Detecting Out-of-distribution Data through In-distribution Class Prior

Xue Jiang [1 2] **Feng Liu** [3] **Zhen Fang** [4] **Hong Chen** [5] **Tongliang Liu** [6 7] **Feng Zheng** [1] **Bo Han** [2]

Abstract

Given a pre-trained *in-distribution* (ID) model, the inference-time *out-of-distribution* (OOD) detection aims to recognize OOD data during the inference stage. However, some representative methods share an unproven assumption that the probability that OOD data belong to every ID class should be the same, i.e., these OOD-to-ID probabilities actually *form a uniform distribution*. In this paper, we show that this assumption makes the above methods incapable when the ID model is trained with *class-imbalanced data*. Fortunately, by analyzing the causal relations between ID/OOD classes and features, we identify several common scenarios where the OOD-to-ID probabilities should be the ID-class-prior distribution and propose two strategies to modify existing inference-time detection methods: 1) *replace* the uniform distribution with the ID-class-prior distribution if they explicitly use the uniform distribution; 2) otherwise, *reweight* their scores according to the similarity between the ID-class-prior distribution and the softmax outputs of the pre-trained model. Extensive experiments show that both strategies can improve the OOD detection performance when the ID model is pre-trained with imbalanced data, reflecting the importance of ID-class prior in OOD detection. The codes are available at https://github.com/tmlr-group/class_prior.

[1]Department of Computer Science and Engineering, Southern University of Science and Technology [2]Department of Computer Science, Hong Kong Baptist University [3]University of Melbourne [4]Australian Artificial Intelligence Institute, University of Technology Sydney [5]College of Informatics, Huazhong Agricultural University [6]Mohamed bin Zayed University of Artificial Intelligence [7]Sydney AI Centre, The University of Sydney. Correspondence to: Feng Zheng <f.zheng@ieee.org>.

Proceedings of the 40^{th} International Conference on Machine Learning, Honolulu, Hawaii, USA. PMLR 202, 2023.

1. Introduction

How to reliably deploy machine learning models into real-world scenarios has been attracting more and more attention (Huang et al., 2021; Liang et al., 2018; Liu et al., 2020). In real-world scenarios, test data usually contain known and unknown classes (Hendrycks & Gimpel, 2017). We expect the deployed model to eliminate the interference of unknown classes while classifying known classes well. Nevertheless, current models tend to be overconfident in the unknown classes (Nguyen et al., 2015), and thus confusing known and unknown classes, which increases the risk of deploying these models in the real world. Especially, if the scenarios are life-critical (e.g., car-driving scenarios), we cannot take the risks of deploying unreliable models in them. This motivates researchers to study *out-of-distribution* (OOD) detection, where we need to identify unknown classes (i.e., OOD classes) and classify known classes (i.e., *in-distribution* (ID) classes) well at the same time (Hendrycks & Gimpel, 2017).

In the OOD detection, a well-known branch is to develop the inference-time/post hoc OOD detection methods (Huang et al., 2021; Liang et al., 2018; Liu et al., 2020; Lee et al., 2018a; Sun et al., 2021), where we are given a pre-trained ID model and then aim to recognize upcoming OOD data well. The key advantage of inference-time OOD detection methods is that the classification performance on ID data will be unaffected since we only use the ID model instead of changing it. A general strategy to design a large-scale-friendly inference-time OOD detection method is to propose a socre function by using the ID model's information. For example, the *maximum softmax probability* (MSP) uses the ID model's outputs (Hendrycks & Gimpel, 2017), and GradNorm uses the ID model's gradients (Huang et al., 2021). If the score of a data point is smaller, then this data point is an OOD data point with a higher probability.

Class imbalance issue is natural in real-world applications (Cao et al., 2019; Park et al., 2022). Models trained on classes with fewer samples are often insufficiently trained, leading to higher inter-class confusion risks. The presence of fragile classification boundaries can affect the OOD detection performance of the model. However, some representative methods (Huang et al., 2021; Hendrycks & Gimpel, 2017) share *an unproven assumption*: the probability that an OOD data point $\mathbf{x}^{\text{out}}$ belongs to each ID class i is always

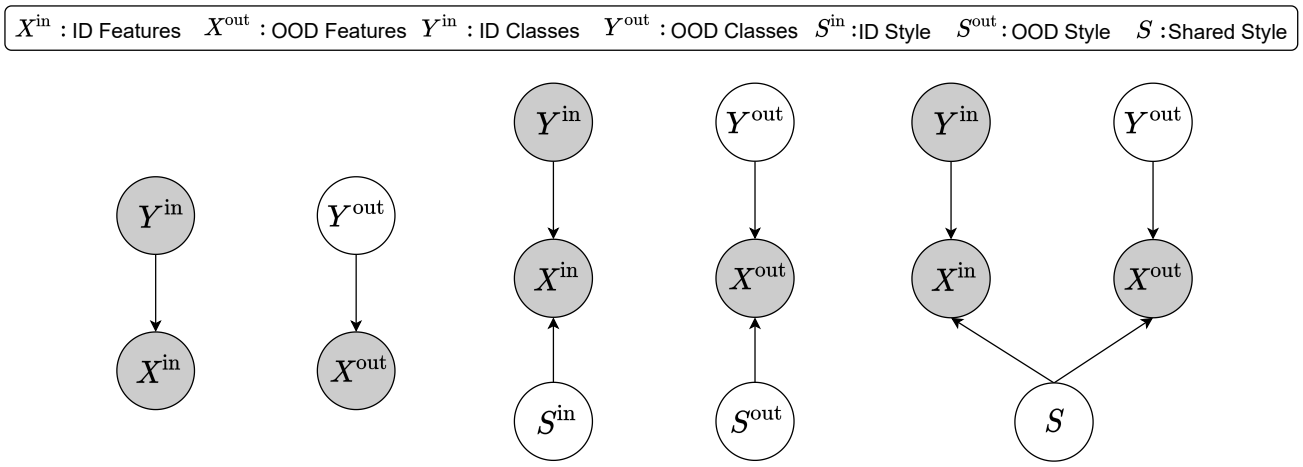

(a) No Styles (b) Different Styles (c) Shared Style

Figure 1: Three common causal graphs in OOD detection. Under these graphs, we prove that probabilities that an OOD data point $\mathbf{x}^{\text{out}}$ belongs to ID classes should be the ID-class-prior distribution $\mathbb{P}_{Y^{\text{in}}}$ (Theorem 3.2). However, some representative OOD detection methods (Huang et al., 2021; Hendrycks & Gimpel, 2017) assume such probabilities to be a uniform distribution $\mathbf{u}$ (e.g., GradNorm in Eq. (2)). In this figure, each node represents a random variable, and gray ones indicate observable variables. X stands for features, Y stands for classes, and S stands for styles. In the three graphs, features are generated by classes (i.e., $Y \to X$) (Gong et al., 2016; Stojanov et al., 2021) or generated by classes and styles (i.e., $Y \to X \leftarrow S$) (Yao et al., 2021). The three causal graphs broadly exist in our common datasets. For example, (a) corresponds to datasets consisting of sketch images, like ImageNet-Sketch (Wang et al., 2019) where ID classes could be cars and OOD classes could be animals; (b) and (c) correspond to datasets consisting of common images, like ImageNet (Deng et al., 2009) and MNIST (LeCun et al., 1998). In (b), the ID classes could be cars in ImageNet, and the OOD classes could be numbers in MNIST (different styles). In (c), the ID classes could be numbers in MNIST, and OOD classes could be classes in Fashion-MNIST (Xiao et al., 2017) (the same style). Through these graphs, it is clear that $Y^{\text{in}} \perp\!\!\!\perp X^{\text{out}}$, i.e., Y^{in} and X^{out} are independent.

the same. Namely, for any $\mathbf{x}^{\text{out}}$,

$$[\mathbb{P}(\text{label of } \mathbf{x}^{\text{out}} \text{ is } 1), \ldots, \mathbb{P}(\text{label of } \mathbf{x}^{\text{out}} \text{ is } K)] = [1/K, \ldots, 1/K]_{1\times K} := \mathbf{u}, \tag{1}$$

where $\mathbf{u}$ is a uniform distribution and K is the number of ID classes. Taking the GradNorm (Huang et al., 2021) as an example[1], let $\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})$ be ID model's output of a data point $\mathbf{x}$, and the score function of GradNorm $S_{\text{GradNorm}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x})$ is

$$\left\| \frac{\partial D_{\text{KL}}(\mathbf{u} \| \operatorname{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})))}{\partial w} \right\|_{L^1}, \tag{2}$$

where $\operatorname{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x}))$ is a vector consisting of predicted probabilities that $\mathbf{x}$ belongs to ID classes, $\mathbf{\Theta}$ represents the parameters of model $\mathbf{f}_{\mathbf{\Theta}}$, w is a component of network parameter $\mathbf{\Theta}$, and $D_{\text{KL}}(\cdot \| \cdot)$ is the Kullback-Leibler divergence function. It is clear that GradNorm considers $\mathbf{u}$ as a *reference distribution* to distinguish between ID and OOD data. If the divergence between $\operatorname{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x}))$ and $\mathbf{u}$ is smaller, then $\mathbf{x}$ is an OOD data point with a higher probability. Nonetheless, since we do not have this assumption proven, we do not know whether it is correct. If not, the $\mathbf{u}$-based score functions (e.g., Eq. (2)) are ill-defined because they cannot guarantee that the lowest score corresponds to the most OOD-ness data.

[1]MSP also follows this assumption, see Section 3.1.

In this paper, we analyze the above assumption (i.e., Eq. (1)) under three common causal graphs (Fig. 1), and find that this assumption holds only when the ID-class prior is $\mathbf{u}$, i.e., the ID model is trained with class-balanced data. In other cases, the reference distribution of OOD data *should be* the ID-class-prior distribution $\mathbb{P}_{Y^{\text{in}}}$ (Theorem 3.2), i.e.,

$$[\mathbb{P}(\text{label of } \mathbf{x}^{\text{out}} \text{ is } 1), \ldots, \mathbb{P}(\text{label of } \mathbf{x}^{\text{out}} \text{ is } K)] = \mathbb{P}_{Y^{\text{in}}}.$$

Specifically, assume that we have K classes in training data (i.e., ID data). Let n_j be the number of samples in class j, then the total number of samples is $N = \sum_j^K n_j$. Thus, we have $\mathbb{P}_{Y^{\text{in}}} = [n_1/N, n_2/N, ..., n_K/N]$.

Empirically, we test the performance of OOD detection methods when the data are not class-balanced (Fig. 2a), i.e., $\mathbb{P}_{Y^{\text{in}}} \neq \mathbf{u}$. We find that the GradNorm will suffer from the imbalanced situation (see cyan and yellow bars in Fig. 2b). Besides, it is interesting to find that Energy (Liu et al., 2020), the other one of representative OOD detection methods that do not explicitly use $\mathbf{u}$, also suffers from this situation (see cyan and yellow bars in Fig. 2c). Based on Theorem 3.2, we propose two effective strategies to modify previous score-based OOD detection methods using the ID-class-prior distribution: *replacing* (RP) strategy and *reweighting* (RW) strategy. In RP strategy, previous methods explicitly use the uniform distribution (like GradNorm), so we can modify them by *replacing* $\mathbf{u}$ with the

ID-class-prior distribution $\mathbb{P}_{Y^{\text{in}}}$. For example, we can modify score functions of GradNorm by replacing $\mathbf{u}$ in Eq. (2) with $\mathbb{P}_{Y^{\text{in}}}$:

$$S_{\text{RP+GradNorm}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) = \left\| \frac{\partial D_{\text{KL}}(\mathbb{P}_{Y^{\text{in}}} \| \operatorname{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})))}{\partial w} \right\|_{L^1}.$$

For the methods that do not explicitly use the uniform distribution to compute scores (like Energy (Liu et al., 2020)), we can use the RW strategy to *reweight* their scores according to the similarity between the ID-class-prior distribution $\mathbb{P}_{Y^{\text{in}}}$ and the softmax outputs of the pre-trained model $\operatorname{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x}))$. Namely, $S_{\text{RW+Method}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x})$ is equal to

$$S_{\text{Method}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) \cdot W(\mathbf{f}_{\mathbf{\Theta}}, \mathbb{P}_{Y^{\text{in}}}),$$

where $S_{\text{Method}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x})$ is a score function proposed in previous studies (like Energy (Liu et al., 2020)). $W(\mathbf{f}_{\mathbf{\Theta}}, \mathbb{P}_{Y^{\text{in}}})$ represents the weight function shown in Eq. (8).

We conduct extensive experiments to verify the effectiveness of RP and RW strategies. After our modification, the results (red bars in Fig. 2) show a significant improvement, which illustrates the effectiveness of our theory. Moreover, when evaluate the OOD detection performance on iNaturalist dataset, our method can achieve $\sim$30% increase on AUROC and $\sim$51% decrease on FPR95, compared with ODIN (Liu et al., 2020) (see Table 1). It further validates that we cannot default the reference distribution of OOD data to the uniform distribution. To improve the generalizability of OOD detection methods, the class-prior distribution of the training data should be taken into account, which might benefit future researches in the OOD detection community.

