

EDGE: Unknown-aware Multi-label Learning by Energy Distribution Gap Expansion

Yuchen Sun^{1,2}, Qianqian Xu^{1*}, Zitai Wang¹, Zhiyong Yang², Junwei He^{1,2}

¹Key Laboratory of Intelligent Information Processing, Institute of Computing Technology, CAS, Beijing, China

²School of Computer Science and Technology, University of Chinese Academy of Sciences, Beijing, China
{sunyuchen22s, xuqianqian, wangzitai}@ict.ac.cn, yangzhiyong21@ucas.ac.cn, hejunwei22s@ict.ac.cn

Abstract

Multi-label Out-Of-Distribution (OOD) detection aims to discriminate the OOD samples from the multi-label In-Distribution (ID) ones. Compared with its multiclass counterpart, it is crucial to model the joint information among classes. To this end, JointEnergy, which is a representative multi-label OOD inference criterion, summarizes the logits of all the classes. However, we find that JointEnergy can produce an imbalance problem in OOD detection, especially when the model lacks enough discrimination ability. Specifically, we find that the samples only related to minority classes tend to be classified as OOD samples due to the ambiguous energy decision boundary. Besides, imbalanced multi-label learning methods, originally designed for ID ones, would not be suitable for OOD detection scenarios, even producing a serious negative transfer effect. In this paper, we resort to auxiliary outlier exposure (OE) and propose an unknown-aware multi-label learning framework to reshape the uncertainty energy space layout. In this framework, the energy score is separately optimized for tail ID samples and unknown samples, and the energy distribution gap between them is expanded, such that the tail ID samples can have a significantly larger energy score than the OOD ones. What's more, a simple yet effective measure is designed to select more informative OE datasets. Finally, comprehensive experimental results on multiple multi-label and OOD datasets reveal the effectiveness of the proposed method.

Introduction

Traditional classification methods mainly focus on In-Distribution (ID) scenarios, where the training set and the test set share the same data distribution. However, in practical applications, it is unavoidable that some test samples are Out Of the Distribution (OOD) of the training ones (Yang et al. 2021). In this case, an ID model will classify the OOD samples into one of the ID classes due to the notorious over-confidence issue. Such wrong predictions seriously degenerate the model performance. Given this, OOD detection has attracted rising attention in recent years, intending to discriminate the OOD samples from the ID ones (Huang and Li 2021; Wang et al. 2022c; Djuricic et al. 2023)

*Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

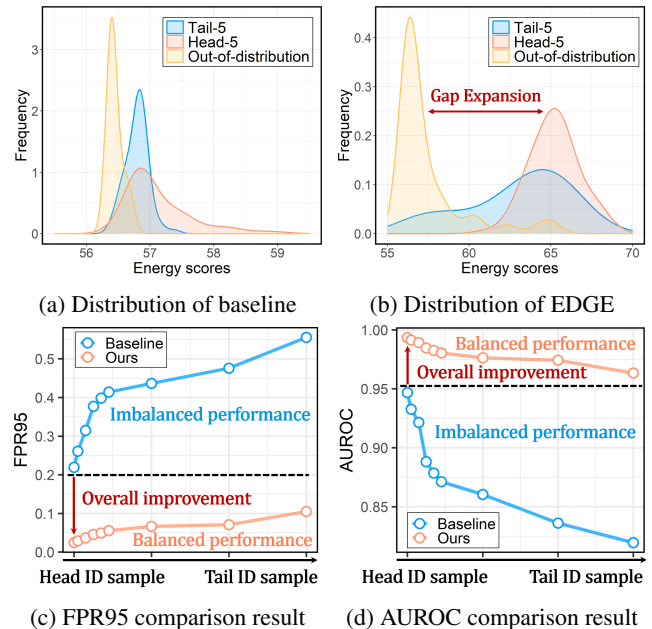


Figure 1: Distribution and performance results with Binary Cross-Entropy loss (BCE) training loss and our framework on Places50 dataset. (a) shows the small gap between known and unknown domains under basic multi-label models, and (b) indicates the improvement of our method. The last two figures show the (c) FPR95 and (d) AUROC imbalanced and balanced results on different in-distribution subsets, where the proportion of in-distribution samples that only relate to minority classes increases with the x-axis value increases.

In this direction, most prior art focuses on scenarios where the ID samples only have a single ground-truth label. Nevertheless, samples are inherently associated with multiple labels in many real-world applications such as autonomous driving (Mao et al. 2023) and medical image (Zhou et al. 2023). For multi-label OOD detection, one intuitive strategy is to adapt the off-the-shelf multiclass algorithms (Du et al. 2022; Wang et al. 2022b). However, Wang et al. (Wang et al. 2021), pioneered in multi-label OOD detection, point out that simply using the largest model output (*i.e.*, MaxLogit (Hendrycks et al. 2022), Max) might be sub-optimal. Considering this, it is crucial to jointly model the information

across different labels for multi-label tasks. To fix this issue, they suggest summing the energy score of all labels as the inference criteria (dubbed JointEnergy), where the energy score is proportional to the logit.

Although JointEnergy is effective for OOD inference, we are surprised to find that it induces a new imbalanced problem. As shown in Figure 1 (a), the energy gap between ID samples and OOD samples is relatively small, especially for the tail ID samples. Notably, in multi-label learning, *tail* samples refer to instances whose relevant classes are all minority ones. In contrast, head samples are relevant to at least one majority class, where minority/majority indicates the class is associated with many/few samples, respectively. The intuition behind this problem is that traditional multi-label optimization methods (like binary cross-entropy loss) treat each class of samples equally while somewhat ignoring the focus on minority classes. Consequently, head samples tend to have large logits and thus are more likely to be recognized as ID samples. By contrast, some hard samples generally have low logits and are misclassified as OOD ones. Due to such a distribution overlap, the OOD performance degenerates and becomes significantly imbalanced across head samples and tail samples. As the blue line shown in Figure 1 (c) & (d), we rank the classes according to the number of associated samples from largest to smallest. When the x-axis values increase, we remove all samples that are associated with the corresponding head categories and observe an obvious degradation in the FPR95 and AUROC for the remaining samples.

