.1. BoostLU

From the modified CNN architecture described in §3.2, define convolutional layers (Convs-1x1Conv) as  $\Phi$ . Given an input image  $\mathbf{x}$ , its class activation map (CAM) is  $\mathbf{M} =$




Figure 3. **Schematic diagram of applying BoostLU.** BoostLU is applied to the model’s CAM output element-wisely to compensate for the attribution scores damaged by false negative labels.

$\Phi(x)$ . Our goal is to make the explanation  $M$  of the model trained with partial labels closer to the explanation of the model trained with full labels, even if we do not have access to the full labels, thus improving the prediction performance.

In the previous section, we observe that when a multi-label classification model is trained with AN loss, the way the model understands images is damaged by false negatives. However, we also find that this damage is mainly focused on a drop in the relatively high attribution scores while the overall spatial structure of CAM is preserved. Based on these findings, we conjecture that if the damaged high attribution scores are scaled up in the model trained with partial labels,  $CAM_{partial}$  will become similar to  $CAM_{full}$ . To achieve this, we devise a piece-wise linear function that boosts the attribution scores that are above a certain threshold:

$$f(x) = \begin{cases} \alpha x + (1 - \alpha)\beta, & x \geq \beta \\ x, & x < \beta \end{cases}, \quad (6)$$

where  $\alpha$  is a scaling factor with  $\alpha > 1$ , and  $\beta$  is a threshold determining whether to boost the score. Since top-ranking attribution scores on CAM tend to have large positive values for positive labels and around zero for negative labels (as seen in Figure 2b), we search for the values of  $\beta$  around zero. Since we empirically observe no significant difference in model performance for different  $\beta$  (these results are reported in the Appendix), we only consider the simplest case of  $\beta = 0$ . Then we can rewrite Equation (6) in a ReLU-like form as

$$\text{BoostLU}(x) = \max(x, \alpha x). \quad (7)$$

By applying BoostLU to each element of CAM, as illustrated in Figure 3, BoostLU boosts positive attribution scores by  $\alpha$  times, which are the main target to be damaged by false negatives, while maintaining the negative scores unchanged. These selectively boosted attribution scores are aggregated through the GAP layer to produce a logit value as

$$g(x) = (\text{GAP} \circ \text{BoostLU} \circ \Phi)(x). \quad (8)$$

From now on, we will consider three different scenarios for applying BoostLU in multi-label classification.

## 5.2. Usage 1: BoostLU in inference only

Since the idea of BoostLU comes from analyzing the CAM of a model which finished training with AN loss, we first propose to apply BoostLU only during the inference phase of that model. Initially, this model produces low logits for categories whose label is positive. However, applying BoostLU increases the corrupted attribution scores and produces higher logits. At the same time, boosting effect is not much for categories whose label is negative; therefore, its logits remain almost the same. As a result, prediction scores are better separated between samples with positive and negative labels, improving average precision.

## 5.3. Usage 2: BoostLU in both training and inference

Next, we consider applying BoostLU during the training phase with AN loss and the inference phase. The gradient of logit  $g$  with respect to the attribution score on CAM at location  $(i, j)$  (i.e.,  $M_{ij}$ ) then becomes

$$\frac{\partial g}{\partial M_{ij}} = \begin{cases} \alpha/HW, & M_{ij} \geq 0 \\ 1/HW, & M_{ij} < 0 \end{cases}. \quad (9)$$

Compared to the case that does not use BoostLU, where every spatial location gets a uniform gradient of  $1/HW$ , the locations with positive attribution scores receive gradients boosted by  $\alpha$  times. Thanks to the boosted gradients, these locations are encouraged to produce higher attribution scores during training when the model receives a positive label. Also, when a true negative label comes in, these locations are encouraged to produce lower attribution scores.

However, in practice, we observe only marginal improvement in model performance. It is because the boosted gradients have an adverse effect when false negatives come in as input. That is, BoostLU also boosts the wrong direction of gradients from false negatives, which can be easily seen by combining Equation (4) and (9):

$$\frac{\partial \mathcal{L}_-}{\partial M_{ij}} = \frac{\partial \mathcal{L}_-}{\partial g} \cdot \frac{\partial g}{\partial M_{ij}} \quad (10)$$

Note that  $\partial \mathcal{L}_- / \partial g$  has a wrong sign for false negatives, and it decreases CAM values.