2. Preliminaries

Let $\mathcal{X}$ and $\mathcal{Y}^{\text{in}} = \{1, ..., K\}$ be the feature space and ID label space. Let $X^{\text{in}} \in \mathcal{X}$, $X^{\text{out}} \in \mathcal{X}$ and $Y^{\text{in}} \in \mathcal{Y}^{\text{in}}$ be the random variables with respect to $\mathcal{X}$ and $\mathcal{Y}^{\text{in}}$. $\mathbb{P}(X^{\text{in}}, Y^{\text{in}})$ is the ID joint distribution, and $\mathbb{P}(X^{\text{out}}, Y^{\text{out}})$ is the OOD joint distribution, where X^{out} is a random variable from $\mathcal{X}$, but Y^{out} is a random variable whose value do not belong to $\mathcal{Y}^{\text{in}}$. $\mathbb{P}_{X^{\text{in}}}$ is the ID marginal distribution, and $\mathbb{P}_{X^{\text{out}}}$ is the OOD marginal distribution.

OOD Detection (Fang et al., 2022). Given the training data

$$\mathcal{D}_{\text{in}}^{\text{train}} = \{(\mathbf{x}^1, y^1), ..., (\mathbf{x}^n, y^n)\} \sim \mathbb{P}(X^{\text{in}}, Y^{\text{in}}),\ i.i.d.,$$

the aim of OOD detection is to learn a classifier h using $\mathcal{D}_{\text{in}}^{\text{train}}$ such that for any test data $\mathbf{x}$ drawn from $\mathbb{P}_{X^{\text{in}}}$ or $\mathbb{P}_{X^{\text{out}}}$: 1) if $\mathbf{x}$ is drawn from $\mathbb{P}_{X^{\text{in}}}$, then h can classify $\mathbf{x}$ into correct ID classes; and 2) if $\mathbf{x}$ is drawn from $\mathbb{P}_{X^{\text{out}}}$, then h can detect $\mathbf{x}$ as OOD data.

Inference-time OOD Detection. A well-known branch of OOD detection methods is to develop the inference-time OOD detection (or post hoc OOD detection) methods (Huang et al., 2021; Liang et al., 2018; Liu et al., 2020; Hendrycks & Gimpel, 2017; Lee et al., 2018a; Sun et al., 2021), where we are given a pre-trained ID model and then aim to recognize upcoming OOD data well. The key advantage of inference-time OOD detection methods is that the classification performance on ID data will be unaffected since we only use the ID model instead of changing it.

Score Functions. Many representative OOD detection methods (Liu et al., 2020) use a score-based strategy: given a threshold γ, an ID model $\mathbf{f}_{\mathbf{\Theta}}$ and a score function S, then $\mathbf{x}$ is detected as ID data if $S(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) \geq \gamma$:

$$G_\gamma(\mathbf{x}) = \begin{cases} \text{ID}, & \text{if } S(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) \geq \gamma \\ \text{OOD}, & \text{if } S(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) < \gamma \end{cases} \tag{3}$$

The performance of OOD detection depends on how to design a score function S to make OOD data obtain lower scores while ID data obtain higher scores.

3. Methodology

Clearly, without any assumptions or conditions, OOD detection cannot be addressed well due to the unavailability of OOD data (Zhang et al., 2021; Fang et al., 2022). Therefore, to investigate the feasibility of OOD detection, in this section, we consider a natural case that ID classes and OOD features do not interfere with each other, i.e.,

Assumption 3.1. Random variables X^{out} and Y^{in} are independent, i.e., $\mathbb{P}(X^{\text{out}}, Y^{\text{in}}) = \mathbb{P}_{X^{\text{out}}} \mathbb{P}_{Y^{\text{in}}}$.

Justification of Assumption 3.1. To justify that Assumption 3.1 is realistic, we conclude three common causal graphs in Fig. 1. These graphs illustrate how the data is generated through the lens of causality. Notably, in Fig. 1c, we can observe that X^{in} and X^{out} are actually dependent, which is very common in our daily life. It seems that the dependence of X^{in} and X^{out} could result in the failure of Assumption 3.1. However, since X^{in} and X^{out} are dependent only because of the same style (S in Fig. 1) instead of classes (Y in Fig. 1) (Yao et al., 2021), the condition that X^{in} and X^{out} are dependent does not conflict with Assumption 3.1. In fact, there exist many practical scenarios which meet the causal structure in Fig. 1c, e.g., MNIST and Fashion-MNIST (Xiao et al., 2017). According to this assumption, we can prove our main theorem, which provides the theoretical foundation for our paper.

Theorem 3.2 (Corollary of Bayesian Rule)**.** *If Assumption 3.1 holds, $\mathbb{P}_{Y^{\text{in}}|X^{\text{out}}}(y|\mathbf{x}) = \mathbb{P}_{Y^{\text{in}}}(y)$, for any $y \in \mathcal{Y}^{\text{in}}$.*

Based on Theorem 3.2, we have $\operatorname{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})) = \mathbb{P}_{Y^{\text{in}}|X^{\text{out}}}(y|\mathbf{x}) = [\mathbb{P}_{Y^{\text{in}}|X^{\text{out}}}(1|\mathbf{x}), ..., \mathbb{P}_{Y^{\text{in}}|X^{\text{out}}}(K|\mathbf{x})] = [\mathbb{P}_{Y^{\text{in}}}(1), ..., \mathbb{P}_{Y^{\text{in}}}(K)] = \mathbb{P}_{Y^{\text{in}}}$, for any OOD data $\mathbf{x}$. $\operatorname{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})) = \mathbb{P}_{Y^{\text{in}}|X^{\text{out}}}(y|\mathbf{x})$ holds because $\operatorname{softmax}(\cdot)$ maps the model predictions into conditional

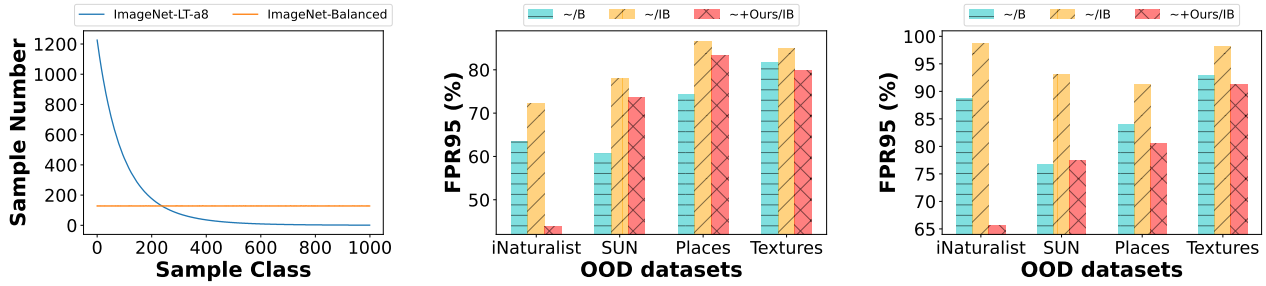

(a) Data Distribution (b) GradNorm (c) Energy

Figure 2: (a) Plot showing the data distribution of balanced and imbalanced datasets. OOD detection performances of (b) GradNorm and (c) Energy. Smaller FPR95 values are better. Cyan (left) bar: the original method on a balanced dataset. Yellow (middle) bar: the original method on an imbalanced dataset. Red (right) bar: the original method with our method on an imbalanced dataset. For a fair comparison, the sample numbers of balanced and imbalanced datasets are the same. More detailed results are shown in Appendix A.1.1.

probability space (see in Bishop & Nasrabadi (2006)). Next, we discuss how to utilize this *novel observation* to improve existing score-based OOD methods. When the labels of the training dataset are not available, we can use the predictions made by the model as an alternative to simulate empirical ID-class-prior distribution $\mathbb{P}_{Y^{\text{in}}}$. More analyses and experiments can be found in the Appendix B.2.

3.1. Rethinking MSP and GradNorm by Theorem 3.2

According to Eq. (3), we discover that the score-based strategy has an *implied assumption* that if a data point $\mathbf{x}$ has a lower score, then the data $\mathbf{x}$ has a higher probability detected as an OOD data point. Based on this assumption, we consider the *ideal* case that if a data point $\mathbf{x}$ has the smallest score, then what will happen? When the score function is MSP, we answer this issue in Theorem 3.3.

Rethinking MSP by Theorem 3.2. We consider the MSP score and answer the above issue by Theorem 3.3.

Theorem 3.3. *Given a data point* $\mathbf{x} \in \mathcal{X}$, *if* $\mathbf{f}^*_{\mathbf{\Theta}}(\mathbf{x}) \in \arg\min_{\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})} S_{\text{MSP}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x})$, *then*

$$\text{softmax}(\mathbf{f}^*_{\mathbf{\Theta}}(\mathbf{x})) = \mathbf{u}, \text{ where } \mathbf{u} = [1/K, \ldots, 1/K] \in \mathbb{R}^K.$$

The proof of Theorem 3.3 is in Appendix B.3. According to the implied assumption, we know that when the data point has the smallest score, then $\mathbf{x}$ has the largest probability detected as an OOD data point. Then, Theorem 3.3 shows that in the ideal case, the output of this data point $\mathbf{x}$ is a uniform distribution $\mathbf{u}$, which is conflict with our observation (i.e., $\text{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})) = \mathbb{P}_{Y^{\text{in}}}$, if $\mathbf{x}$ is an OOD data point) in the ID class-imbalance case. Hence, to avoid the contradiction, we replace the uniform distribution $\mathbf{u}$ in MSP as follows:

$$\begin{aligned} &S_{\text{RP+MSP}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) \\ &= \max_{i \in \{1,\ldots,K\}} (\text{softmax}_i(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})) - \mathbb{P}_{Y^{\text{in}}}(i)). \end{aligned} \tag{4}$$

Rethinking GradNorm by Theorem 3.2. Here we discuss how to adjust the GradNorm score. By Eq. (2), in the ideal case, for any OOD data $\mathbf{x}$, we can conclude that $\text{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})) \approx \mathbf{u}$, i.e.,

$$\lim_{\gamma \to 0} \text{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})) = \mathbf{u}, \tag{5}$$

where $\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})$ satisfies $S_{\text{GradNorm}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) < \gamma$.

Therefore, Eq. (5) is inconsistent with our observation (i.e., $\text{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})) = \mathbb{P}_{Y^{\text{in}}}$, if $\mathbf{x}$ is an OOD data point) in the ID class-imbalance case. Similar to the MSP scenario, the basic idea is to use the ID-class-prior distribution $\mathbb{P}_{Y^{\text{in}}} = [\mathbb{P}_{Y^{\text{in}}}(1), ..., \mathbb{P}_{Y^{\text{in}}}(K)]$ to replace the uniform distribution $\mathbf{u}$, i.e., $S_{\text{RP+GradNorm}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x})$, as

$$\left\| \frac{\partial D_{\text{KL}}(\mathbb{P}_{Y^{\text{in}}} \| \, \text{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})))}{\partial w} \right\|_{L^1}. \tag{6}$$

3.2. Replacing and Reweighting Strategies

Replacing (RP) Strategy. For those methods (e.g., MSP and GradNorm) whose score functions are related to the uniform distribution $\mathbf{u}$, the simple and straight way to modify them is to replace the uniform distribution $\mathbf{u}$ with the ID-class-prior distribution $\mathbb{P}_{Y^{\text{in}}}$. As mentioned in Section 3.1, we modify MSP's score, i.e., $S_{\text{RP+MSP}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x})$ to Eq. (4). and GradNorm's score, i.e., $S_{\text{RP+GradNorm}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x})$ to Eq. (6).

Reweighting (RW) Strategy. For the methods that have no obvious correlations with the uniform distribution $\mathbf{u}$ (e.g., ODIN (Liang et al., 2018) and Energy (Liu et al., 2020)),

we design the RW strategy as a complementary to the RP strategy. RW strategy *reweights* their scores according to a similarity between the ID-class-prior distribution $\mathbb{P}_{Y^{\text{in}}}$ and $\text{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x}))$. Here, we expect that the weights do not impact on the OOD scores seriously. In this paper, we use the cosine function as the weight function, which is one of the most popular distances and similarity functions in contrastive learning (Chen et al., 2020; Grill et al., 2020; He et al., 2020). The main reason we choose cosine function is that cosine is a bounded function and suitable as a weighting parameter after normalization. Specifically, $S_{\text{RW+Method}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x})$ is equal to

$$S_{\text{RW+Method}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) = S_{\text{Method}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) \cdot W(\mathbf{f}_{\mathbf{\Theta}}, \mathbb{P}_{Y^{\text{in}}}), \quad (7)$$

where $W(\mathbf{f}_{\mathbf{\Theta}}, \mathbb{P}_{Y^{\text{in}}})$ is the weight function, i.e.,

$$\begin{aligned} W(\mathbf{f}_{\mathbf{\Theta}}, \mathbb{P}_{Y^{\text{in}}}) =& (1 - \cos(\text{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})), \mathbb{P}_{Y^{\text{in}}})) \\ =& (1 - \frac{\text{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})) \cdot \mathbb{P}_{Y^{\text{in}}}^{\top}}{\|\text{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x}))\| \cdot \|\mathbb{P}_{Y^{\text{in}}}\|}), \end{aligned} \quad (8)$$

where we add a constant 1 to the cosine function to make the weight function be non-negative. $S_{\text{Method}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x})$ is a score function proposed in previous studies, e.g., ODIN and Energy. Next, we introduce the details about the reweighted ODIN and reweighted Energy.