To tackle this challenge, we propose a simple yet effective unknown-aware multi-label learning framework named **Energy Distribution Gap Expansion (EDGE)**. Specifically, this method could be divided into three parts. Firstly, we enhance the generalization of models by learning known distribution. Then, we adopt the Outlier Exposure (OE) strategy, *i.e.*, introducing auxiliary OOD data, to lower the logit of OOD samples and separate them from known classes. However, merely optimizing distribution space from the perspective of logit might suffer from a trivial solution. Thus thirdly, we further expand the positive energy gap between tail ID and OOD samples to establish a reliable energy boundary (see Figure 1 (b)). Compared with post-hoc methods, this framework could achieve a balanced performance, as the orange line shown in Figure 1 (c) & (d). To sum up, the whole learning framework is shown in Figure 2, and the main contributions of this paper are three-fold:

- To the best of our knowledge, we are the pioneer in discovering the induced imbalance issue existing in the energy-based method for multi-label OOD detection.
- A novel unknown-aware multi-label learning framework **EDGE** is proposed to reshape the uncertainty energy space of ID and OOD data, as well as a feature-based pre-processing module to guide informative outlier exposure data selection.
- We conduct comprehensive experiments on three large-scale real-world datasets and various OOD datasets, where the test results consistently show the superiority of the EDGE framework.

Related Work

Multi-label Classification

Multi-label classification (MLC) focuses on the instances simultaneously assigned with multiple labels (Wang et al. 2024; Liu et al. 2022; Huang et al. 2020). This topic mainly falls into two categories: **problem transformation** and **algorithm adaptation**. For the former, most studies tend to simplify MLC into multi-class problems (Jernite, Choromanska, and Sontag 2017; Wallach et al. 2019). In contrast, algorithm adaptation leverages traditional machine-learning methods to model data (Gong et al. 2014; Wang et al. 2016; Chen et al. 2019). Unlike traditional MLC work, this paper is dedicated to exploring OOD detection in multi-label learning, which is still an underexplored problem.

Out-of-distribution Detection

The study toward OOD detection is first suggested in (Nguyen, Yosinski, and Clune 2015), and now mainly has two directions, *i.e.*, **inference detection**, and **outlier exposure**. Representatively, **inference-based approaches** leverage informative outputs, like features and logits, to detect OOD inputs. For example, Energy Score (Liu et al. 2020) function identifies OOD samples from the view of energy; MSP (Hendrycks and Gimpel 2017) detects OOD samples by maximum predicted softmax probability; and ODIN (Liang, Li, and Srikant 2018) pre-processes the inputs and rescales the logits with temperature. Recently, a few efficient methods like ViM (Wang et al. 2022a) and DML (Zhang and Xiang 2023) decouple outputs for mining meaningful information. In multi-label scenarios, JointEnergy (Wang et al. 2021) is the only inference technique, which jointly estimates the OOD uncertainty across all labels. While **outlier exposure** uses certain auxiliary data that disjoint with in-distribution samples for model training (Hendrycks, Mazeika, and Dietterich 2019). Liu *et al.* (Liu et al. 2020) propose a framework for fine-tuning energy-based classifiers. Besides, some other methods like VOS (Du et al. 2022) and POEM (Ming, Fan, and Li 2022) adaptively sample virtual outliers to regularize decision boundary. These studies assume that feature representation forms a class-conditional Gaussian distribution. Considering the underexplored outlier exposure and the problem of JointEnergy for multi-label settings, we propose an energy-based learning framework to improve OOD detection.

Long-Tailed OOD Detection

Traditional long-tailed learning advocates re-balancing imbalanced distributions with undersampling (Peng et al. 2019), oversampling (Mullick, Datta, and Das 2019), or re-weighting (Alshammari et al. 2022; Wang et al. 2023b; Yang et al. 2024). However, these methods proved to be inapplicable in (1) OOD detection (Wang et al. 2022b) and (2) multi-label settings (Duarte, Rawat, and Shah 2021), which could not be simply transferred to long-tailed OOD detection. Also, the research for imbalanced multi-label learning (Guo and Wang 2021; Wu et al. 2020) is not well-suitable. Fortunately, Wang *et al.* (Wang et al. 2022b) propose the first learning framework to instruct OOD detection on multiclass

long-tailed scenarios. Inspired by this work, this paper attempts to explore the study of long-tailed OOD detection for multi-label scenarios.

Motivation

As illustrated in Figure 1, JointEnergy could simply discriminate the OOD sample from head ones due to the large energy gap between them. However, this success collapses when encountering the samples related to the tail classes. Due to the trivial underfitting issue on tail classes, models do not have great confidence to assign these hard samples with high logits, thereby inducing an inferior inference result. Meanwhile, (Wang et al. 2022b) conveys that simply applying long-tailed recognition methods to OOD detection is not quite feasible. We further show that the imbalanced problem could not be solved by simply applying imbalanced multi-label learning methods. The results are shown in Table 1. Specifically, we apply three popular imbalanced multi-label recognition methods: ASL (Ridnik et al. 2021), Focal loss (Lin et al. 2017), and DB Loss (Wu et al. 2020). Except for the BCE method, all other methods degrade the OOD detection performance. This phenomenon shows that simply applying imbalanced multi-label optimization methods is not feasible for multi-label OOD detection. Thus, this problem motivates us to seek a new OOD detection method for multi-label scenarios. In the following sections, we will present a learning framework to alleviate this issue by enlarging the energy difference between tail ID samples and OOD samples.

Preliminaries

Problem Definition

OOD detection could be formed as a binary classification problem. We use \mathcal{D}_{in} to denote the in-distribution and \mathcal{D}_{out} as out-of-distribution dataset. Furthermore, \mathcal{D}_{in} consists of $\mathcal{D}_{\text{in}}^{\text{train}}$ and $\mathcal{D}_{\text{in}}^{\text{test}}$, i.e. $\mathcal{D}_{\text{in}} = \{\mathcal{D}_{\text{in}}^{\text{train}}, \mathcal{D}_{\text{in}}^{\text{test}}\}$, where $\mathcal{D}_{\text{in}}^{\text{train}}$ refers to the training dataset and $\mathcal{D}_{\text{in}}^{\text{test}}$ represents the test dataset. For out-of-distribution data \mathcal{D}_{out} , we term the outlier exposure samples as \mathcal{D}_{OE} and the test OOD dataset as \mathcal{D}_{OOD} , respectively. In practice, two parts of distribution should meet $\mathcal{D}_{\text{in}} \cap \mathcal{D}_{\text{out}} = \emptyset$. The goal for the OOD detector $g_\gamma(\mathbf{x})$ is to identify whether a sample \mathbf{x} belongs to unknown domains or not, which could be done with the following decision function:

$$g_\gamma(\mathbf{x}) = \begin{cases} \text{ID} & S(\mathbf{x}) \geq \gamma, \\ \text{OOD} & S(\mathbf{x}) < \gamma, \end{cases} \quad (1)$$

where $S(\mathbf{x})$ represents the score function and γ refers to threshold. The sample with a high confidence score will be classified as ID and vice versa.