## 5.4. Usage 3: Combination with Large Loss Modification

To alleviate the problem mentioned above, we propose combining our BoostLU with recent studies [2, 21] that detect and treat suspicious false negatives while training multi-label classification models. We especially adopt three methods, i.e., LL-R, LL-Ct, and LL-Cp [21], since they work on several partial label settings. When these methods are combined with BoostLU, they suppress the side effects caused by false negatives. As a result, the model can

take full advantage of the boosted gradients from the positive labels during training. Moreover, because these combined methods consider samples with relatively high prediction scores among unobserved labels as false negatives, BoostLU helps the model detect more false negatives by boosting their logit values.

## 6. Experiments

To validate the efficacy of our proposed method, we report our experimental results on two partial label settings: single positive label (§6.1) and large-scale partial label (§6.2). In both sections, we adopt mean Average Precision (mAP) as an evaluation metric and report the performance on a test set using the model weight with the highest mAP in the validation set. We fix our hyperparameters as  $\alpha = 5$ ,  $\beta = 0$ . Next, we present analysis results on §6.3.

### 6.1. Single positive label

**Datasets.** We target four multi-label classification datasets: PASCAL VOC 2012 [14], MS COCO 2014 [28], NUSWIDE [10], and CUB [38]. Each dataset is annotated for 20 classes, 80 classes, 81 concepts, and 312 attributes. Since they are fully annotated, we only keep one positive label and drop the rest of the labels for every training image to build a single positive label setting identical to [11].

**Hyperparameter settings.** For a fair comparison, we set the same search space as [11, 50]:  $\{8, 16\}$  for batch size and  $\{10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\}$  for learning rate. We train the model for 10 epochs with Adam optimizer [22]. LL-R, LL-Ct, and LL-Cp [21] have a hyperparameter  $\Delta_{rel}$  that controls the slope of increase in the modification rate. We set  $\Delta_{rel} = 0.5$  for LL-R, 0.2 for LL-Ct, and 0.1 for LL-Cp, respectively. We set a 10x learning rate for the last 1x1Conv layer.

**Implementation details.** We also follow identical configurations as [11, 21, 50]. Specifically, 20% of the original training set is used for validation. ResNet-50 [16] CNN backbone pre-trained on ImageNet [44] is used as a feature extractor. Each image is resized to 448x448 before being fed to CNN, and only random horizontal flipping is used for data augmentation during training. Note that some categories do not have positive labels in the CUB dataset on a generated single positive label setting. In these categories, we do not apply BoostLU when training as the benefit from the boosted gradient becomes weakened.

**Results of ablation study.** We first conduct ablation studies on PASCAL VOC and COCO datasets. Its results are reported in Table 1. First, we show the performance of the model trained with AN loss in the first row. In the second row, it can be seen that when BoostLU is applied during inference of this model, its test performance is improved even without additional training. It confirms the property of BoostLU that compensates for the damaged attribution

BoostLU in inference	BoostLU in training	LL-R in training	Performance	
			VOC	COCO
			86.10	64.58
✓			87.31	66.27
✓	✓		86.73	65.33
		✓	88.24	70.60
	✓	✓	87.18	68.45
✓		✓	88.90	70.87
✓	✓	✓	<b>89.27</b>	<b>72.82</b>

Table 1. **Ablation study on BoostLU and LL-R.** We test seven combinations of using BoostLU and LL-R [21] on VOC and COCO datasets. Training a model with both LL-R and BoostLU and applying BoostLU during inference shows the best mAP.