Compared to MSP, the main improvement of ODIN is the use of a temperature scaling strategy. We can modify ODIN as follows:[2] for a temperature $T > 0$, $S_{\text{RW+ODIN}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x})$ is

$$\max_{i \in \{1,\dots,K\}} \frac{\exp(\mathbf{f}_i(\mathbf{x})/T)}{\sum_{j=1}^{K} \exp(\mathbf{f}_j(\mathbf{x})/T)} \cdot W(\mathbf{f}_{\mathbf{\Theta}}, \mathbb{P}_{Y^{\text{in}}}). \quad (9)$$

Energy (Liu et al., 2020) proposes to replace the softmax function with the energy function for OOD detection. The energy function has a property that is highly correlated with the distribution: the system with a more concentrated probability distribution has lower energy, while the system with a more divergent probability distribution (more similar to the uniform distribution) has higher energy (LeCun et al., 2006). Thus, the energy of ID data is smaller than OOD data. Based on Eq. (7), we modify Energy as follows:

$$-T \cdot \log \sum_{i=1}^{K} e^{\mathbf{f}_i(\mathbf{x})/T} \cdot W(\mathbf{f}_{\mathbf{\Theta}}, \mathbb{P}_{Y^{\text{in}}}). \quad (10)$$

In this paper, we mainly realize our strategies using Eq. (4), Eq. (6), Eq. (9) and Eq. (10).

[2] In fact, ODIN uses the modified softmax function with temperature T, which is also related to the uniform distribution, so we can also modify ODIN with RP strategy. We can map the class-prior distribution to the same feature space with ODIN's OOD scores by temperature T. However, if following the default setting ($T = 1000$) in ODIN, $\|\mathbb{P}_{Y^{\text{in}}} - \mathbf{u}\|/T \approx 0$. Thus, RP+ODIN may not work. We will discuss this issue in Appendix B.4.

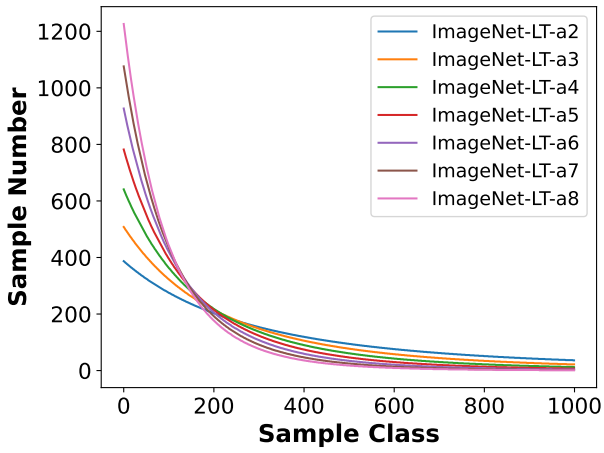

Figure 3: Data distribution with different tail index a.

4. Experiments

In this section, we construct a series of imbalanced ID datasets whose data are sampled from the ImageNet-1K (Deng et al., 2009). Then, we train the ID classifiers on them as pre-trained ID models, and use large-scale ImageNet OOD detection benchmark (Huang & Li, 2021) to evaluate our methods, i.e., RP+MSP (Eq. (4)), RP+GradNorm (Eq. (6)), RW+ODIN (Eq. (9)), and RW+Energy (Eq. (10)).

4.1. Experiment Setup

Training Dataset. Following Liu et al. (2019), we construct a series of imbalanced datasets that are sampled by the Pareto distribution in ImageNet-1K dataset. The definition of Pareto distribution is in Eq. (11).

$$p(x) = \frac{am^a}{x^{a+1}}. \quad (11)$$

In Appendix B.5, we have shown that the parameter m does not affect the level of imbalance. Thus, we set $m = 1$. Additionally, we note that the level of imbalance depends on the tail index a (see Fig. 3), thus, to evaluate the performance of our methods in different imbalanced cases, we take different tail index a. The frequency distributions of classes of the sampled datasets are shown in Fig. 3. As the increase of the tail index a, the sampled datasets become more imbalanced, thus, the ImageNet-LT-a8 dataset is the most imbalanced.

OOD Datasets. In the inference time, we use the large-scale benchmark proposed by Huang & Li (2021). In this benchmark, the OOD datasets include the subsets of iNaturalist (Horn et al., 2018), SUN (Xiao et al., 2010), Places (Zhou et al., 2018), and Textures (Cimpoi et al., 2014). Note that, there are no overlapping classes between ID datasets and OOD datasets (Huang & Li, 2021).

Table 1: OOD detection performance comparison with other competitive score-based OOD detection methods. All methods are based on ResNet101 trained on ImageNet-LT-a8. ↑ indicates larger values are better and ↓ indicates smaller values are better. All values are percentages. The bold indicates the best performance while the underline indicates the second.

Method	iNaturalist		SUN		Places		Textures		Average	
	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
MSP (Hendrycks & Gimpel, 2017)	63.95	97.72	66.60	93.13	66.84	92.11	42.74	98.79	60.03	95.44
ODIN (Liang et al., 2018)	60.14	98.70	70.63	93.13	70.14	91.96	41.83	98.30	60.69	95.52
Mahalanobis (Lee et al., 2018a)	60.72	95.87	56.79	94.50	55.27	93.78	49.43	86.99	55.55	92.78
Energy (Liu et al., 2020)	55.99	98.74	71.12	93.11	70.24	91.30	42.38	98.07	59.93	95.30
GradNorm (Huang et al., 2021)	82.51	72.19	74.57	78.10	70.67	86.58	57.31	84.95	71.26	80.45
Dice (Sun & Li, 2022)	85.80	58.96	73.17	76.90	67.83	87.89	58.43	80.59	71.31	76.11
RP+MSP (Ours)	64.95	96.44	67.39	91.79	67.46	91.16	43.05	98.51	60.71	94.48
RW+ODIN (Ours)	90.24	47.52	**81.33**	75.44	**77.90**	**79.90**	61.43	88.03	**77.73**	72.72
RW+Energy (Ours)	83.12	65.67	80.13	77.51	77.12	80.49	51.70	91.31	73.02	78.75
RP+GradNorm (Ours)	**91.23**	**43.87**	77.36	**73.53**	72.67	83.29	**62.94**	**79.80**	76.05	**70.12**

Evaluation Metrics. We use two common metrics to evaluate OOD detection methods (Huang et al., 2021): the false positive rate that OOD data are classified as ID data when 95% of ID data are correctly classified (FPR95) (Provost et al., 1998) and the *area under the receiver operating characteristic curve* (AUROC) (Huang et al., 2021).

Baselines. In order to verify the effectiveness of our strategies, we select MSP (Hendrycks & Gimpel, 2017), ODIN (Liang et al., 2018), Energy (Liu et al., 2020), GradNorm (Huang et al., 2021) and Dice (Sun & Li, 2022) as the baselines. Following the setting of (Huang et al., 2021; Liang et al., 2018), the temperature parameter T in ODIN is set to be 1000 and in GradNorm is 1. The norm function in GradNorm is set to be L1 norm.

Models and Hyperparameters. We use *mmclassification*[3] (Contributors, 2020) with Apache-2.0 license to train ID models. The training details of ResNet (He et al., 2016) and MobileNet (Howard et al., 2019) follow the default setting in *mmclassification*. Note that, all methods are realized by Pytorch 1.6.0 with CUDA 10.2, where we use several NVIDIA Tesla V100 GPUs.

4.2. Experimental Results and Analysis

Verification of Two Strategies. Our strategies are applicable to various score functions used by OOD detection methods. The performance of our methods and baselines are shown in Table 1. Overall, after modifying previous methods using our strategies, their performance are significantly improved, indicating the effectiveness of our strategies. Specifically, RW+ODIN achieves the highest AUROC and RP+GradNorm achieves lowest FPR95 compared to all methods. As a highlight, RW+ODIN shows the most significant performance improvement on all four datasets: ∼51% FPR95 decrease in iNaturalist, ∼18% FPR95 decrease in SUN, ∼12% FPR95 decrease in Places and ∼10% FPR95 decrease in Textures. Experimental results have shown that our strategies can outperform the baselines in the ID-class-imbalanced scenarios.

Analysis of Detection Results on Different ID Classes. Since the training dataset is imbalanced, we follow Liu et al. (2019) to divide all ID classes into three categories (ID Head, ID Mid, ID Tail) for further analysis. In detail, ID Head category includes the classes containing more than 100 samples in the training dataset;ID Mid category includes the classes whose number of samples is between 20 and 100 in the training dataset; and ID Tail category includes the classes containing less than 20 samples in the training dataset. Then, we evaluate the OOD detection performance on three datasets: ID Head+OOD, ID Mid+OOD and ID Tail+OOD (details can be found in Appendix A.1.3). If the performance of one method on ID Tail+OOD is better than that on ID Head+OOD, then this method performs better when facing tailed classes and OOD data.

In the case of GradNorm, experiment results in Fig. 4 show that our method RP+GradNorm improves the performance on the above three datasets (ID Head+OOD, ID Mid+OOD, and ID Tail+OOD). When we take a closer look at the performance improvement, we notice that the overall performance improvement of RP+GradNorm is mainly due to the significant improvement on the Tail+OOD dataset. This result might indicate that the previous method, like GradNorm, confuses OOD data and ID tailed classes, which hinders their OOD detection performance. The results show that our strategies can overcome this issue. More detailed results are shown in Appendix A.1.3.

4.3. Ablation Study

Analysis regarding Tail Index a. Here, we report the performance of our method and baselines when changing the

[3] https://github.com/open-mmlab/mmclassification

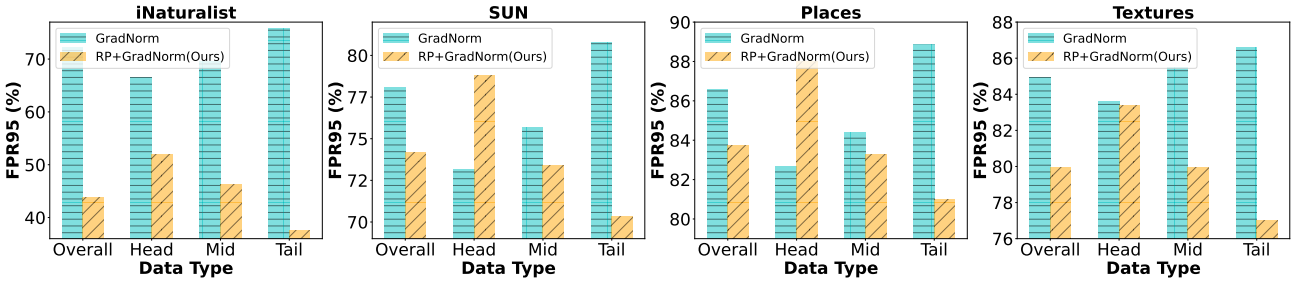

Figure 4: Performance comparison of different data type. The figures shows the OOD detection performance of GradNorm and RP+GradNorm in four OOD datasets.

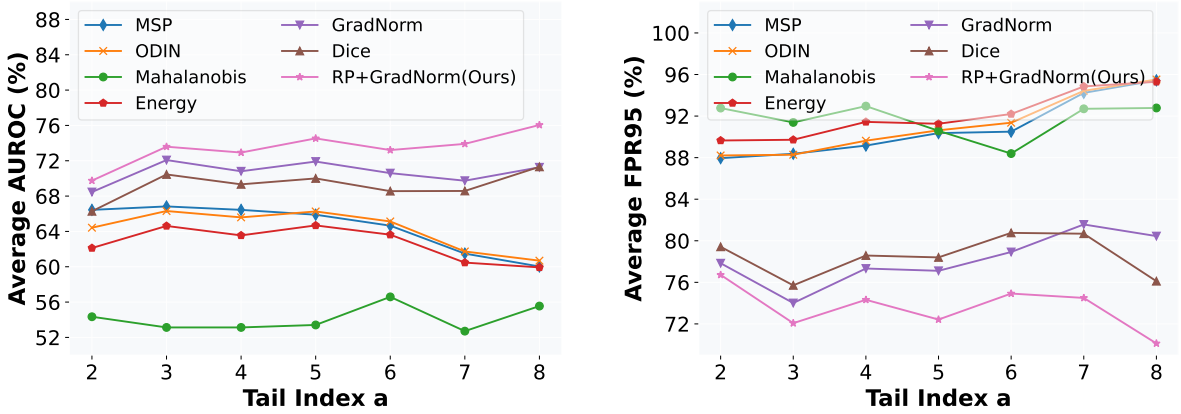

(a) AUROC↑ (b) FPR95↓

Figure 5: OOD detection performance with ResNet101 trained on different imbalanced ID datasets.

tail index $a \in \{2, 3, \ldots, 8\}$. We conduct repeated experiments on these seven datasets (ImageNet-LT-a2, ImageNet-LT-a3,. . . , ImageNet-LT-a8), and the results are shown in Fig. 5. Overall, our method RP+GradNorm always outperforms other baselines with different imbalance degrees. More importantly, the performance improvement between RP+GradNorm and each baseline gradually increases, as the increase of the imbalance degree. This indicates that RP+GradNorm can handle different imbalanced scenarios better. More detailed results are in Appendix A.1.2.