In multi-label OOD detection, we denote the ID input space as $\mathcal{X} \in \mathbb{R}^d$ and label space with C classes as $\mathcal{Y} = \{0, 1\}^C$ with 0 for negative and 1 for positive. Let $\mathcal{D}_{\text{in}} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$ be the ID dataset with N data points which are all sampled *i.i.d* from the joint data distribution $\mathcal{P}_{\mathcal{X}\mathcal{Y}}$. Each label vector could be represented as $\mathbf{y} = [y_1, \dots, y_C]$. For OOD dataset, we define $\mathcal{D}_{\text{out}} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^{N'}$ includes N'

Datasets	Methods	FPR95	AUROC	AUPRC	mAP
PASCAL	DBLoss	86.90	77.49	83.37	86.69
	Focal	32.13	95.25	96.21	86.04
	ASL	32.06	94.90	95.78	87.65
	BCE	4.50	98.75	98.69	87.16
COCO	DBLoss	99.96	50.09	83.51	74.11
	Focal	89.79	73.21	95.58	72.54
	ASL	81.12	78.96	96.67	74.52
	BCE	33.81	95.61	96.77	74.25

Table 1: OOD detection results using DenseNet121 with naive multi-label learning methods. All results are evaluated on the `Texture`. The best results are shown in bold.

samples with only one label. Assume $\mathbf{z} = (\mathbf{x}, \mathbf{y}) \in \mathcal{P}_{\mathcal{X}\mathcal{Y}}$, we define $f: \mathcal{X} \rightarrow \mathcal{Y}$ as a mapping from sample input to its logit output. On top of the mapping function f , we define a standard pre-trained deep neural model as follows:

$$f(\mathbf{x}; \theta) = \mathbf{w}_{\text{cls}}^\top \cdot h(\mathbf{x}), \quad (2)$$

where θ is the parameter of model, $\mathbf{w}_{\text{cls}} \in \mathbb{R}^{d_1 \times C}$ is the classifier weight, and $h(\mathbf{x}) \in \mathbb{R}^{d_1}$ is the penultimate layer feature *w.r.t* sample \mathbf{x} . For MaxLogit, $S(\mathbf{x}) = \max f_{y_i}(\mathbf{x})$, which only considers the maximum logit as OOD discrimination criteria. For multi-label OOD detection, JointEnergy (Wang et al. 2021) is the unique energy-based score function, whose specific formula is presented as follows:

$$\mathcal{E}(\mathbf{z}) = \sum_{i=1}^C \log(1 + e^{f_{y_i}(\mathbf{x})}), \quad (3)$$

where $f_{y_i}(\mathbf{x})$ refers to the logit output on i -th class of \mathbf{x} . Unlike previous studies concerning the maximum logit, this method jointly estimates the OOD uncertainty by summing the label-wise energy across all labels, accumulating all label confidence to amplify the confidence difference between ID and OOD data effectively. However, this method would break down when encountering imbalanced multi-label data.

Methodology

In this section, we will concretely introduce our proposed method EDGE in three parts. Then, a feature-based outlier selection module is presented to guide auxiliary OOD sample selection.

Energy Distribution Gap Expansion

The key to discriminating the OOD samples from ID ones lies in the ways of expanding the energy difference. Without knowing any open-set information, it is hard to determine a precise energy boundary for the whole feature space. Thus, we establish an effective energy-based learning framework in the multi-label setting, which consists of three main parts as follows:

Acquire Known Distribution Information. In order to enhance detection generalization, the first thing is to acquire useful knowledge from known domains. Here we incorporate the Binary Cross-Entropy (BCE), which is a traditional

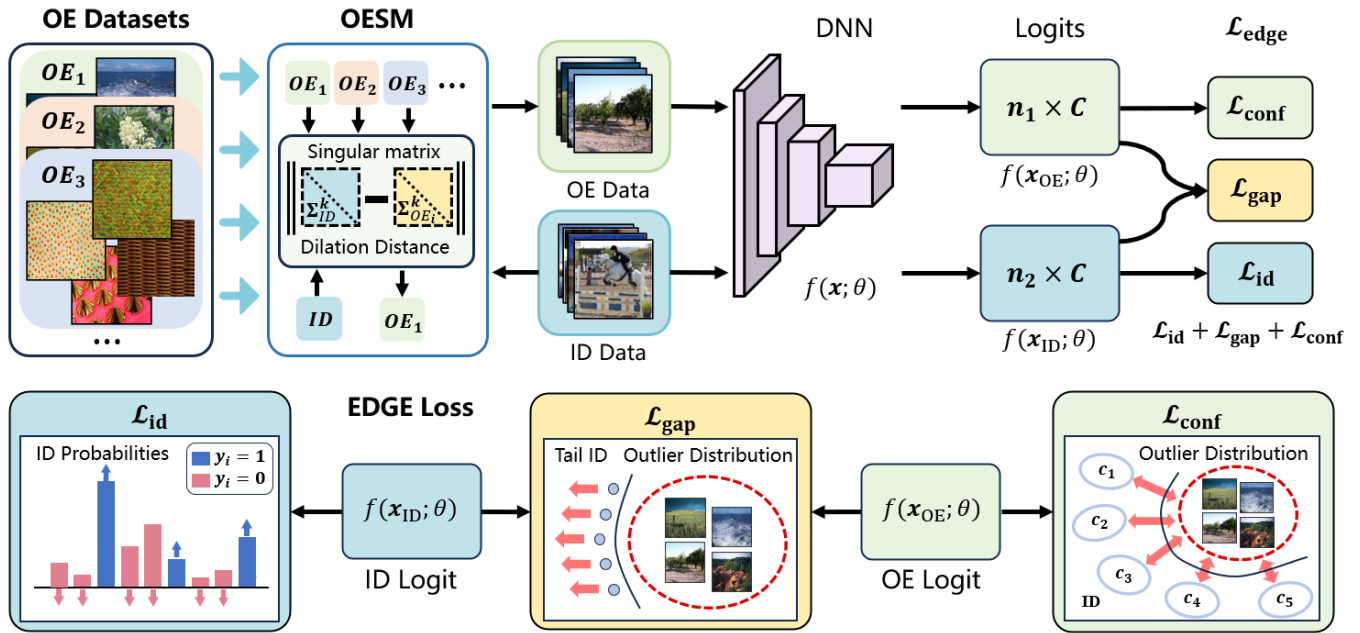


Figure 2: Overview of EDGE learning framework. This process could be divided into two parts, *i.e.*, auxiliary OE data selection, and EDGE loss optimization. Firstly, the candidate OE datasets are fed into the OE Selection Module (OESM, shown in Fig. 3). Then, the deep model is provided with both in-distribution data and an informative auxiliary dataset to conduct unknown-aware multi-label learning with the two-part EDGE loss.