Methods	VOC	COCO	NUS	CUB
Full Label	89.42	76.78	52.08	30.90
AN	85.89	64.92	42.27	18.31
LS [30]	87.90	67.15	43.77	16.26
ASL [33]	87.76	68.78	46.93	18.81
ROLE [11]	87.77	67.04	41.63	13.66
ROLE + LI [11]	88.26	69.12	45.98	14.86
EM [50]	89.09	70.70	47.15	20.85
EM + APL [50]	89.19	70.87	47.59	<b>21.84</b>
LL-R [21]	88.27	70.70	48.76	19.56
+ BoostLU (Ours)	<b>89.29</b>	<b>72.89</b>	<b>49.59</b>	19.80
LL-Ct [21]	87.79	70.29	48.08	19.06
+ BoostLU (Ours)	88.61	71.78	48.37	19.25
LL-Cp [21]	87.44	70.27	47.92	19.21
+ BoostLU (Ours)	87.81	71.41	48.61	19.34

Table 2. **Experimental results on various datasets with single positive label setting.** Each number shows the average of mAP in three experiments. A bold number means the best performance. Results of methods except for LL-R, LL-Ct, and LL-Cp are taken from [50]. We report the reimplemented results for LL-R, LL-Ct, and LL-Cp with the same hyperparameter search space as [11, 50].

score. However, if we further apply BoostLU while training (third row), the performance improvement is lower than when BoostLU is applied only during inference. We can observe the side effect of BoostLU that the gradient received by the region with a positive attribution score is boosted even for false negative labels.

In the fourth row, we show the performance of LL-R, which rejects large losses during training. We then train the model by applying both LL-R and BoostLU during training and BoostLU during inference. Its performance is reported in the final row, and its improvement is much more significant than the case where LL-R is not applied (+0.63

Methods	Group 1	Group 2	Group 3	Group 4	Group 5	All Classes
CNN-RNN [39]	68.76	69.70	74.18	78.52	84.61	75.16
Curriculum Labeling [13]	70.37	71.32	76.23	80.54	86.81	77.05
IMCL [17]	70.95	72.59	77.64	81.83	87.34	78.07
P-ASL [2]	73.19	78.61	85.11	87.70	90.61	83.03
LL-R [21]	77.76	79.07	81.94	84.51	89.36	82.53
+ BoostLU (Ours)	79.28	80.81	83.32	85.63	90.27	83.86
LL-Ct [21]	77.76	79.18	81.97	84.46	89.51	82.58
+ BoostLU (Ours)	79.43	80.75	83.41	85.70	90.41	83.94
LL-Cp [21]	77.49	79.22	81.89	84.51	89.18	82.46
+ BoostLU (Ours)	79.53	81.04	83.40	85.85	90.39	<b>84.04</b>

Table 3. **Experimental results on a OpenImages V3 dataset.** Each group includes 1,000 classes without overlapping. Group 1 has the smallest annotations, and Group 5 has the most. The number of annotations increases as the group number increases. LL-R, LL-Ct, and LL-Cp are reimplemented and the other results are borrowed from [2]. A bold number shows the best performance.

v.s. +1.03 on PASCAL, and +0.75 v.s. +2.22 on COCO). Thanks to LL-R filtering out false negatives, the side effect of the boosted gradient becomes minimized. Moreover, since our BoostLU helps LL-R detect false negatives, its advantage is further amplified. We also find in the last three rows that when we combine BoostLU with LL-R, applying BoostLU either during training or during inference results in a performance drop compared to applying it during both phases. This shows that BoostLU plays a vital role both in training and inference, together with large loss modification methods. In particular, it is crucial to apply BoostLU during inference to achieve high performance. Additional discussion about this is described in the Appendix.

From now on, we will only report the experimental results using the configuration of the last row (BoostLU in inference + BoostLU in training + LL-R in training).

**Comparison with prior arts.** We compare our results with recent state-of-the-art: Label Smoothing (LS) [30], Asymmetric loss (ASL) [33], ROLE (with LinearInit) [11], and Entropy-maximization loss (EM) with Asymmetric Pseudo-Labeling (APL) [50]. We train the network three times and report the average test performance.

The results are shown in Table 2. We find that applying BoostLU in both training and inference consistently improves the performance of LL-R, LL-Ct, and LL-Cp in all datasets, only with little additional computational cost. It achieves +1.02, +0.82, and +0.37 mAP improvement in VOC, as well as +2.19, +1.49, and +1.14 mAP improvement in COCO, respectively. Especially the performance of LL-R + BoostLU shows the most significant increase, achieving state-of-the-art performance and reaching closest to the full label performance on VOC, COCO, and NUSWIDE. It also surpasses the previous state-of-the-art method EM+APL which does not use AN assumption on these datasets. However, the performance improvement is not that large in CUB.