Analysis regarding Network Architecture. We evaluate all methods on a different network architecture, MobileNet-V3 (Howard et al., 2019). Experiment results in Table 2 show that our methods (RP+GradNorm and RW+ODIN) still outperform baselines even when we change the network architecture. In addition, RP+GradNorm has a better performance on FPR95 while RW+ODIN has higher AUROC, corresponding to the performances of GradNorm and ODIN.

Analysis regarding Model Size. We provide an experiment about the model size of RP+GradNorm. We compare ResNet50, ResNet101 and ResNet152 trained on ImageNet-LT-a8 datasets. The results are shown in Table 3. The optimal model is the smallest one (ResNet50), and we observe that as the increase of the model size, the performance decreases. One possible reason is that small models are more difficult to overfit and thus more suitable for OOD detection in imbalanced scenarios.

More Experiments and Exploration. First, we can regard the cosine similarity weights in the RW strategy as a score function, and conduct several experiments in Appendix A.1.4. We notice that the cosine similarity also achieves an improvement compared with baselines but the cosine similarity is sensitive to the ID data distribution.

Second, to evaluate the stability of our strategies, we conduct 10 independent replicate experiments in Appendix A.1.5. The results show that our proposed method RP+GradNorm is stable under differnet trails with randomly picked tail classes.

Third, to further explore the combination ways between existing methods and the proposed strategies, we conducted a series of experiments involving RW+MSP, RW+GradNorm,

Table 2: OOD detection performance with MobileNet trained on ImageNet-LT-a8. ↑ indicates larger values are better and ↓ indicates smaller values are better. All values are percentages. The bold indicates the best performance while the underline indicates the second.

Method	iNaturalist		SUN		Places		Textures		Average	
	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
MSP (Hendrycks & Gimpel, 2017)	63.47	92.38	67.27	85.62	64.16	89.62	59.88	89.06	63.69	89.17
ODIN (Liang et al., 2018)	64.68	93.78	74.29	79.42	69.94	89.70	69.06	82.23	69.49	86.28
Energy (Liu et al., 2020)	63.42	96.37	74.95	77.86	70.30	90.50	70.69	80.83	69.84	86.39
GradNorm (Huang et al., 2021)	70.87	78.12	69.70	67.59	66.00	85.75	63.09	74.89	67.41	76.59
Dice (Sun & Li, 2022)	65.61	86.40	69.35	66.38	65.95	88.42	68.85	68.19	67.44	77.35
RP+MSP (Ours)	63.76	91.76	67.56	85.11	64.37	89.41	60.02	88.62	63.93	88.73
RW+ODIN (Ours)	**82.51**	**66.06**	**80.08**	69.12	**74.33**	**79.41**	69.58	78.07	**76.63**	73.16
RW+Energy (Ours)	73.70	80.28	79.32	69.73	74.31	82.30	**72.24**	76.01	74.89	77.08
RP+GradNorm (Ours)	77.25	68.61	72.49	**66.02**	68.56	82.30	64.69	**71.86**	70.75	**72.20**

Table 3: OOD detection performance with model size increases. The RP+GradNorm method is trained on ImageNet-LT-a8. All values are percentages. ↑ indicates larger values are better and ↓ indicates smaller values are better.

Model	iNaturalist		SUN		Places		Textures		Average	
	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
ResNet50	89.85	50.03	**80.73**	**64.52**	**74.69**	**78.18**	**63.31**	**77.73**	**77.14**	**67.62**
ResNet101	**91.23**	**43.87**	77.36	73.53	72.67	83.29	62.94	79.80	76.05	70.12
ResNet152	88.24	53.14	73.45	78.96	68.61	87.14	59.41	83.71	72.43	75.74

and RW+RP+MSP/GradNorm and the details can be found in Appendix A.1.6. The results illustrate that MSP and RW are the best match while GradNorm matches RP well, in line with the scenario where RW and RP are suitable for implicit and explicit distribution respectively.

Besides, we also conduct more experiments on various benchmarks and settings to widely investigate the performances. Appendix A.1.7 shows that our methods still work well when models are trained with long-tailed learning strategies (Cao et al., 2019; Park et al., 2022) during the training phase. The results of detecting OOD samples under balanced training are reported in Appendix A.1.8. We explore more distance metrics for data distribution and report the results in Appendix A.1.9. We evaluate our strategies on near OOD benchmark (Fort et al., 2021) in Appendix A.2.

5. Related works

OOD Detection. OOD detection is a crucial problem for reliably deploying machine learning models into real-world scenarios. OOD detection can be divided into two categories according to whether the classifier will be re-trained for OOD detection or not.

1) Inference-time/post hoc OOD Detection: Some methods (Huang et al., 2021; Liang et al., 2018; Liu et al., 2020; Hendrycks & Gimpel, 2017; Lee et al., 2018a; Sun et al., 2021) focus on designing OOD score functions for OOD detection in the inference time and are easy to use without changing the model's parameters. This property is important for deploying OOD detection methods in real-world scenarios where the cost of re-training is prohibitively expensive and time-consuming. MSP (Hendrycks & Gimpel, 2017) directly takes the maximum value of the model's prediction as the OOD score function. Based on MSP, ODIN (Liang et al., 2018) uses a temperature scaling strategy and input perturbation to improve OOD detection performance. Moreover, (Liu et al., 2020) and (Wang et al., 2021a) propose to replace the softmax function with the energy functions for OOD detection. Recently, GradNorm (Huang et al., 2021) uses the similarity of the model-predicted probability distribution and the uniform distribution to improve OOD detection and achieve state-of-the-art performance. In this paper, we mainly work on the inference-time OOD detection methods and aim at improving the generalizability of OOD detection in real-world scenarios.

2) Training-time OOD detection: Other methods (Hsu et al., 2020; Hein et al., 2019; Bitterwolf et al., 2020; Wang et al., 2021b; Zhu et al., 2023) will complete ID tasks and OOD detection simultaneously in the training time. (Bitterwolf et al., 2020) uses adversarial learning to process OOD data in training time and make the model predict lower confidence scores for them. (Wang et al., 2021b) generates pseudo OOD data by adversarial learning to re-training a K+1 model for OOD detection. These methods usually require auxiliary OOD data available in the training process.

Thus, the model will be affected by both ID data and OOD data. It is important for these method to explore an inherent trade-off (Liu et al., 2019; Vaze et al., 2022; Yang et al., 2021) between ID tasks and OOD detection.

Wang et al. (2022) is training-time OOD detection and uses OOD data to train the model. After being finetuned, the model can deal with the imbalanced issue and OOD problem. The problem is similar to our paper, but the setting is completely different with our paper (inference-time OOD detection). In our paper, we do not change any parameters of the model and design methods to deal with the imbalanced issue on OOD detection. Note that our work and this work are not comparable due to the different problem settings.

Open Set Recognition. In open set recognition, machine learning models (Huang & Li, 2021; Lee et al., 2018b; Perera & Patel, 2019; Perera et al., 2020; Shalev et al., 2018; Radford et al., 2021; Fort et al., 2021) are required to both correctly classify the known data (ID) from the closed set and detect unknown data (OOD) from the open set. Some works (Lee et al., 2018b; Huang & Li, 2021) use the information in the label space for OOD detection, and they divide the large semantic space into multiple levels for models to easily understand. (Perera & Patel, 2019) designs two parallel networks training on different dataset and use the membership loss to encourage high activations for ID data while reducing activations for OOD data. (Perera et al., 2020) uses self-supervision and data augmentation to improve the network's ability to detect OOD data. Input images are augmented with the representation obtained from a generative model. In this paper, we consider a more complex open set, large scale and imbalanced, to achieve OOD detetcion.

6. Conclusion

This paper theoretically and empirically shows that the unproven assumption of uniform distribution in previous methods is not valid when the training dataset is imbalanced. Moreover, by analyzing the causal relations between ID/OOD classes and features, we point out that the best reference distribution for OOD data is the ID-class-prior distribution. Based on this, we propose two simple and effective strategies to modify the uniform distribution assumption in previous inference-time OOD detection methods. RP strategy is suitable for the methods that directly use the uniform distribution to design the OOD score function, while RW strategy is designed for methods that potentially use the assumption. Extensive experiments show that both strategies can improve the performance of OOD detection on large-scale image classification benchmarks.

7. Acknowledgement

This work was supported by the National Key R&D Program of China (Grant NO. 2022YFF1202903) and the National Natural Science Foundation of China (Grant NO. 62122035, 61972188 and 12071166). BH was supported by NSFC Young Scientists Fund No. 62006202, and Guangdong Basic and Applied Basic Research Foundation No. 2022A1515011652.

References

Bishop, C. M. and Nasrabadi, N. M. *Pattern recognition and machine learning*, volume 4. Springer, 2006.

Bitterwolf, J., Meinke, A., and Hein, M. Certifiably adversarially robust detection of out-of-distribution data. In *NeurIPS*, 2020.

Cao, K., Wei, C., Gaidon, A., Arechiga, N., and Ma, T. Learning imbalanced datasets with label-distribution-aware margin loss. In *NeurIPS*, 2019.

Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. E. A simple framework for contrastive learning of visual representations. In *ICML*, 2020.

Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., and Vedaldi, A. Describing textures in the wild. In *CVPR*, 2014.

Contributors, M. Openmmlab's image classification toolbox and benchmark. https://github.com/open-mmlab/mmclassification, 2020.

Deng, J., Dong, W., Socher, R., Li, L., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.

Fang, Z., Li, Y., Lu, J., Dong, J., Han, B., and Liu, F. Is out-of-distribution detection learnable? In *NeurIPS*, 2022.

Fort, S., Ren, J., and Lakshminarayanan, B. Exploring the limits of out-of-distribution detection. In *NeurIPS*, 2021.

Gong, M., Zhang, K., Liu, T., Tao, D., Glymour, C., and Schölkopf, B. Domain adaptation with conditional transferable components. In *ICML*, 2016.

Grill, J., Strub, F., Altché, F., Tallec, C., Richemond, P. H., Buchatskaya, E., Doersch, C., Pires, B. Á., Guo, Z., Azar, M. G., Piot, B., Kavukcuoglu, K., Munos, R., and Valko, M. Bootstrap your own latent - A new approach to self-supervised learning. In *NeurIPS*, 2020.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *CVPR*, 2016.

He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020.

Hein, M., Andriushchenko, M., and Bitterwolf, J. Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In *CVPR*, 2019.

Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *ICLR*, 2017.

Horn, G. V., Aodha, O. M., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., and Belongie, S. J. The inaturalist species classification and detection dataset. In *CVPR*, 2018.

Howard, A., Sandler, M., Chu, G., Chen, L.-C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., Le, Q. V., and Adam, H. Searching for mobilenetv3. In *ICCV*, 2019.

Hsu, Y., Shen, Y., Jin, H., and Kira, Z. Generalized ODIN: detecting out-of-distribution image without learning from out-of-distribution data. In *CVPR*, 2020.

Huang, R. and Li, Y. MOS: towards scaling out-of-distribution detection for large semantic space. In *CVPR*, 2021.

Huang, R., Geng, A., and Li, Y. On the importance of gradients for detecting distributional shifts in the wild. In *NeurIPS*, 2021.

LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., and Huang, F. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.

Lee, K., Lee, K., Lee, H., and Shin, J. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *NeurIPS*, 2018a.

Lee, K., Lee, K., Min, K., Zhang, Y., Shin, J., and Lee, H. Hierarchical novelty detection for visual object recognition. In *CVPR*, 2018b.

Liang, S., Li, Y., and Srikant, R. Enhancing the reliability of out-of-distribution image detection in neural networks. In *ICLR*, 2018.

Liu, W., Wang, X., Owens, J. D., and Li, Y. Energy-based out-of-distribution detection. In *NeurIPS*, 2020.

Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., and Yu, S. X. Large-scale long-tailed recognition in an open world. In *CVPR*, 2019.

Nguyen, A. M., Yosinski, J., and Clune, J. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *CVPR*, 2015.

Park, S., Hong, Y., Heo, B., Yun, S., and Choi, J. Y. The majority can help the minority: Context-rich minority oversampling for long-tailed classification. In *CVPR*, 2022.

Perera, P. and Patel, V. M. Deep transfer learning for multiple class novelty detection. In *CVPR*, 2019.

Perera, P., Morariu, V. I., Jain, R., Manjunatha, V., Wigington, C., Ordonez, V., and Patel, V. M. Generative-discriminative feature representations for open-set recognition. In *CVPR*, 2020.

Provost, F. J., Fawcett, T., and Kohavi, R. The case against accuracy estimation for comparing induction algorithms. In *ICML*, 1998.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. Learning transferable visual models from natural language supervision. In *ICML*, 2021.

Shalev, G., Adi, Y., and Keshet, J. Out-of-distribution detection using multiple semantic label representations. In *NeurIPS*, 2018.

Stojanov, P., Li, Z., Gong, M., Cai, R., Carbonell, J. G., and Zhang, K. Domain adaptation with invariant representation learning: What transformations to learn? In *NeurIPS*, 2021.

Sun, Y. and Li, Y. Dice: Leveraging sparsification for out-of-distribution detection. In *ECCV*, 2022.

Sun, Y., Guo, C., and Li, Y. React: Out-of-distribution detection with rectified activations. In *NeurIPS*, 2021.

Sun, Y., Ming, Y., Zhu, X., and Li, Y. Out-of-distribution detection with deep nearest neighbors. *ICML*, 2022.