loss function in multi-label learning, as our foundational in-distribution objective, which is termed as \mathcal{L}_{id} :

$$\mathcal{L}_{id} = \mathbb{E}_{z \sim \mathcal{D}_{in}^{train}} \left[\frac{1}{C} \sum_{i=1}^C -y_i \mathcal{L}_+(\mathbf{x}) - (1 - y_i) \mathcal{L}_-(\mathbf{x}) \right], \quad (4)$$

where $\mathcal{L}_+(\mathbf{x}) = \log \sigma(\mathbf{x})$, $\mathcal{L}_-(\mathbf{x}) = \log(1 - \sigma(\mathbf{x}))$, and $\sigma(\mathbf{x}) = 1/(1 + \exp(-f(\mathbf{x}; \theta)))$ is the sigmoid function *w.r.t* sample z .

Segregate Unknown Samples. To mitigate the issue of OOD samples yielding excessively high scores, a straightforward approach is to suppress their logits across all categories. Motivated by this, we propose the confidence optimization objective \mathcal{L}_{conf} as follows:

$$\mathcal{L}_{conf} = \mathbb{E}_{z' \sim \mathcal{D}_{OE}} \left[\frac{1}{C} \sum_{i=1}^C -\log(1 - \sigma(\mathbf{x}')) \right], \quad (5)$$

where z' refers to an OE instance. Considering that outlier exposure data do not relate to any known class, it is reasonable for models to assign a negative label on all known classes. Consequently, we have formulated this entropy-based loss function to refine the optimization process. This objective focuses on pushing OOD samples away from the known cluster centers.

Enlarge Distribution Distance. The reduction in logits for OE samples allows for the accurate identification of most OOD samples. Yet, a subset of challenging cases persists among tail ID samples and semantically similar near-boundary OOD samples (Wang et al. 2023a). Our analysis

suggests that this problem stems from a critically narrow gap in energy distribution, as depicted in Figure 1 (a). The considerable overlap in distribution areas culminates in ambiguous distinctions. A type of partial contrastive learning (Wang et al. 2022b) proves to be a potent method for distancing tail ID samples from their OOD counterparts. Nonetheless, its assumption of class-conditional independence is untenable for the unknown joint distributions encountered in multi-label settings. In response to these problems, we opt to treat the k ID samples with the lowest scores collectively in contrast with all OE sample, and propose an energy-based loss function \mathcal{L}_{gap} to expand the energy gap:

$$\mathcal{L}_{gap} = \mathbb{E}_{z' \sim \mathcal{D}_{OE}} \left[\mathcal{E}(z') - \mathcal{E}_k^\downarrow(z) + m \right]_+, \quad (6)$$

where \mathcal{E} refers to the JointEnergy, $[a]_+ = \max\{0, a\}$, and m is the margin hyper-parameter. Assume there are n samples in \mathcal{D}_{in}^{train} , and we define $\mathcal{E}_{[k]}$ as top- k JointEnergy scores among all ID samples (Lyu et al. 2022). Then, the average bottom- k JointEnergy score of ID samples $\mathcal{E}_k^\downarrow(z)$ could be denoted as:

$$\mathcal{E}_k^\downarrow(z) = \frac{1}{k} \sum_{i=0}^{k-1} \mathcal{E}_{[n-i]}(z). \quad (7)$$

Training Objective. Benefiting the above improvements, we finally formalize our EDGE objective based on entropy-based and energy-based learning and present it as follows:

$$\mathcal{L}_{edge} = \mathcal{L}_{id} + \alpha \cdot \mathcal{L}_{conf} + \beta \cdot \mathcal{L}_{gap}, \quad (8)$$

where α, β are hyper-parameters to control weights.

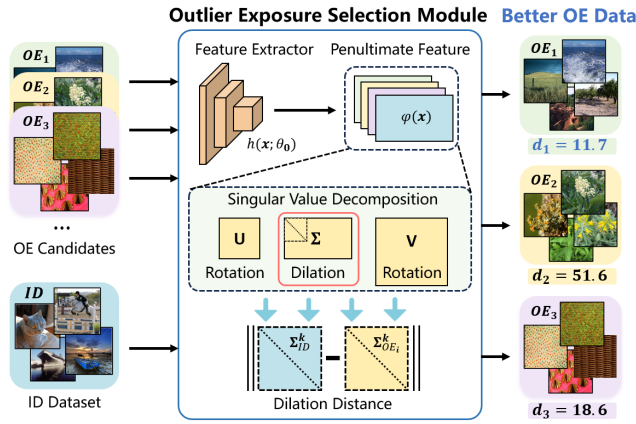


Figure 3: Illustration of Outlier Exposure Selection Module.

OE Selection via Feature-Based Distance

It has been proved that a certain amount of auxiliary out-of-distribution data would extremely boost the discrimination of OOD detectors (Hendrycks, Mazeika, and Dietterich 2019). In this case, the model establishes a sufficiently strong decision boundary by previewing information from the real distribution. However, the detection effect is uncertain when using different \mathcal{D}_{OE} for training, which indicates that the success of outlier exposure is quite dependent on proper auxiliary data. Thus, a key challenge in our framework is *how to select a more reliable auxiliary dataset?* By observing the feature space, it could be found the samples with similar semantic information usually assemble (Du et al. 2022). This trivial nature results in a serious failure to distinguish the samples with strong semantic correlation. Thus, we argue that these OOD samples could play a decisive role in improving discrimination. Based on this discovery, we propose a reasonable feature-based dilation distance to measure the suitability of outlier samples:

$$d_{\text{dilation}} = \left\| \Sigma_{\text{in}}^k - \Sigma_{\text{out}}^k \right\|_F, \quad (9)$$

where Σ_{in}^k and Σ_{out}^k are the diagonal matrix with top- k singular values of ID and OOD penultimate layer feature, respectively. Motivated by current works (Du et al. 2022; Ming, Fan, and Li 2022), we think the feature representation space covers rich information, where prime singular values have a vital effect on the dilation degrees at different dimensions. Thus, an intuitive idea is that “the closer the singular vector of two couple of samples, the more similar their feature embedding is”. Based on this finding, we could select a more informative outlier exposure dataset. To be specific, (1) we first set the batch size of ID data equal to that of OE data to ensure that the output dimensions are the same; (2) Then, solve the singular matrix separately for the two-part feature matrices; (3) Subtract the two singular matrices obtained from multiple batches, and calculate their average Frobenius norm; (4) Finally, calculate the mean of all norm results and choose the data with smaller d_{dilation} for training. We present our selection process in Figure 3 and the whole training process in Algorithm 1.