Since CUB has an annotation for attributes, the number of false negative labels is much higher, and this may increase the side effect of BoostLU when applied during training.

## 6.2. Large-scale partial label

**Dataset.** We target a partially annotated OpenImages V3 [23] dataset which consists of 3.4M training images, 41,620 validation images, and 125,436 test images with 5,000 trainable classes (having more than 30 human-verified samples in the training set and 5 in the valid or test sets). We sort these classes in ascending order by the number of annotations in the training set and divide them into five groups of equal size 1,000. We report the mAP score averaged within each group and the entire 5,000 classes.

**Implementation details.** We use ImageNet [44] pre-trained ResNet-101 [16] as a feature extractor, the same as prior works. We follow [21] to set the learning rate as  $2 \times 10^{-5}$  and batch size as 288. We train the model for 20 epochs and set  $\Delta_{rel} = 0.005$ . We resize every image to 224x224 resolution and perform a random horizontal flip during training. We set a 10x learning rate for the last 1x1Conv layer.

**Results.** We compare our results with prior works: CNN-RNN [39], Curriculum Labeling [13], IMCL [17], and P-ASL [2]. As shown in Table 3, BoostLU also works well in a real partial label scenario. Combined with LL-R, LL-Ct, and LL-Cp, it boosts their performance by a large margin: improvement of +1.33, +1.36, and +1.58 mAP, respectively. All of the combined methods surpass other previous methods and achieve state-of-the-art performance. In particular, LL-Cp + BoostLU shows the highest 84.04 mAP.

## 6.3. Analysis

**Qualitative results.** Figure 4 visualizes the CAM results from four different methods. The category corresponding to the CAM is shown above the image. The prediction score,

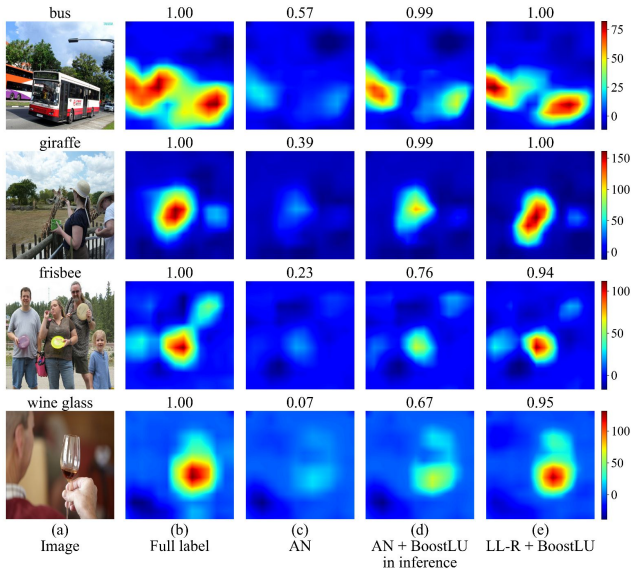


Figure 4. **Qualitative results.** Categories and their corresponding prediction scores are displayed above the images and CAM results. LL-R + BoostLU is the closest to the explanation and prediction score of the model trained with full labels.

obtained by averaging attribution scores on CAM and applying sigmoid activation, is shown above each CAM. First, column (c) shows that a model trained with AN loss gives low prediction scores due to the damage of false negatives. Although this model highlights similar regions for a given input image, the attribution scores of the corresponding regions are considerably shrunk compared to column (b).

When we perform inference by attaching BoostLU to this model, it can be seen in column (d) that BoostLU successfully recovers the model’s explanation, yielding high prediction scores. For LL-R + BoostLU in column (e), its model explanation is further improved due to the role of LL-R and BoostLU during training which further accelerates the improvement of the attribution score of the highlighted region. It is the most similar to the explanation of the model trained with full annotation (column (b)) compared to other methods.