Vaze, S., Han, K., Vedaldi, A., and Zisserman, A. Open-set recognition: A good closed-set classifier is all you need. In *ICLR*, 2022.

Wang, H., Ge, S., Lipton, Z. C., and Xing, E. P. Learning robust global representations by penalizing local predictive power. In *NeurIPS*, 2019.

Wang, H., Liu, W., Bocchieri, A., and Li, Y. Can multi-label classification networks know what they don't know? In *NeurIPS*, 2021a.

Wang, H., Zhang, A., Zhu, Y., Zheng, S., Li, M., Smola, A. J., and Wang, Z. Partial and asymmetric contrastive learning for out-of-distribution detection in long-tailed recognition. In *ICML*, 2022.

Wang, Y., Li, B., Che, T., Zhou, K., Liu, Z., and Li, D. Energy-based open-world uncertainty modeling for confidence calibration. In *ICCV*, 2021b.

Xiao, H., Rasul, K., and Vollgraf, R. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms, 2017.

Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., and Torralba, A. SUN database: Large-scale scene recognition from abbey to zoo. In *CVPR*, 2010.

Yang, J., Wang, H., Feng, L., Yan, X., Zheng, H., Zhang, W., and Liu, Z. Semantically coherent out-of-distribution detection. In *ICCV*, 2021.

Yao, Y., Liu, T., Gong, M., Han, B., Niu, G., and Zhang, K. Instance-dependent label-noise learning under a structural causal model. In *NeurIPS*, 2021.

Zhang, L. H., Goldstein, M., and Ranganath, R. Understanding failures in out-of-distribution detection with deep generative models. In *ICML*, 2021.

Zhou, B., Lapedriza, À., Khosla, A., Oliva, A., and Torralba, A. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464, 2018.

Zhu, J., Li, H., Yao, J., Liu, T., Xu, J., and Han, B. Unleashing mask: Explore the intrinsic out-of-distribution detection capability. In *ICML*, 2023.

A. Further Experiments

A.1. Evaluation on ImageNet Benchmark

A.1.1. Evaluation on Imbalanced Data and Balanced Data

We randomly sample a balanced dataset from ImageNet-1K dataset, which has the same sample numbers with the imbalanced datasets. We conduct experiments on the balanced data as shown in Fig. 6 and Fig. 7. All methods shows a similar trend, i.e., the performance drop a lot when the training dataset becomes imbalanced (from cyan bars to yellow bars). Moreover, our method shows a significant improvement with previous methods on all evaluation tasks (from yellow bars to red bars).

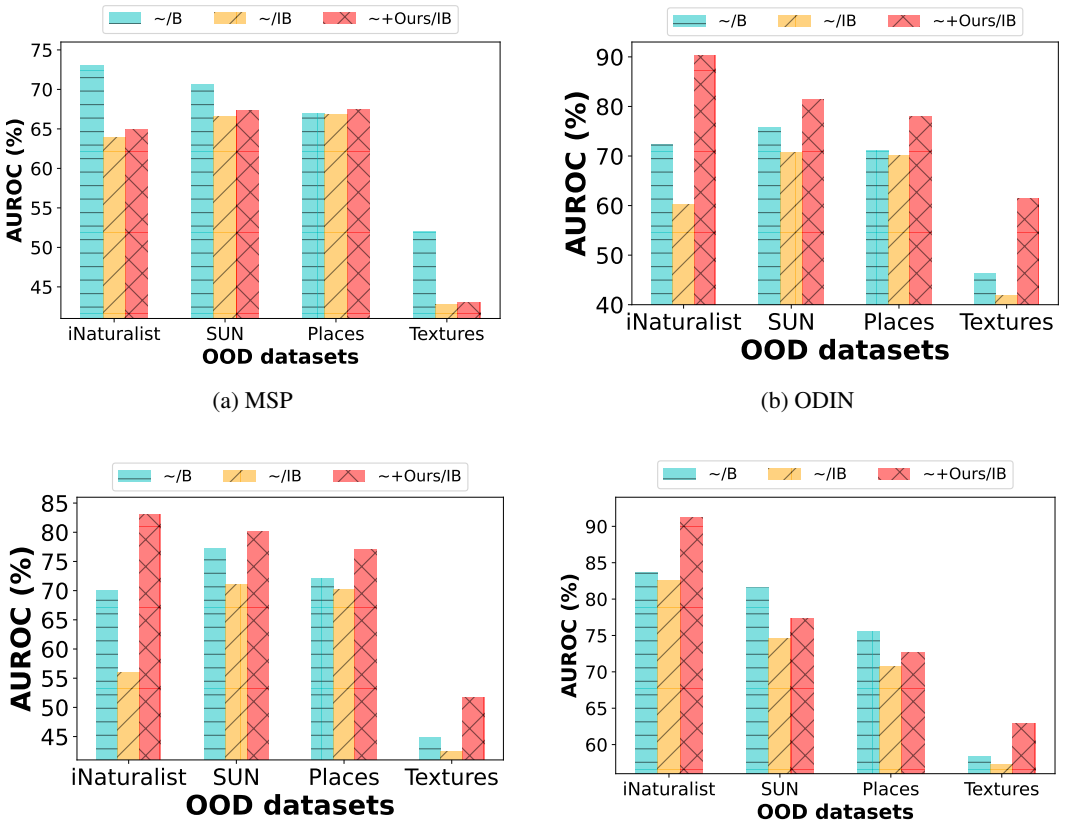

(a) MSP (b) ODIN

(c) Energy (d) GradNorm

Figure 6: OOD detection performance (AUROC) of (a) MSP (b) ODIN (c) Energy and (d) GradNorm. Larger AUROC values are better. Cyan (left) bar: the original method on balanced dataset. Yellow (middle) bar: the original method on imbalanced dataset. Red (right) bar: the original method with our method on imbalanced dataset.

A.1.2. Analysis regarding Tail Index

We can obtain sampled datasets with different levels of imbalance based on the Pareto distribution. ImageNet-LT-a8 dataset is the most imbalanced, while ImageNet-LT-a2 is the most balanced. We conduct repeated experiments on these seven datasets and results are shown in Table 4. Obviously, the more imbalanced the training ID dataset becomes, the more our methods (RP+GradNorm and Cosine Similarity) demonstrates their superior performance of OOD detection, compared to other methods on all evaluation tasks.

It is noticeable that the detection performance of GradNorm is relatively stable no matter how imbalanced the ratio changes, compared with other existing methods (such as MSP, ODIN, Energy). These methods explicitly/implicitly use the discrepancy between the classifier's output and the uniform distribution. Thus, they will be affected a lot if the prior distribution changes from the uniform distribution to an imbalanced/tailed one.

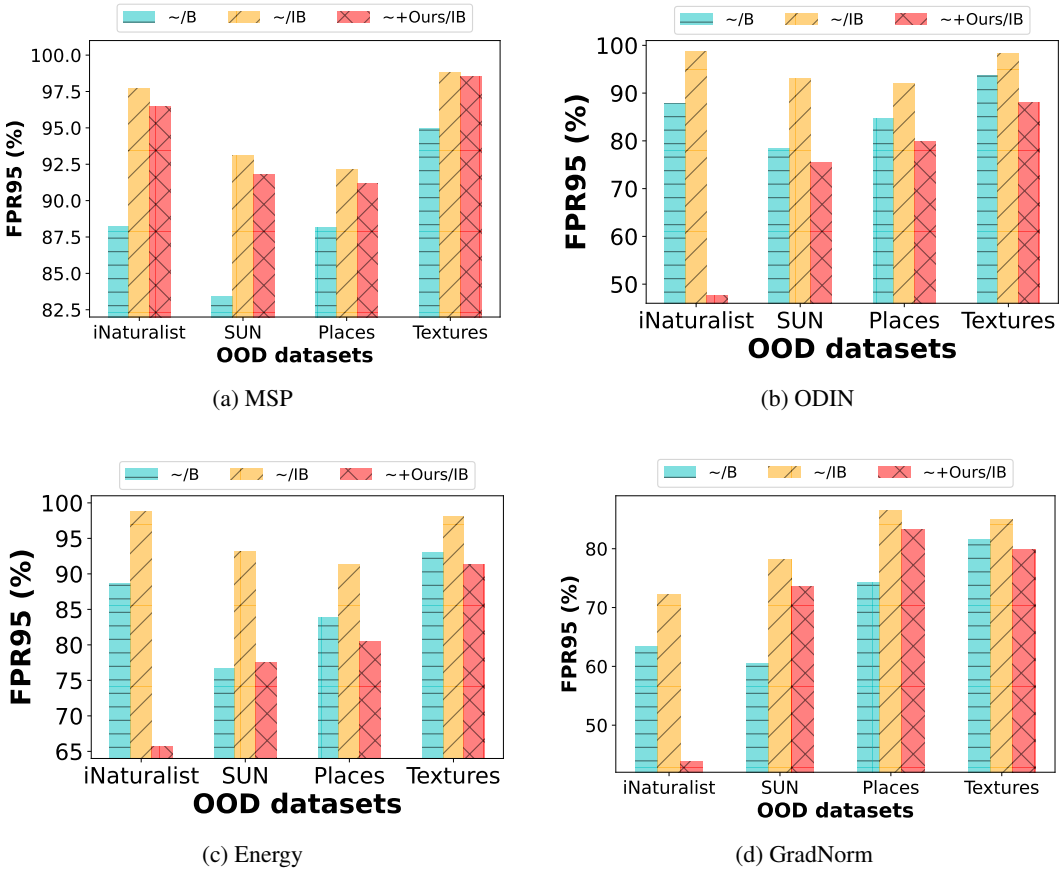

(a) MSP (b) ODIN

(c) Energy (d) GradNorm

Figure 7: OOD detection performance (FPR95) of (a) MSP (b) ODIN (c) Energy and (d) GradNorm. Smaller FPR95 values are better. Cyan (left) bar: the original method on balanced dataset. Yellow (middle) bar: the original method on imbalanced dataset. Red (right) bar: the original method with our method on imbalanced dataset.

As for GradNorm, we conjecture that considering the gradient space might be robust to the changes of priors (such as from a uniform prior to an imbalanced prior). To verify this conjecture, we conduct the experiment that KL divergence is directly used to measure the discrepancy between the output of the classifier and the uniform distribution (i.e., GradNorm without gradient-norm process). The results are shown in Table 5. Obviously, this KL-based method is also significantly affected by the imbalanced situation, then we can verify this conjecture. Thus, we confirm that GradNorm's robustness of the imbalanced ratio depends on the gradient space.

A.1.3. Analysis of Detection Results on Different ID Classes

We calculate the evaluation metrics for the three categories (ID-Head, ID-Mid, ID-Tail) by randomly sampling OOD data in equal proportions corresponding to the number of samples in each category. For example, we have 50000 ID samples and 10000 OOD samples in total. If the number of samples in category Head is 10000 and accordingly we will sample 2000 OOD samples, then we use the 12000 samples to calculate AUROC and FPR95. The results reflects the confusion degree between ID Head data and OOD data in the view of OOD detection methods. We conduct experiments to analyze the performance of different data types, as shown in Table 6.

Moreover, we visualize the OOD score distributions in Fig. 8-Fig. 11. Obviously, the results and figures show that the previous methods tend to confuse OOD data and the minority classes, which hinders their performance of OOD detection. And our strategies can reduce the confusion to improve OOD detection performance.

Table 4: OOD detection performance with ResNet101 trained on different imbalanced ID datasets. ↑ indicates larger values are better and ↓ indicates smaller values are better. All values are percentages.