Algorithm 1: Energy Distribution Gap Expansion (EDGE)

Input: ID training set $\mathcal{D}_{\text{in}}^{\text{train}}$, OE dataset \mathcal{D}_{OE} ;
Parameter: model parameter θ , epoch T , transform epoch τ , hyper-parameter α, β, k, m ;

- 1: Compute *dilation distance* between $\mathcal{D}_{\text{in}}^{\text{train}}$ and each individual set \mathcal{D}_{OE} by Eq. (9).
- 2: **for** $t = 0, 1, \dots, T$ **do**
- 3: Sample a batch of ID and OE training data from $\mathcal{D}_{\text{in}}^{\text{train}}$ and \mathcal{D}_{OE} respectively.
- 4: **if** $t < \tau$ **then**
- 5: Set $\beta = 0$.
- 6: **end if**
- 7: Calculate each loss using Eq. (4), (5), and (6).
- 8: Update parameter θ by minimizing (8).
- 9: **end for**

\mathcal{D}_{out}	\mathcal{D}_{in}		
	PASCAL	MS-COCO	NUS-WIDE
Texture	51.56	49.82	43.06
MNIST	67.97	67.70	68.13
iSUN	18.60	22.39	21.00
Places50	14.88	23.05	17.13
LSUN	34.88	40.33	36.38
SVHN	39.24	42.34	41.28
iNaturalist	15.50	20.40	16.20
Img-20-R	16.90	23.05	13.60
Img-20-G	11.75	9.55	/

Table 2: Feature-based Distance between ID and OOD dataset.

Experiments

In this section, we present and analyze the overall OOD detection performance of EDGE framework. **Due to the space limitation, we present partial vital details and results in the main paper and put others in Appendix.**

Experimental Settings

In-distribution Datasets. We adopt three real-world benchmark multi-label image annotation datasets as our training sets, *i.e.*, PASCAL-VOC (Everingham et al. 2015), MS-COCO (Lin et al. 2014) and NUS-WIDE (Chua et al. 2009). PASCAL-VOC consists of 22,531 images with 20 categories; MS-COCO contains 122,218 annotation images with 80 categories, where 82,783 images are for training, 40,504 for validation, and 40,775 for testing; and NUS-WIDE covers 269,648 real-world images with 81 categories, where we reserve 119,986 images of the training set and 80,283 images of the test set, respectively.

Out-of-distribution Datasets. To fully evaluate the robustness of the model on out-of-distribution examples, we follow previous literature and choose the following benchmark OOD datasets as main experimental candidates: (a subset of) ImageNet-22K (Deng et al. 2009), Textures (Cimpoi et al. 2014), Places50 (Zhou et al. 2018), iSUN (Xu et al. 2015), iNaturalist (Horn et al. 2018), LSUN-C (Yu et al. 2015), and SVHN (Netzer et al. 2011), where the

Training	Inference	LSUN			Texture			iSUN			mAP ↓
		FPR95 ↓	AUROC ↑	AUPR ↑	FPR95 ↓	AUROC ↑	AUPR ↑	FPR95 ↓	AUROC ↑	AUPR ↑	
BCE	MaxLogit	1.08	99.63	99.90	11.19	96.97	99.44	29.80	89.93	96.02	73.79
	MSP	22.41	95.13	98.65	45.51	87.64	97.62	60.58	79.99	92.89	
	ODIN	1.31	99.23	99.81	11.51	96.31	99.28	33.42	88.70	95.71	
	Mahalanobis	10.18	98.02	99.54	18.21	96.28	99.42	97.78	63.25	88.26	
	Energy Score	0.75	99.72	99.92	9.11	97.33	99.49	27.51	90.28	96.16	
	ReAct	3.38	98.98	99.69	8.58	97.78	99.62	27.39	89.88	95.60	
	ViM	0.27	99.91	99.98	5.76	98.31	99.71	36.41	89.01	96.04	
	ASH	1.38	98.88	98.90	8.24	97.82	99.61	30.45	90.21	96.81	
	MaxEnergy	1.08	99.63	99.90	11.19	96.97	99.44	29.80	89.83	96.02	
JointEnergy	0.48	99.80	99.95	7.20	98.15	99.71	<u>26.05</u>	<u>92.04</u>	<u>97.47</u>		
OE	MaxEnergy	5.79	98.78	99.72	9.38	98.09	99.72	25.42	93.41	97.95	68.47
	JointEnergy	69.19	65.80	81.33	60.30	71.41	89.58	57.67	78.84	88.62	
EnergyOE	MaxEnergy	18.01	97.28	99.34	27.85	92.92	98.69	49.29	86.36	95.99	72.37
	JointEnergy	33.51	81.38	88.01	38.60	80.49	92.37	46.29	81.52	89.35	
EDGE	MaxEnergy	0.07	99.95	99.99	4.18	98.90	99.82	20.15	93.90	97.84	73.50
	JointEnergy	0.01	99.91	99.98	3.09	99.17	99.88	18.45	94.86	98.40	

Table 3: Comparison results on MS-COCO with ResNet-50. The **best** and **runner-up** performance are highlighted, and the best inference methods are underlined. **Better result of OE methods** are bold. ↑ means the larger the better, and ↓ is opposite.


categories of all samples do not overlap with those in used ID datasets and ImageNet-1K (Deng et al. 2009). Owing to the outlier exposure selection module, we present the distance between each ID and OOD dataset in Table 2, and select Img-20-G as \mathcal{D}_{OE} . Note that NUS-WIDE overlaps some classes of Img-20-G , we thus follow (Wang et al. 2021) and adopt another different subset Img-20-R . More details about OOD datasets can be found in Appendix.