**Synergy effect of BoostLU and large loss modification methods during training.** We train LL-R and LL-R + BoostLU on the COCO dataset with the same  $\Delta_{rel} = 0.5$  to make both models reject the same number of samples during training. We then inspect how many of the rejected samples are false negative labels. Figure 5 shows the number of false negative labels rejected by each model per epoch. It can be seen that after the warmup phase (first epoch), LL-R + BoostLU rejects more false negatives than LL-R in every epoch. It is because BoostLU boosts the logit value of false negative samples, thus boosting the large loss modification methods’ ability to detect false negatives. At the same time, it also reduces the number of true negative sam-

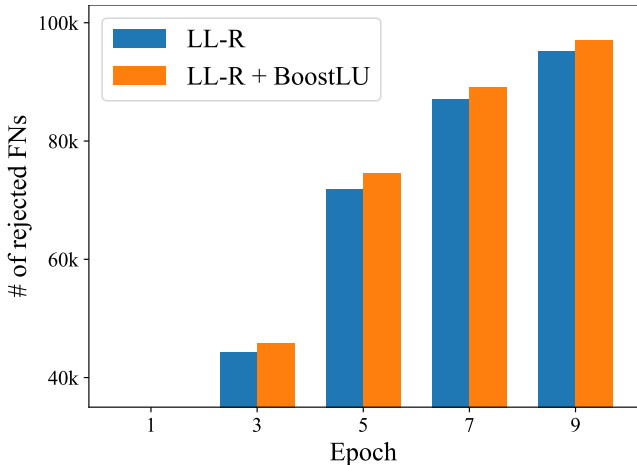


Figure 5. **Comparison of the number of rejected false negative labels.** BoostLU helps LL-R detect more false negative labels in every epoch.

ples that the model incorrectly rejects, further contributing to performance improvement.

## 7. Conclusion

In this paper, we studied the effect of false negative labels on model explanation when assuming unobserved labels as negatives in a partially annotated multi-label classification situation. We found that the overall spatial shape of the explanation tends to be preserved, but the scale of attribution scores is significantly affected. Based on these findings, we proposed a conceptually simple piece-wise linear function BoostLU that compensates for the damaged attribution scores. Through several experiments, we confirmed that BoostLU successfully contributed to bridging the explanation of the model closer to the explanation of the model trained with full labels. Furthermore, combined with large loss modification methods, it achieved state-of-the-art performance on several multi-label datasets.

**Acknowledgements.** Youngwook Kim thanks Junghyun Lee and Youngmin Ro for their valuable help. Jae Myung Kim thanks the European Laboratory for Learning and Intelligent Systems (ELLIS) PhD program and the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support. This work is in part supported by National Research Foundation of Korea (NRF, 2021R1A4A1030898(10%)), Institute of Information & communications Technology Planning & Evaluation (IITP, 2021-0-00106 (50%), 2021-0-01059 (20%), 2021-0-00180 (20%)) grant funded by the Ministry of Science and ICT (MSIT), INMAC, and BK21-plus. This work is also supported by DFG project number 276693517, by BMBF FKZ: 01IS18039A, by the ERC (853489 - DEXIM), and by EXC number 2064/1 – project number 390727645.