Datasets	Method	iNaturalist		SUN		Places		Textures		Average	
		AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
ImageNet-LT-a2	MSP (Hendrycks & Gimpel, 2017)	71.57	89.71	73.74	80.64	70.76	85.17	49.68	96.28	66.44	87.95
	ODIN (Liang et al., 2018)	65.92	93.17	**76.32**	79.60	**72.77**	84.29	42.67	95.76	64.42	88.21
	Mahalanobis (Lee et al., 2018a)	61.12	92.51	52.79	95.47	52.26	94.96	51.19	88.14	54.34	92.77
	Energy (Liu et al., 2020)	60.26	95.63	75.57	81.89	72.00	85.90	40.67	95.18	62.12	89.65
	GradNorm (Huang et al., 2021)	75.82	75.86	75.33	**68.03**	70.13	**80.61**	52.58	86.86	68.46	77.84
	Dice (Sun & Li, 2022)	72.17	77.31	73.94	70.28	68.88	83.69	50.12	**86.38**	66.28	79.42
	RP+GradNorm (Ours)	79.66	70.18	75.70	68.50	70.33	81.34	**53.30**	86.88	**69.75**	**76.72**
	Cosine Similarity (Ours)	**81.84**	**67.06**	73.26	74.27	68.78	81.75	47.94	92.18	67.95	78.82
ImageNet-LT-a3	MSP (Hendrycks & Gimpel, 2017)	73.97	89.10	73.38	80.91	69.91	87.20	50.07	96.28	66.83	88.37
	ODIN (Liang et al., 2018)	72.04	92.06	75.92	79.71	**71.49**	86.30	45.79	95.04	66.31	88.28
	Mahalanobis (Lee et al., 2018a)	48.51	94.60	53.81	92.63	53.85	91.84	56.36	86.48	53.13	91.39
	Energy (Liu et al., 2020)	67.94	94.42	75.24	82.36	70.48	88.07	44.81	94.02	64.62	89.72
	GradNorm (Huang et al., 2021)	83.41	65.27	76.93	**67.10**	70.94	**81.43**	57.03	**82.25**	72.08	74.01
	Dice (Sun & Li, 2022)	81.72	65.63	75.60	70.30	69.81	84.44	54.68	82.48	70.45	75.71
	RP+GradNorm (Ours)	87.13	56.16	**77.47**	67.71	71.42	81.96	**58.35**	82.45	**73.59**	**72.07**
	Cosine Similarity (Ours)	**87.81**	**50.66**	74.15	73.80	68.87	82.92	52.98	87.73	70.96	72.78
ImageNet-LT-a4	MSP (Hendrycks & Gimpel, 2017)	72.41	89.92	73.00	83.34	70.38	86.97	49.96	96.42	66.44	89.16
	ODIN (Liang et al., 2018)	70.35	92.76	74.66	82.95	**71.35**	87.30	46.00	95.50	65.59	89.63
	Mahalanobis (Lee et al., 2018a)	49.37	96.48	56.58	93.36	54.98	93.75	51.59	85.45	53.13	92.26
	Energy (Liu et al., 2020)	66.27	95.05	73.17	85.85	69.63	89.71	45.13	95.14	63.55	91.44
	GradNorm (Huang et al., 2021)	80.87	71.75	74.64	**71.60**	69.86	84.24	57.87	81.72	70.81	77.33
	Dice (Sun & Li, 2022)	78.60	72.35	74.02	73.34	68.94	86.45	55.75	82.18	69.33	78.58
	RP+GradNorm (Ours)	86.20	60.40	**75.29**	71.93	70.40	84.43	**59.85**	**80.51**	**72.93**	74.32
	Cosine Similarity (Ours)	**89.21**	**48.91**	74.13	74.53	70.06	**82.41**	57.25	86.60	72.66	**73.11**
ImageNet-LT-a5	MSP (Hendrycks & Gimpel, 2017)	72.18	91.15	72.43	85.42	70.38	87.73	48.59	97.13	65.89	90.36
	ODIN (Liang et al., 2018)	69.53	94.68	75.92	84.34	73.07	87.04	46.49	96.47	66.26	90.63
	Mahalanobis (Lee et al., 2018a)	48.35	96.03	58.07	91.81	57.20	90.52	50.00	83.95	53.41	90.58
	Energy (Liu et al., 2020)	64.85	95.98	75.38	85.29	72.25	87.89	46.25	95.89	64.68	91.26
	GradNorm (Huang et al., 2021)	84.01	67.28	75.30	73.25	69.82	85.38	58.53	82.52	71.91	77.11
	Dice (Sun & Li, 2022)	82.88	65.51	73.14	76.80	66.94	89.04	57.04	82.25	70.00	78.40
	RP+GradNorm (Ours)	89.59	51.01	76.55	72.93	70.91	84.68	**61.06**	**81.05**	74.53	72.42
	Cosine Similarity (Ours)	**91.53**	**41.44**	**78.57**	**71.83**	**74.41**	**79.89**	60.13	86.74	**76.16**	**69.97**
ImageNet-LT-a6	MSP (Hendrycks & Gimpel, 2017)	70.99	91.04	71.63	86.67	70.50	86.62	45.49	97.66	64.65	90.50
	ODIN (Liang et al., 2018)	70.82	92.57	74.04	87.81	72.18	88.25	43.49	96.83	65.13	91.36
	Mahalanobis (Lee et al., 2018a)	57.62	89.82	60.18	89.38	59.07	89.46	49.49	84.92	56.59	88.39
	Energy (Liu et al., 2020)	67.93	93.88	72.62	89.16	70.44	89.81	43.53	95.94	63.63	92.20
	GradNorm (Huang et al., 2021)	83.49	67.76	74.02	**76.58**	69.12	85.87	55.75	85.46	70.60	78.92
	Dice (Sun & Li, 2022)	82.31	68.27	71.70	79.35	66.20	89.82	54.04	85.60	68.56	80.76
	RP+GradNorm (Ours)	89.07	51.55	75.09	77.47	69.89	86.28	58.80	**84.38**	73.21	74.92
	Cosine Similarity (Ours)	**91.22**	**43.01**	**78.35**	78.35	**74.64**	**82.94**	**60.08**	88.63	**76.07**	**73.23**
ImageNet-LT-a7	MSP (Hendrycks & Gimpel, 2017)	65.60	96.81	68.50	90.68	67.22	91.26	44.65	98.16	61.49	94.23
	ODIN (Liang et al., 2018)	63.13	97.72	71.29	91.17	69.16	91.20	43.33	97.64	61.73	94.43
	Mahalanobis (Lee et al., 2018a)	50.52	97.86	54.18	92.36	54.19	92.31	51.96	88.28	52.71	92.70
	Energy (Liu et al., 2020)	59.46	98.20	70.56	91.89	68.07	91.90	43.77	97.41	60.47	94.85
	GradNorm (Huang et al., 2021)	80.38	77.04	73.07	77.43	66.86	88.60	58.65	83.21	69.74	81.57
	Dice (Sun & Li, 2022)	81.61	70.42	71.19	78.93	63.79	91.40	57.72	81.95	68.58	80.68
	RP+GradNorm (Ours)	88.59	55.40	75.17	**75.35**	68.78	86.41	**63.08**	**80.83**	73.91	74.50
	Cosine Similarity (Ours)	**90.12**	**49.49**	**78.62**	77.87	**73.94**	**82.90**	62.43	87.54	**76.28**	**74.45**
ImageNet-LT-a8	MSP (Hendrycks & Gimpel, 2017)	63.95	97.72	66.60	93.13	66.84	92.11	42.74	98.79	60.03	95.44
	ODIN (Liang et al., 2018)	60.14	98.70	70.63	93.13	70.14	91.96	41.83	98.30	60.69	95.52
	Mahalanobis (Lee et al., 2018a)	60.72	95.87	56.79	94.50	55.27	93.78	49.43	86.99	55.55	92.78
	Energy (Liu et al., 2020)	55.99	98.74	71.12	93.11	70.24	91.30	42.38	98.07	59.93	95.30
	GradNorm (Huang et al., 2021)	82.51	72.19	74.57	78.10	70.67	86.58	57.31	84.95	71.26	80.45
	Dice (Sun & Li, 2022)	85.80	58.96	73.17	76.90	67.83	87.89	58.43	80.69	71.31	76.11
	RP+GradNorm (Ours)	**91.23**	**43.87**	77.36	73.53	72.67	83.29	**62.94**	**79.80**	76.05	**70.12**
	Cosine Similarity (Ours)	89.81	51.42	**81.55**	**73.25**	**77.47**	**78.38**	62.41	87.59	**77.65**	72.66

A.1.4. Cosine Similarity as A Score Function

We can also regard the cosine similarity weights in the RW strategy as a score function, and conduct several experiments in Table 4 and Table 7. We notice that the cosine similarity also achieves a significant improvement compared with baselines in main evaluation tasks. Yet we notice that the cosine similarity is sensitive to the ID data distribution, since the performance in random sampling experiments (see Table 7) is not good enough compared with RP+GradNorm.

Table 5: OOD detection performance with ResNet101 trained on different imbalanced ID datasets. KL stands for using only KL divergence as the OOD detection function. ↑ indicates larger values are better and ↓ indicates smaller values are better. All values are percentages.

Datasets	Method	iNaturalist		SUN		Places		Textures		Average	
		AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
ImageNet-LT-a2	KL	60.66	95.51	75.21	83.06	71.94	86.15	40.36	95.25	62.04	89.99
ImageNet-LT-a3	KL	68.14	94.26	74.91	83.22	70.32	88.10	44.55	94.10	64.48	89.92
ImageNet-LT-a4	KL	66.58	94.78	72.97	86.35	69.66	89.62	44.97	95.25	63.54	91.50
ImageNet-LT-a5	KL	65.14	95.88	75.02	86.18	72.12	88.03	45.96	96.01	64.56	91.53
ImageNet-LT-a6	KL	68.14	93.77	72.30	89.82	70.44	89.86	43.10	96.15	63.50	92.40
ImageNet-LT-a7	KL	59.71	98.10	70.32	92.18	68.12	91.81	43.49	97.41	60.41	94.88
ImageNet-LT-a8	KL	56.33	98.72	71.09	93.29	70.43	91.13	42.27	98.12	60.03	95.32

Table 6: Performance comparison of different data type. All methods are based on ResNet101 trained on ImageNet-LT-a8. All values are percentages.

Method	Data Type	iNaturalist		SUN		Places		Textures	
		AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
GradNorm (Huang et al., 2021)	Overall	82.51	72.19	74.57	78.10	70.67	86.58	57.31	84.95
	Head	82.68	66.48	73.13	73.18	68.86	82.66	51.61	83.59
	Mid	83.85	69.50	75.88	75.68	71.91	84.41	55.99	85.48
	Tail	82.51	75.77	75.13	80.78	71.64	88.85	58.62	86.62
RP+GradNorm (Ours)	Overall	91.23(+8.72)	43.83(-28.36)	77.05(+2.48)	74.20(-3.90)	72.36(+1.69)	83.73(-2.85)	63.05(+5.74)	79.95(-5.00)
	Head	87.56(+4.88)	51.94(-14.54)	68.13(-5.00)	78.77(+5.59)	61.71(-7.15)	88.02(+5.36)	52.56(+0.94)	83.39(-0.20)
	Mid	91.42(+7.56)	42.26(-27.24)	76.87(+0.99)	73.40(-2.29)	72.12(+0.21)	83.30(-1.11)	63.32(+7.33)	79.92(-5.57)
	Tail	93.38(+10.87)	37.58(-38.19)	81.82(+6.69)	70.33(-10.45)	78.51(+6.87)	80.98(-7.87)	69.40(+10.78)	77.01(-9.61)

Table 7: Performance comparison under random sampling. All methods are based on ResNet101 trained on different imbalanced ID dataset with tail index $a = 8$. The results are means $\pm$ standard errors among ten randomly sampled datasets. ↑ indicates larger values are better and ↓ indicates smaller values are better. All values are percentages.

Method	iNaturalist		SUN		Places		Textures	
	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
MSP (Hendrycks & Gimpel, 2017)	65.19±0.95	93.77±0.58	66.18±2.91	91.49±0.62	59.83±1.19	93.44±0.47	46.94±0.81	97.16±0.26
ODIN (Liang et al., 2018)	59.30±0.86	96.99±0.31	66.82±3.36	94.2±0.48	59.84±1.35	94.65±0.40	39.79±0.62	98.32±0.10
Mahalanobis (Lee et al., 2018a)	52.70±1.05	94.56±0.72	50.12±5.09	95.79±1.32	49.72±3.50	95.66±1.04	54.79±3.94	91.36±2.52
Energy (Liu et al., 2020)	53.96±0.77	98.20±0.17	65.44±3.72	96.26±0.49	58.67±1.46	95.54±0.51	37.09±0.76	99.07±0.11
GradNorm (Huang et al., 2021)	81.03±0.56	68.73±1.40	76.03±1.77	70.39±1.74	70.87±0.77	83.39±1.18	61.78±0.83	79.80±0.80
RP+GradNorm (Ours)	**83.48±1.03**	**61.93±2.26**	**78.74±2.43**	**62.66±1.81**	**74.20±0.78**	**77.35±1.42**	**64.11±0.73**	**76.80±0.69**
Cosine Similarity (Ours)	71.99±7.05	84.29±13.22	71.72±3.87	82.04±3.06	66.63±4.70	87.99±3.45	51.46±4.44	92.95±1.93

Table 8: OOD detection performance with model size increasing. The RP+GradNorm method is trained on ImageNet-LT-a8. All values are percentages.