Implementation Details. All the experiments are run on NVIDIA GeForce RTX 3090 implemented by PyTorch and conducted on the ResNet-50 (He et al. 2016) and VGG16 (Simonyan and Zisserman 2015) pre-trained on ImageNet-1K (Deng et al. 2009). We randomly horizontally flip all input images and resize them to 256×256 . We use stochastic gradient descent (SGD) (Sutskever et al. 2013) with a momentum of 0.9 and a weight decay of 1×10^{-4} . For trivial BCE models, we use the Adam (Kingma and Ba 2015) to optimize. The main hyper-parameters α , β , and m take value from $\{1e+1, 1e-1, 1e-2, 1e-3\}$, $\{1, 1e-1, 1e-2, 1e-3, 1e-4\}$, and $\{0, 1, 2, 3, 4, 5\}$.

Competitors. Totally, we selected 10 inference-based approaches as comparisons, including MaxLogit (Hendrycks et al. 2022), MSP (Hendrycks and Gimpel 2017), ODIN (Liang, Li, and Srikant 2018), Energy (Liu et al. 2020), Mahalanobis (Lee et al. 2018), ReAct (Sun, Guo, and Li 2021), ViM (Wang et al. 2022a), ASH (Djurisic et al. 2023), MaxEnergy, and JointEnergy (Wang et al. 2021). Note that all results obtained by inference methods are based on the model optimizing with BCE loss. We adopt OE (Hendrycks, Mazeika, and Dietterich 2019) and EnergyOE (Liu et al. 2020) as outlier exposure comparisons.

Main Results

Overall Performance. We present the overall comparative results on various ID and OOD datasets in Table 3. We test various inference methods as well as different outlier ex-



(a) w/o EDGE, Texture (b) w/ EDGE, Texture
Figure 4: UMAP visualization of feature embedding results.

posure methods under the same training criterion and pre-trained model. We get some observations below: (1) Our proposed EDGE significantly outperforms all comparative methods; (2) Our EDGE achieves a significant improvement in OOD detection with a less loss of ID performance, promising a relatively balanced result between known and unknown domains; (3) The complexity of the \mathcal{D}_{in} and \mathcal{D}_{OE} directly affects both ID classification and OOD detection. In larger-scale NUS-WIDE, most comparative methods produce serious degeneration, especially on iSUN, where this dataset contains some similar scene images, like the jungle, with the images in NUS-WIDE. However, our method still maintains promising performance while resisting hard samples; (4) An interesting point is that both OE and EnergyOE achieve an unstable performance, even poorer than some post-hoc methods. It indicates that simply transferring general outlier exposure methods is not always feasible; (5) OE and EnergyOE perform better when utilizing MaxEnergy while not JointEnergy. A reasonable explanation is that traditional methods are applied in multi-class scenarios and prefer to identify OOD samples with their maximum logit. This nature makes them function poorly in multi-label scenarios.

Visualization. To further show the expansion effect, we visualize the results of feature embedding output in Figure

\mathcal{D}_{in}	\mathcal{L}_{bce}	\mathcal{L}_{conf}	\mathcal{L}_{gap}	Places50			Texture			iSUN			mAP \uparrow
				FPR95 \downarrow	AUROC \uparrow	AUPR \uparrow	FPR95 \downarrow	AUROC \uparrow	AUPR \uparrow	FPR95 \downarrow	AUROC \uparrow	AUPR \uparrow	
MS-COCO	✓			0.48	99.80	99.95	7.20	98.15	99.71	26.05	92.04	97.46	73.79
	✓	✓		0.02	99.90	99.95	5.38	98.69	99.77	20.39	93.03	97.54	68.18
	✓		✓	0.01	99.93	99.98	8.46	97.73	99.61	43.14	86.93	95.61	57.98
	✓	✓	✓	0.01	99.93	99.98	3.09	99.17	99.88	18.45	94.86	98.40	73.50

Table 4: Ablative results on different combinations of loss functions, where \checkmark means the utilization of the corresponding loss term. The **best** and **runner-up** OOD performance is highlighted.

\mathcal{D}_{OE}	$d_{dilation} \downarrow$	\mathcal{D}_{OOD}	Texture	iSUN	Img-20-R	Places50	iNaturalist	MNIST	SVHN	Average
Texture	43.06	F \downarrow	/	70.92	30.00	67.83	34.31	9.70	0.20	35.49
iSUN	21.00		10.62	/	15.73	5.01	8.28	5.78	2.25	7.95
Img-20-R	13.61		2.46	9.10	/	26.36	0.79	0.00	0.00	6.45
Texture	43.06	AR \uparrow	/	67.65	92.47	69.20	90.74	97.23	99.65	86.16
iSUN	21.00		97.16	/	97.03	94.40	98.02	97.69	99.02	97.21
Img-20-R	13.61		99.34	96.92	/	89.35	99.83	99.85	99.91	97.53

Table 5: FPR95 (F \downarrow) and AUROC (AR \uparrow) Performance of OOD detection and ID classification when using different datasets as \mathcal{D}_{OE} and \mathcal{D}_{OOD} on NUS-WIDE. The **best results** on the same \mathcal{D}_{OOD} are bold.

4. The blue and orange parts represent the cluster of ID data and OOD data, respectively. We qualitatively analyze the discriminative effect when model training with/without our EDGE framework. Compared with general models, the EDGE-based model could accurately separate OOD samples from ID data and cluster them. This demonstrates the superiority of our proposed method.

Performance on Tail Samples. Figure 5 shows several groups of comparative results on tail samples on PASCAL-VOC. On top of the overall improvement, our method achieves significant tail improvement. As the proportion of tail samples increases, the performance of EDGE is less negatively impacted. These results indicate the solving of the problem that is illustrated in the motivation.

Effectiveness of OE Selection. We record the performance of OOD detection based on different outlier exposure data in Table 5. The overall OOD performance on Img-20-R is obviously better than others, where we observe it also holds a relatively small $d_{dilation}$. Moreover, for any two datasets, selecting the one with a smaller $d_{dilation}$ as \mathcal{D}_{OE} could achieve better performance. These intuitive results are consistent with our elaboration in the previous section and validate the effectiveness of the selection method.

Sensitivity Analysis. We conducted a series of hyperparameter sensitivity analyses on PASCAL-VOC. Due to space constraints, we include all results in Appendix.