## References

- [1] Devansh Arpit, Stanisław Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *ICML*, pages 233–242. PMLR, 2017. **2**
- [2] Emanuel Ben-Baruch, Tal Ridnik, Itamar Friedman, Avi Ben-Cohen, Nadav Zamir, Asaf Noy, and Lihi Zelnik-Manor. Multi-label classification with partial annotations using class-aware selective loss. In *CVPR*, pages 4764–4772, 2022. **1, 2, 5, 7**
- [3] Lucas Beyer, Olivier J Hénaff, Alexander Kolesnikov, Xiaohua Zhai, and Aäron van den Oord. Are we done with imagenet? *arXiv preprint arXiv:2006.07159*, 2020. **1**
- [4] Serhat Selcuk Bucak, Rong Jin, and Anil K Jain. Multi-label learning with incomplete class assignments. In *CVPR*, pages 2801–2808. IEEE, 2011. **1**
- [5] Ricardo Cabral, Fernando Torre, Joao P Costeira, and Alexandre Bernardino. Matrix completion for multi-label image classification. *NeurIPS*, 24, 2011. **2**
- [6] Minmin Chen, Alice Zheng, and Kilian Weinberger. Fast image tagging. In *ICML*, pages 1274–1282. PMLR, 2013. **1**
- [7] Pengfei Chen, Junjie Ye, Guangyong Chen, Jingwei Zhao, and Pheng-Ann Heng. Beyond class-conditional assumption: A primary attempt to combat instance-dependent label noise. In *AAAI*, volume 35, pages 11442–11450, 2021. **2**
- [8] Tianshui Chen, Tao Pu, Hefeng Wu, Yuan Xie, and Liang Lin. Structured semantic transfer for multi-label recognition with partial labels. In *AAAI*, volume 36, pages 339–346, 2022. **2**
- [9] Junsuk Choe, Seong Joon Oh, Seungho Lee, Sanghyuk Chun, Zeynep Akata, and Hyunjung Shim. Evaluating weakly supervised object localization methods right. In *CVPR*, pages 3133–3142, 2020. **3**
- [10] Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yantao Zheng. Nus-wide: a real-world web image database from national university of singapore. In *Proceedings of the ACM international conference on image and video retrieval*, pages 1–9, 2009. **2, 6**
- [11] Elijah Cole, Oisín Mac Aodha, Titouan Llorca, Pietro Perona, Dan Morris, and Nebojsa Jojic. Multi-label learning from single positive labels. In *CVPR*, pages 933–942, 2021. **1, 2, 6, 7**
- [12] Elijah Cole, Kimberly Wilber, Grant Van Horn, Xuan Yang, Marco Fornoni, Pietro Perona, Serge Belongie, Andrew Howard, and Oisín Mac Aodha. On label granularity and object localization. *ECCV*, 2022. **3**
- [13] Thibaut Durand, Nazanin Mehrasa, and Greg Mori. Learning a deep convnet for multi-label classification with partial labels. In *CVPR*, pages 647–657, 2019. **1, 2, 7**
- [14] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, June 2010. **2, 6**
- [15] Andrew Goldberg, Ben Recht, Junming Xu, Robert Nowak, and Jerry Zhu. Transduction with matrix completion: Three birds with one stone. *NeurIPS*, 23, 2010. **2**
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. **3, 6, 7**
- [17] Dat Huynh and Ehsan Elhamifar. Interactive multi-label cnn learning with partial labels. In *CVPR*, pages 9423–9432, 2020. **1, 2, 7**
- [18] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, pages 2304–2313. PMLR, 2018. **2**
- [19] Ashish Kapoor, Raajay Viswanathan, and Prateek Jain. Multilabel classification using bayesian compressed sensing. *NeurIPS*, 25:2645–2653, 2012. **2**
- [20] Jae Myung Kim, Junsuk Choe, Zeynep Akata, and Seong Joon Oh. Keep calm and improve visual feature attribution. In *ICCV*, pages 8350–8360, 2021. **2**
- [21] Youngwook Kim, Jae Myung Kim, Zeynep Akata, and Jungwoo Lee. Large loss matters in weakly supervised multi-label classification. In *CVPR*, pages 14156–14165, 2022. **1, 2, 3, 5, 6, 7**
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. **6**
- [23] Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Andreas Veit, Serge Belongie, Victor Gomes, Abhinav Gupta, Chen Sun, Gal Chechik, David Cai, Zheyun Feng, Dhyanesh Narayanan, and Kevin Murphy. Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://github.com/openimages>*, 2017. **2, 7**
- [24] Kaustav Kundu and Joseph Tighe. Exploiting weakly supervised visual patterns to learn from partial annotations. *NeurIPS*, 33:561–572, 2020. **1, 2**
- [25] Jungbeom Lee, Eunji Kim, and Sungroh Yoon. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *CVPR*, pages 4071–4080, 2021. **3**
- [26] Seungho Lee, Minhyun Lee, Jongwuk Lee, and Hyunjung Shim. Railroad is not a train: Saliency as pseudo-pixel supervision for weakly supervised semantic segmentation. In *CVPR*, pages 5495–5505, 2021. **3**
- [27] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013. **3**
- [28] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755. Springer, 2014. **2, 3, 6**
- [29] Sheng Liu, Kangning Liu, Weicheng Zhu, Yiqiu Shen, and Carlos Fernandez-Granda. Adaptive early-learning correction for segmentation from noisy annotations. In *CVPR*, pages 2606–2616, 2022. **3**
- [30] Michal Lukasik, Srinadh Bhojanapalli, Aditya Menon, and Sanjiv Kumar. Does label smoothing mitigate label noise? In *ICML*, pages 6448–6458. PMLR, 2020. **6, 7**
- [31] Tao Pu, Tianshui Chen, Hefeng Wu, and Liang Lin. Semantic-aware representation blending for multi-label image recognition with partial labels. *AAAI*, 2022. **2**