Model	iNaturalist		SUN		Places		Textures		Average	
	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
ResNet50	89.18	51.42	81.55	73.25	77.47	78.38	62.41	87.59	77.65	72.66
ResNet101	90.34	47.35	81.30	75.55	77.84	80.01	62.04	88.05	77.88	72.74
ResNet152	89.19	50.08	79.53	79.20	76.25	82.63	61.41	89.17	76.59	75.27
MobileNet	82.55	65.76	79.73	69.31	73.84	79.47	69.09	78.21	76.30	73.19

A.1.5. Performance Evaluation under Random Sampling

To further evaluate the performance of our method RP+GradNorm, we conduct experiments in 10 different ID datasets, which are generated randomly by the Pareto distribution with $a = 2$ and $a = 8$ from ImageNet-1K. The results on ImageNet-a8 dataset are reported in Table 7. From this table, it can be observed that, our method RP+GradNorm can outperform baseline methods on all evaluation tasks. As a highlight, RP+GradNorm reduces FPR95 from 70.39% to 62.66%. All these results

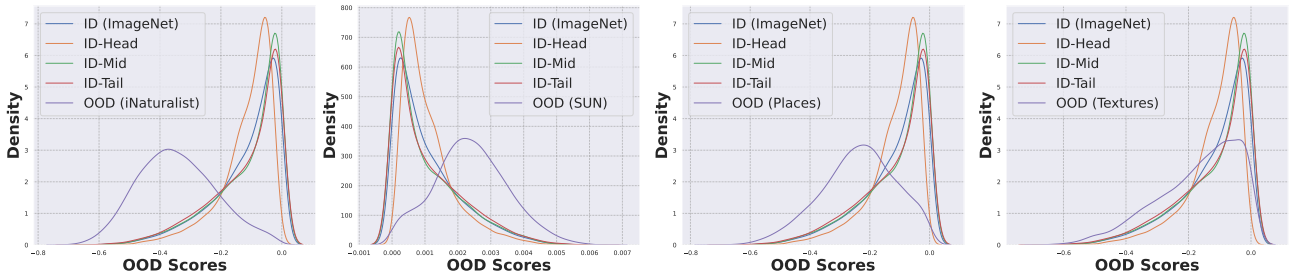

(a) MSP

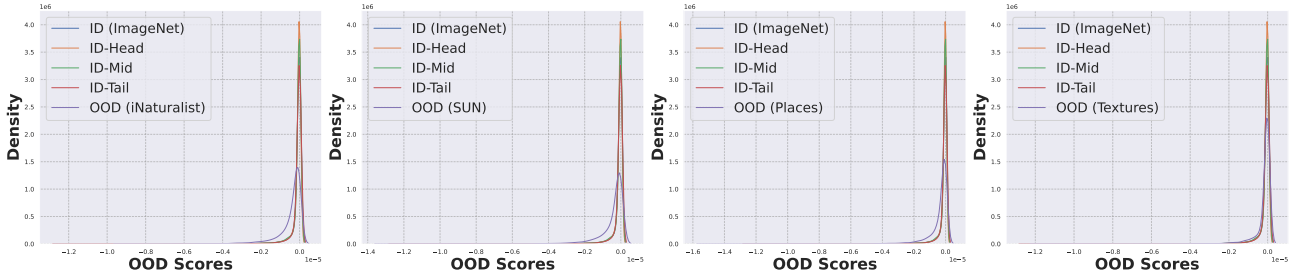

(b) RW+MSP

Figure 8: OOD score distribution of (a) MSP and (b) RP+MSP.

Table 9: Performance comparison under random sampling. All methods are based on ResNet101 trained on different imbalanced ID dataset with tail index $a = 2$. The results are means $\pm$ standard errors among ten randomly sampled datasets. $\uparrow$ indicates larger values are better and $\downarrow$ indicates smaller values are better. All values are percentages.

Method	iNaturalist		SUN		Places		Textures	
	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
MSP (Hendrycks & Gimpel, 2017)	72.55±0.39	87.64±0.59	68.23±0.35	87.33±0.27	64.82±0.25	90.21±0.27	50.58±0.48	96.02±0.26
ODIN (Liang et al., 2018)	70.70±0.63	89.41±0.75	72.03±0.43	85.31±0.55	67.26±0.38	89.28±0.46	44.86±0.69	95.57±0.28
Energy (Liu et al., 2020)	67.36±0.89	91.14±0.83	72.52±0.63	85.27±0.88	67.33±0.54	89.67±0.61	43.19±0.74	95.42±0.27
GradNorm (Huang et al., 2021)	83.66±0.93	61.90±2.25	78.40±0.76	66.78±1.31	72.09±0.67	79.47±1.09	61.80±0.72	80.21±0.88
RP+GradNorm (Ours)	**83.77±0.95**	**61.83±2.30**	**78.84±0.77**	**65.79±1.40**	**72.50±0.69**	**78.65±1.17**	**61.91±0.74**	**80.12±0.89**

show the our method RP+GradNorm still outperforms all baselines under random sampling. ImageNet-a2 dataset is more similar to the balanced dataset. The results are reported in Table 9. It can be observed that, our method RP+GradNorm can still outperform baseline methods on all evaluation tasks.

A.1.6. Ablation Study between Proposed Strategies

To further explore the combination ways between existing methods and the proposed strategies, we conducted a series of experiments involving RW+MSP, RW+GradNorm, and RW+RP+MSP/GradNorm. The details of these experiments can be found in Table 10. The results illustrate that MSP and RW are the best match while GradNorm matches RP well, in line with the scenario where RW and RP are suitable for implicit and explicit distribution respectively.

Our recommendation for practitioners is to prioritize the RP strategy, as it does not alter the original method's functioning and can revert to the original method when the ID data is balanced. The RW strategy is a viable option to mitigate the impact of data imbalance when some methods do not explicitly assume uniform distribution. Moreover, as demonstrated in Appendix A.1.8, when the data is balanced, the RW strategy has minimal influence on the original method's performance.

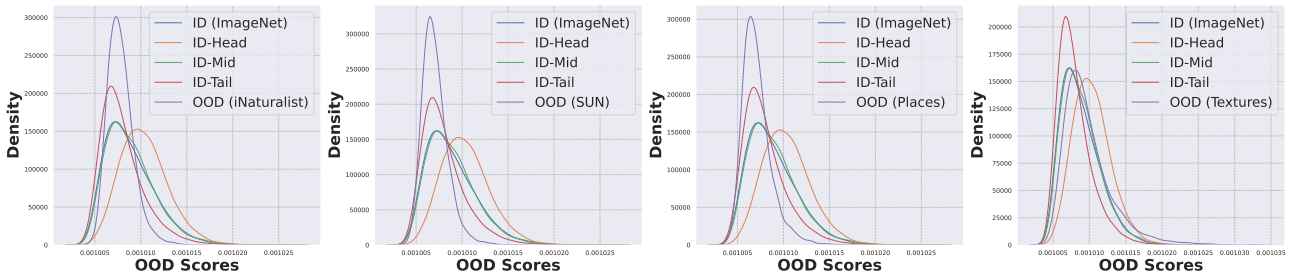

(a) ODIN

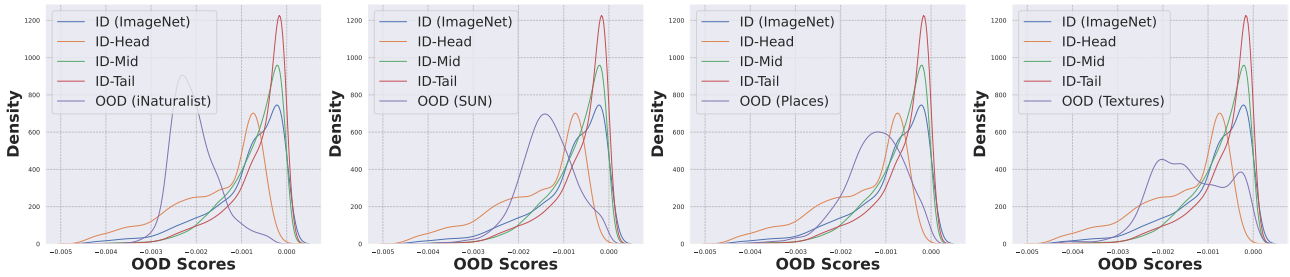

(b) RW+ODIN

Figure 9: OOD score distribution of (a) ODIN and (b) RW+ODIN.

Table 10: Ablation study between proposed strategies. All methods are trained on ImageNet-LT-a8 dataset with ResNet101.

Method	Strategy		iNaturalist		SUN		Places		Textures		Average	
	RP	RW	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
MSP	w/o	w/o	63.95	97.72	66.60	93.13	66.84	92.11	42.74	98.79	60.03	95.44
	w/	w/o	64.95	96.44	67.39	91.79	67.46	91.16	43.05	98.51	60.71	94.48
	w/o	w/	**71.29**	**88.46**	**70.19**	**87.59**	**69.66**	**87.84**	**45.16**	**96.90**	**64.08**	**90.20**
	w/	w/	67.40	94.79	68.23	91.12	68.30	90.34	43.71	98.09	61.91	93.58
GradNorm	w/o	w/o	82.51	72.19	74.57	78.10	70.67	86.58	57.31	84.95	71.26	80.45
	w/	w/o	**91.23**	**43.87**	77.36	**73.53**	72.67	83.29	**62.94**	**79.80**	**76.05**	**70.12**
	w/o	w/	88.31	55.53	**78.49**	73.60	**74.60**	**81.00**	58.75	84.96	75.04	73.77
	w/	w/	82.56	85.92	77.89	89.06	76.30	88.23	58.73	96.37	73.87	89.89

A.1.7. OOD Detection with Long-tailed Learning

We train ResNet50 with some long-tailed methods, like LDAM (Cao et al., 2019) and CMO (Park et al., 2022). Then we evaluate it with different OOD detection methods and our strategies, as shown in Table 11. The models with LDAM loss do perform better than those with CrossEntropy loss, but after applying our strategies, there are also significant improvements in all methods. However, CMO does not bring the performance improvement of OOD detection as LDAM does, and even performs worse than CrossEntropy. We think that this phenomenon indicates that not all long-tailed training methods are helpful to improve the OOD detector. But the results show that our strategies still works well while the models try to overcome the class imbalance in training time.

At last, We would like to reiterate our view on class-imbalanced OOD detection:

- Data imbalance is a common phenomenon, and even a slight imbalance (like the ImageNet-LT-a2 dataset) can still lead to a decrease in the performance of the OOD detector. After applying our strategies, this phenomenon can be improved.
- Developers do not necessarily use strategies to overcome data imbalance during the training phase of the model, depending on whether developers need to pay more attention to the minority in specific applications.
- Even if developers use strategies to overcome data imbalance during the training time, it is very hard to obtain a

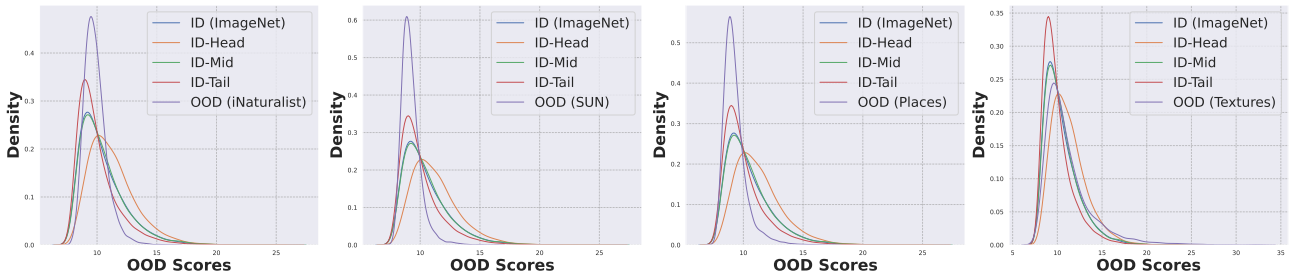

(a) Energy

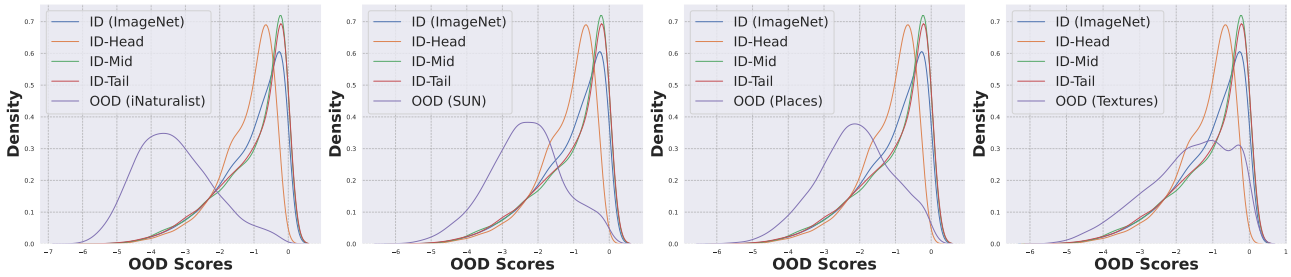

(b) RW+Energy

Figure 10: OOD score distribution of (a) Energy and (b) RW+Energy.

class-balanced classifier. Experiment results show that our method can achieve performance improvements with or without a strategy to overcome data imbalance.

A.1.8. Evaluation on the Balanced Dataset

As for RP strategy, when the training dataset is balanced (the class-prior distribution is uniform distribution in Eqs. (4) and (6)), RP+Method will be same as the original method.

For RW strategy, we conduct experiments on full ImageNet dataset (note that it is balanced) and the results are shown in Table 12. The performances of Energy and ODIN are quite close to the performnce of Energy and ODIN using RW strategy on the balanced dataset.

A.1.9. Ablation on the Similarity Metric

The performance of RW does indeed depend on the choice of measures used to evaluate the difference between the model outputs and the ID-class prior. Therefore, We explore how different measures influence the performance of RW in Table 13.

Here, we implement KL and JS divergence for RW, and the results are presented below. Specifically, as the range of KL divergence is $[0, +\infty)$, we used the sigmoid function to normalize it to $[0, 1)$. Additionally, KL divergence is an asymmetrical function, so we calculate both KL(pred, target) and KL(target, pred). Besides, as the range of JS divergence is $[0, 1]$, so we make it work as Eq. (7) in our paper.

The experimental results show that the results obtained using different distance metric functions are similar. This suggests that the effectiveness of RW is due to its rebalancing behavior, and that the choice of distance metric function is relatively robust.

A.2. Evaluation on Near OOD Benchmark

We follow the benchmark setting in the paper (Fort et al., 2021) to train ResNet50 on Imbalaned Cifar10 / Cifar100 and evaluate the model on Cifar100 / Cifar10. The bold in Table 14 below means better performance. Basically, our method can obtain some improvements, although not too much. We think that near-OOD benchmark is more challenging.