Ablation Studies

To further test the impact of each part in EDGE framework, we carry out the ablation study on three benchmarks and report the results in Table 4. From this table, we could obtain the following observations: (1) Whichever of \mathcal{L}_{conf} and \mathcal{L}_{gap} is combined with \mathcal{L}_{id} could lead to a significant improvement, demonstrating the improvement of increments. (2) While \mathcal{L}_{conf} and \mathcal{L}_{gap} take turns in the lead, the former has a larger effect on keeping the performance of ID classification, indicating that optimizing the logit information plays

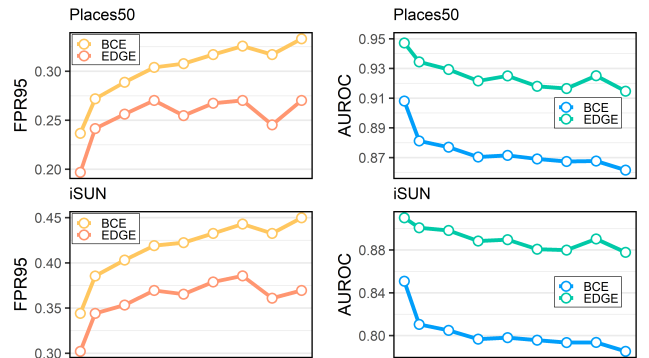


Figure 5: Performance change curve on PASCAL-VOC. The closer to the right, the larger the proportion of tail samples.

a dominant role. (3) The combination of all loss terms could achieve not only a relatively comprehensive improvement in OOD detection, but also a better ID classification performance compared to other OE methods. In general, each module explicitly displays a certain positive effect.

Conclusion

In this paper, we discover the shortcomings in multi-label OOD detection, *i.e.*, a serious long-tailed phenomenon caused by trivial optimization methods. To fix this issue, we introduce outlier exposure into multi-label classification. On top of this, we propose a simple yet effective learning framework dubbed EDGE. This method first separately optimizes ID samples and OE samples to establish a reliable performance of ID classification. It then expands the energy gap between ID and OOD samples by an energy-boundary learning strategy. Besides, a feature-based selection module is devised to boost the performance of model learning. Finally, extensive experiments consistently validate the effectiveness of the proposed method on both ID classification and OOD detection.

Acknowledgments

This work was supported in part by the National Key R&D Program of China under Grant 2018AAA0102000, in part by National Natural Science Foundation of China: 62236008, U21B2038, U23B2051, 61931008, 62122075, 62206264, and 92370102, in part by Youth Innovation Promotion Association CAS, in part by the Strategic Priority Research Program of the Chinese Academy of Sciences, Grant No. XDB0680000, in part by the Innovation Funding of ICT, CAS under Grant No.E000000, in part by the China National Postdoctoral Program for Innovative Talents under Grant BX20240384.

References

- Alshammari, S.; Wang, Y.; Ramanan, D.; and Kong, S. 2022. Long-Tailed Recognition via Weight Balancing. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6887–6897.
- Chen, Z.; Wei, X.; Wang, P.; and Guo, Y. 2019. Multi-Label Image Recognition With Graph Convolutional Networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 5177–5186.
- Chua, T.; Tang, J.; Hong, R.; Li, H.; Luo, Z.; and Zheng, Y. 2009. NUS-WIDE: a real-world web image database from National University of Singapore. In *ACM International Conference on Image and Video Retrieval*.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing Textures in the Wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3606–3613.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 248–255.
- Djurisic, A.; Bozanic, N.; Ashok, A.; and Liu, R. 2023. Extremely Simple Activation Shaping for Out-of-Distribution Detection. In *International Conference on Learning Representations*.
- Du, X.; Wang, Z.; Cai, M.; and Li, Y. 2022. VOS: Learning What You Don’t Know by Virtual Outlier Synthesis. In *International Conference on Learning Representations*.
- Duarte, K.; Rawat, Y. S.; and Shah, M. 2021. PLM: Partial Label Masking for Imbalanced Multi-Label Classification. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2739–2748.
- Everingham, M.; Eslami, S. M. A.; Gool, L. V.; Williams, C. K. I.; Winn, J. M.; and Zisserman, A. 2015. The Pascal Visual Object Classes Challenge: A Retrospective. *International Journal of Computer Vision*, 111(1): 98–136.
- Gong, Y.; Jia, Y.; Leung, T.; Toshev, A.; and Ioffe, S. 2014. Deep Convolutional Ranking for Multilabel Image Annotation. In *International Conference on Learning Representations*.
- Guo, H.; and Wang, S. 2021. Long-Tailed Multi-Label Visual Recognition by Collaborative Training on Uniform and Re-Balanced Samplings. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Conference on Computer Vision and Pattern Recognition*, 770–778.
- Hendrycks, D.; Basart, S.; Mazeika, M.; Zou, A.; Kwon, J.; Mostajabi, M.; Steinhardt, J.; and Song, D. 2022. Scaling Out-of-Distribution Detection for Real-World Settings. In *International Conference on Machine Learning*, volume 162, 8759–8773.
- Hendrycks, D.; and Gimpel, K. 2017. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. In *International Conference on Learning Representations*.
- Hendrycks, D.; Mazeika, M.; and Dietterich, T. G. 2019. Deep Anomaly Detection with Outlier Exposure. In *International Conference on Learning Representations*.
- Horn, G. V.; Aodha, O. M.; Song, Y.; Cui, Y.; Sun, C.; Shepard, A.; Adam, H.; Perona, P.; and Belongie, S. J. 2018. The INaturalist Species Classification and Detection Dataset. In *IEEE Conference on Computer Vision and Pattern Recognition*, 8769–8778.
- Huang, J.; Xu, L.; Wang, J.; Feng, L.; and Yamanishi, K. 2020. Discovering Latent Class Labels for Multi-Label Learning. In *International Joint Conference on Artificial Intelligence*, 3058–3064.
- Huang, R.; and Li, Y. 2021. MOS: Towards Scaling Out-of-Distribution Detection for Large Semantic Space. In *IEEE Conference on Computer Vision and Pattern Recognition*, 8710–8719.
- Jernite, Y.; Choromanska, A.; and Sontag, D. A. 2017. Simultaneous Learning of Trees and Representations for Extreme Classification and Density Estimation. In *International Conference on Machine Learning*, volume 70, 1665–1674.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*.
- Lee, K.; Lee, K.; Lee, H.; and Shin, J. 2018. A Simple Unified Framework for Detecting Out-of-Distribution Samples and Adversarial Attacks. In *Advances in Neural Information Processing Systems*, 7167–7177.
- Liang, S.; Li, Y.; and Srikant, R. 2018. Enhancing The Reliability of Out-of-distribution Image Detection in Neural Networks. In *International Conference on Learning Representations*.
- Lin, T.; Goyal, P.; Girshick, R. B.; He, K.; and Dollár, P. 2017. Focal Loss for Dense Object Detection. In *IEEE International Conference on Computer Vision*, 2999–3007.
- Lin, T.; Maire, M.; Belongie, S. J.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common Objects in Context. In *European Conference on Computer Vision*, 740–755.
- Liu, W.; Wang, H.; Shen, X.; and Tsang, I. W. 2022. The Emerging Trends of Multi-Label Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11): 7955–7974.