- [32] Sai Rajeswar, Pau Rodriguez, Soumye Singhal, David Vazquez, and Aaron Courville. Multi-label iterated learning for image classification with label ambiguity. In *CVPR*, pages 4783–4793, 2022. 1
- [33] Tal Ridnik, Emanuel Ben-Baruch, Nadav Zamir, Asaf Noy, Itamar Friedman, Matan Protter, and Lihi Zelnik-Manor. Asymmetric loss for multi-label classification. In *ICCV*, pages 82–91, 2021. 1, 2, 6, 7
- [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 *ICCV*, pages 618–626, 2017. 2
- [35] Vaishal Shankar, Rebecca Roelofs, Horia Mania, Alex Fang, Benjamin Recht, and Ludwig Schmidt. Evaluating machine accuracy on imagenet. In *ICML*, pages 8634–8644. PMLR, 2020. 1
- [36] Yu-Yin Sun, Yin Zhang, and Zhi-Hua Zhou. Multi-label learning with weak label. In *AAAI*, 2010. 1
- [37] Deepak Vasisht, Andreas Damianou, Manik Varma, and Ashish Kapoor. Active learning for sparse bayesian multi-label classification. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 472–481, 2014. 2
- [38] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 6
- [39] Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, and Wei Xu. Cnn-rnn: A unified framework for multi-label image classification. In *CVPR*, pages 2285–2294, 2016. 7
- [40] Qifan Wang, Bin Shen, Shumiao Wang, Liang Li, and Luo Si. Binary codes embedding for fast image tagging with incomplete labels. In *ECCV*, pages 425–439. Springer, 2014. 1
- [41] Pingyu Wu, Wei Zhai, and Yang Cao. Background activation suppression for weakly supervised object localization. In *CVPR*, pages 14228–14237. IEEE, 2022. 3
- [42] Jinheng Xie, Jianfeng Xiang, Junliang Chen, Xianxu Hou, Xiaodong Zhao, and Linlin Shen. C2am: Contrastive learning of class-agnostic activation map for weakly supervised object localization and semantic segmentation. In *CVPR*, pages 989–998, 2022. 3
- [43] Miao Xu, Rong Jin, and Zhi-Hua Zhou. Speedup matrix completion with side information: Application to multi-label learning. In *NeurIPS*, pages 2301–2309, 2013. 2
- [44] Sangdoon Yun, Seong Joon Oh, Byeongho Heo, Dongyoon Han, Junsuk Choe, and Sanghyuk Chun. Re-labeling imagenet: from single to multi-labels, from global to localized labels. In *CVPR*, pages 2340–2350, 2021. 1, 6, 7
- [45] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021. 2, 3
- [46] Dong Zhang, Hanwang Zhang, Jinhui Tang, Xian-Sheng Hua, and Qianru Sun. Causal intervention for weakly-supervised semantic segmentation. *NeurIPS*, 33:655–666, 2020. 3
- [47] Xiaolin Zhang, Yunchao Wei, Jiashi Feng, Yi Yang, and Thomas S Huang. Adversarial complementary learning for weakly supervised object localization. In *CVPR*, pages 1325–1334, 2018. 2, 3
- [48] Yuhang Zhang, Chengrui Wang, Xu Ling, and Weihong Deng. Learn from all: Erasing attention consistency for noisy label facial expression recognition. *ECCV*, 2022. 3
- [49] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, pages 2921–2929, 2016. 2, 3
- [50] Donghao Zhou, Pengfei Chen, Qiong Wang, Guangyong Chen, and Pheng-Ann Heng. Acknowledging the unknown for multi-label learning with single positive labels. *ECCV*, 2022. 2, 6, 7