- Liu, W.; Wang, X.; Owens, J. D.; and Li, Y. 2020. Energy-based Out-of-distribution Detection. In *Annual Conference on Neural Information Processing Systems*.
- Lyu, S.; Fan, Y.; Ying, Y.; and Hu, B. 2022. Average Top-k Aggregate Loss for Supervised Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1): 76–86.
- Mao, J.; Shi, S.; Wang, X.; and Li, H. 2023. 3D Object Detection for Autonomous Driving: A Comprehensive Survey. *International Journal of Computer Vision*, 131(8): 1909–1963.
- Ming, Y.; Fan, Y.; and Li, Y. 2022. POEM: Out-of-Distribution Detection with Posterior Sampling. In *International Conference on Machine Learning*, volume 162, 15650–15665.
- Mullick, S. S.; Datta, S.; and Das, S. 2019. Generative Adversarial Minority Oversampling. In *IEEE/CVF International Conference on Computer Vision*, 1695–1704.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. In *Reading digits in natural images with unsupervised feature learning*.
- Nguyen, A. M.; Yosinski, J.; and Clune, J. 2015. Deep neural networks are easily fooled: high confidence predictions for unrecognizable images. In *IEEE Conference on Computer Vision and Pattern Recognition*, 427–436.
- Peng, M.; Zhang, Q.; Xing, X.; Gui, T.; Huang, X.; Jiang, Y.; Ding, K.; and Chen, Z. 2019. Trainable Undersampling for Class-Imbalance Learning. In *AAAI Conference on Artificial Intelligence*, 4707–4714.
- Ridnik, T.; Baruch, E. B.; Zamir, N.; Noy, A.; Friedman, I.; Protter, M.; and Zelnik-Manor, L. 2021. Asymmetric Loss For Multi-Label Classification. In *IEEE/CVF International Conference on Computer Vision*, 82–91.
- Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *International Conference on Learning Representations*.
- Sun, Y.; Guo, C.; and Li, Y. 2021. ReAct: Out-of-distribution Detection With Rectified Activations. In *Advances in Neural Information Processing Systems*, 144–157.
- Sutskever, I.; Martens, J.; Dahl, G. E.; and Hinton, G. E. 2013. On the importance of initialization and momentum in deep learning. In *International Conference on Machine Learning*, volume 28, 1139–1147.
- Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E. B.; and Garnett, R. 2019. Multilabel reductions: what is my loss optimising? In *Advances in Neural Information Processing Systems*, 10599–10610.
- Wang, H.; Li, Z.; Feng, L.; and Zhang, W. 2022a. ViM: Out-Of-Distribution with Virtual-logit Matching. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4911–4920.
- Wang, H.; Liu, W.; Bocchieri, A.; and Li, Y. 2021. Can multi-label classification networks know what they don't know? In *Advances in Neural Information Processing Systems*, 29074–29087.
- Wang, H.; Zhang, A.; Zhu, Y.; Zheng, S.; Li, M.; Smola, A. J.; and Wang, Z. 2022b. Partial and Asymmetric Contrastive Learning for Out-of-Distribution Detection in Long-Tailed Recognition. In *International Conference on Machine Learning*, volume 162, 23446–23458.
- Wang, J.; Yang, Y.; Mao, J.; Huang, Z.; Huang, C.; and Xu, W. 2016. CNN-RNN: A Unified Framework for Multi-label Image Classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2285–2294.
- Wang, Z.; Xu, Q.; Yang, Z.; He, Y.; Cao, X.; and Huang, Q. 2022c. OpenAUC: Towards AUC-Oriented Open-Set Recognition. In *Advances in Neural Information Processing Systems*.
- Wang, Z.; Xu, Q.; Yang, Z.; He, Y.; Cao, X.; and Huang, Q. 2023a. Optimizing Partial Area Under the Top-k Curve: Theory and Practice. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4): 5053–5069.
- Wang, Z.; Xu, Q.; Yang, Z.; He, Y.; Cao, X.; and Huang, Q. 2023b. A Unified Generalization Analysis of Re-Weighting and Logit-Adjustment for Imbalanced Learning. In *Advances in Neural Information Processing Systems*.
- Wang, Z.; Xu, Q.; Yang, Z.; Wen, P.; He, Y.; Cao, X.; and Huang, Q. 2024. Top-K Pairwise Ranking: Bridging the Gap Among Ranking-Based Measures for Multi-Label Classification. *CoRR*, abs/2407.06709.
- Wu, T.; Huang, Q.; Liu, Z.; Wang, Y.; and Lin, D. 2020. Distribution-Balanced Loss for Multi-label Classification in Long-Tailed Datasets. In *European Conference on Computer Vision*, volume 12349, 162–178.
- Xu, P.; Ehinger, K. A.; Zhang, Y.; Finkelstein, A.; Kulkarini, S. R.; and Xiao, J. 2015. TurkerGaze: Crowdsourcing Saliency with Webcam based Eye Tracking. *CoRR*, abs/1504.06755.
- Yang, J.; Zhou, K.; Li, Y.; and Liu, Z. 2021. Generalized Out-of-Distribution Detection: A Survey. *CoRR*, abs/2110.11334.
- Yang, Z.; Xu, Q.; Wang, Z.; Li, S.; Han, B.; Bao, S.; Cao, X.; and Huang, Q. 2024. Harnessing Hierarchical Label Distribution Variations in Test Agnostic Long-tail Recognition. In *International Conference on Machine Learning*.
- Yu, F.; Zhang, Y.; Song, S.; Seff, A.; and Xiao, J. 2015. LSUN: Construction of a Large-scale Image Dataset using Deep Learning with Humans in the Loop.
- Zhang, Z.; and Xiang, X. 2023. Decoupling MaxLogit for Out-of-Distribution Detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3388–3397.
- Zhou, B.; Lapedriza, À.; Khosla, A.; Oliva, A.; and Torralba, A. 2018. Places: A 10 Million Image Database for Scene Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6): 1452–1464.
- Zhou, H.; Lu, C.; Chen, C.; Yang, S.; and Yu, Y. 2023. A Unified Visual Information Preservation Framework for Self-supervised Pre-Training in Medical Image Analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(7): 8020–8035.