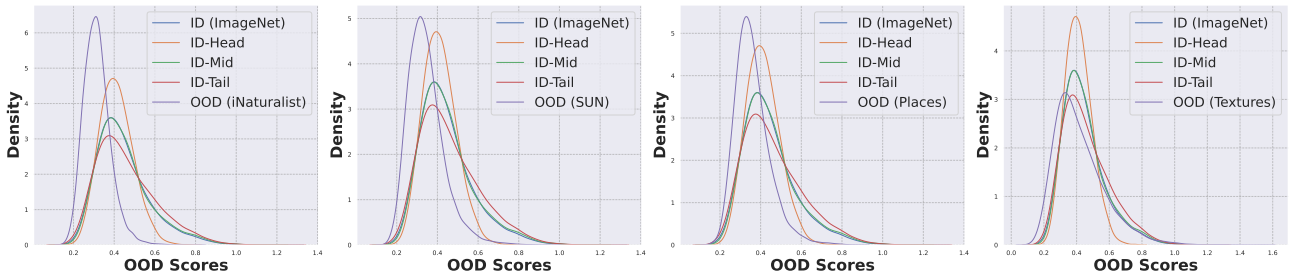

(a) GradNorm

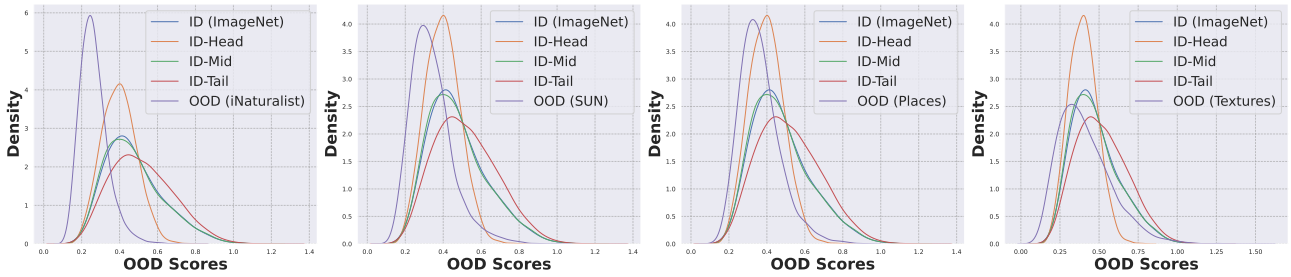

(b) RP-GradNorm

Figure 11: OOD score distribution of (a) GradNorm and (b) RW+GradNorm.

B. Further Analysis

B.1. Proof of Theorem 3.2

Theorem 1. If Assumption 3.1 holds, then $\mathbb{P}_{Y^{\text{in}}|X^{\text{out}}}(y|\mathbf{x}) = \mathbb{P}_{Y^{\text{in}}}(y)$, for any $y \in \mathcal{Y}^{\text{in}}$.

Proof. Using Assumption 3.1 in the second equation, we have

$$\mathbb{P}_{Y^{\text{in}}|X^{\text{out}}}(y|\mathbf{x}) = \frac{\mathbb{P}(Y^{\text{in}} = y \wedge X^{\text{out}} = \mathbf{x})}{\mathbb{P}(X^{\text{out}} = \mathbf{x})} = \frac{\mathbb{P}(Y^{\text{in}} = y)\mathbb{P}(X^{\text{out}} = \mathbf{x})}{\mathbb{P}(X^{\text{out}} = \mathbf{x})} = \mathbb{P}_{Y^{\text{in}}}(y).$$

□

B.2. Alternative Choice for ID-class-prior Distribution

When the labels of the training dataset are not available, we can use the predictions made by the model as an alternative to simulate empirical ID-class-prior distribution. Specifically, for each sample $\mathbf{x_i}$ in training dataset, the prediction made by the model is $\text{softmax}(\mathbf{f_\Theta}(\mathbf{x_i}))$. Thus, we have $\mathbb{P}_{Y^{\text{in}}} = 1/N * \sum_{i=0}^{N} \text{softmax}(\mathbf{f_\Theta}(\mathbf{x_i}))$.

We also conduct experiments to confirm the assumption, and the results are shown in Table 15. Noticeably, OOD detection performances with two kinds of ID-class-prior distribution are similar.

B.3. Proof of Theorem 3.3

Proof of Theorem 3.3. According to the definition of softmax function, it is clear that

$$\sum_{i=1}^{K} \text{softmax}_i(\mathbf{f_\Theta}(\mathbf{x})) = 1 \;\text{ and }\; \text{softmax}_i(\mathbf{f_\Theta}(\mathbf{x})) \geq 0, \;\text{ for } \forall\, i = 1, ..., K.$$

Existence. If we assume that for all $i = 1, ..., K$, $\text{softmax}_i(\mathbf{f_\Theta}(\mathbf{x})) < \frac{1}{K}$, then

$$\sum_{i=1}^{K} \text{softmax}_i(\mathbf{f_\Theta}(\mathbf{x})) < 1, \;\text{ which is conflict with } \sum_{i=1}^{K} \text{softmax}_i(\mathbf{f_\Theta}(\mathbf{x})) = 1.$$

Table 11: OOD detection performance with long-tailed learning methods.

Method	Long-tailed Method	iNaturalist		SUN		Places		Textures		Average	
		AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
MSP	CrossEntropy	62.50	97.30	67.46	92.04	66.89	91.49	42.64	98.44	59.87	94.82
RP+MSP	CrossEntropy	63.45	95.88	67.83	90.85	67.41	90.47	42.74	98.26	60.36	93.87
MSP	LDAM	68.12	95.22	70.04	89.39	69.72	87.71	43.04	98.07	62.73	92.60
RP+MSP	LDAM	68.95	93.66	**70.67**	**88.07**	**70.19**	**86.87**	43.27	97.78	**63.27**	91.60
MSP	CMO	**71.13**	**83.77**	59.35	93.81	60.35	92.77	**49.13**	**93.58**	59.99	**90.98**
RP+MSP	CMO	70.82	84.59	58.92	94.62	59.96	93.49	48.99	93.85	59.67	91.64
ODIN	CrossEntropy	59.34	98.59	72.57	91.68	70.85	90.83	42.00	98.16	61.19	94.81
RW+ODIN	CrossEntropy	84.04	76.58	68.25	95.94	65.69	96.43	53.00	92.16	67.75	90.28
ODIN	LDAM	64.08	98.28	74.40	**90.79**	**72.93**	**89.01**	39.74	98.37	62.79	94.11
RW+ODIN	LDAM	**85.94**	**71.28**	**68.59**	96.92	66.77	96.63	50.30	93.53	**67.90**	89.59
ODIN	CMO	73.67	81.36	58.94	93.39	59.57	92.54	49.34	92.41	60.38	89.93
RW+ODIN	CMO	77.19	81.55	53.21	95.26	56.84	93.90	**57.03**	**83.67**	61.07	**88.60**
Energy	CrossEntropy	56.25	98.95	73.60	91.56	71.32	90.37	42.68	98.10	60.96	94.75
RW+Energy	CrossEntropy	90.57	42.68	80.09	76.41	76.12	80.78	**64.68**	**86.17**	77.87	71.51
Energy	LDAM	59.93	98.80	74.74	91.99	72.74	90.39	39.51	98.37	61.73	94.89
RW+Energy	LDAM	**92.32**	**35.76**	**82.38**	**74.25**	**78.61**	**79.08**	63.32	86.35	**79.16**	**68.86**
Energy	CMO	73.24	85.77	55.50	95.64	55.64	95.54	48.93	91.56	58.33	92.13
RW+Energy	CMO	79.75	71.54	54.32	97.12	53.63	96.51	52.31	90.32	60.00	88.87
GradNorm	CrossEntropy	80.61	74.33	78.73	68.77	72.78	82.15	58.12	81.97	72.56	76.80
RP+GradNorm	CrossEntropy	89.85	50.03	80.73	64.52	74.69	78.18	**63.31**	**77.73**	77.14	67.62
GradNorm	LDAM	87.05	58.11	81.20	66.00	76.01	78.42	55.95	81.86	75.05	71.10
RP+GradNorm	LDAM	**92.87**	**36.75**	**82.70**	**63.74**	**77.41**	**77.21**	60.50	79.01	**78.37**	**64.18**
GradNorm	CMO	78.56	77.25	57.91	95.74	56.94	96.06	52.86	89.79	61.57	89.71
RP+GradNorm	CMO	84.90	63.05	63.65	92.25	61.24	94.42	59.17	86.06	67.24	83.95

Table 12: OOD detection performances on full ImageNet dataset.

Method	RW Strategy	iNaturalist		SUN		Places		Textures		Average	
		AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
ODIN	w/o	89.81	45.74	83.01	63.84	82.50	66.48	81.31	65.51	84.16	60.39
	w/	89.87	45.37	83.05	63.59	82.54	66.32	81.36	65.27	84.20	60.14
Energy	w/o	93.28	37.62	88.60	49.48	87.36	54.61	86.80	53.40	89.01	48.78
	w/	93.48	35.86	88.48	50.13	87.29	54.80	86.72	53.24	88.99	48.51

Therefore, there is at least one i such that $\mathrm{softmax}_i(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})) \geq \frac{1}{K}$, which implies that

$$\min_{\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})} S_{\mathrm{MSP}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) \geq \frac{1}{K}.$$

Note that when $\mathrm{softmax}(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})) = \mathbf{u}$, $S_{\mathrm{MSP}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) = \frac{1}{K}$, which implies that there exists $\tilde{\mathbf{f}}_{\mathbf{\Theta}}(\mathbf{x}) \in \arg\min_{\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})} S_{\mathrm{MSP}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x})$ such that

$$\mathbf{u} = \mathrm{softmax}(\tilde{\mathbf{f}}_{\mathbf{\Theta}}(\mathbf{x})).$$

Uniqueness. If there is $\mathbf{f}^*_{\mathbf{\Theta}}(\mathbf{x}) \in \arg\min_{\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})} S_{\mathrm{MSP}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x})$ such that $\mathrm{softmax}(\mathbf{f}^*_{\mathbf{\Theta}}(\mathbf{x})) \neq \mathbf{u}$, it is clear that

$$\sum_{i=1}^{K} \mathrm{softmax}_i(\mathbf{f}^*_{\mathbf{\Theta}}(\mathbf{x})) < 1,$$

which is conflict with $\sum_{i=1}^{K} \mathrm{softmax}_i(\mathbf{f}_{\mathbf{\Theta}}(\mathbf{x})) = 1$. Therefore, $\mathrm{softmax}(\mathbf{f}^*_{\mathbf{\Theta}}(\mathbf{x})) = \mathbf{u}$.

Combining the results in existence and uniqueness, we have completed this proof. □

B.4. Discussion about RP+ODIN

ODIN (Liang et al., 2018) is an enhanced version of MSP, whose main improvement is the introduction of a temperature scaling strategy. The temperature parameter T smoothes the prediction distribution of the softmax function and thus making

Table 13: Ablation on the similarity metric

Method	RW	iNaturalist		SUN		Places		Textures		Average	
		AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
Energy	cos	83.12	65.67	80.13	77.51	77.12	80.49	51.70	91.31	73.02	78.75
	KL(pred, target)	76.52	74.40	79.10	73.60	75.40	77.40	47.06	89.80	69.52	78.80
	KL(target, pred)	70.92	83.60	77.95	77.60	74.45	78.90	46.27	91.00	67.40	82.78
	JS	83.82	64.20	81.83	69.70	77.71	75.10	51.55	88.30	73.73	74.33
ODIN	cos	90.24	47.52	81.33	75.44	77.90	79.90	61.43	88.03	77.73	72.72
	KL(pred, target)	89.98	44.00	82.04	67.70	77.84	75.00	55.34	85.20	76.30	67.98
	KL(target, pred)	86.10	65.50	81.16	70.70	76.72	76.30	53.14	87.20	74.28	74.93
	JS	89.10	52.00	82.95	68.20	78.75	74.90	56.19	86.00	76.75	70.28

Table 14: Evaluation on near OOD benchmark.

ID Dataset	OOD Dataset	Method	AUROC↑	FPR95↓
Imbalanced Cifar10	Cifar 100	MSP	**81.76**	78.45
		RP+MSP	74.75	**72.90**
		ODIN	83.53	67.65
		RW+ODIN	**85.14**	**65.85**
		Energy	83.61	66.68
		RW+Energy	**85.40**	**64.23**
		GradNorm	29.21	99.45
		RP+GradNorm	**40.62**	**93.63**
Imbalanced Cifar100	Cifar10	MSP	**68.49**	87.91
		RP+MSP	67.75	**87.86**
		ODIN	70.83	**86.42**
		RW+ODIN	**70.90**	86.43
		Energy	70.89	**86.15**
		RW+Energy	**71.03**	86.57
		GradNorm	34.89	**99.06**
		RP+GradNorm	**35.72**	99.15

the prediction sparser and more similar to the uniform distribution.

$$S_{\text{ODIN}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) = \max_i \frac{\exp\left(\mathbf{f}_i(\mathbf{x})/T\right)}{\sum_{j=1}^{C} \exp\left(\mathbf{f}_j(\mathbf{x})/T\right)} \tag{12}$$

Since ODIN maps the prediction distribution of the softmax layer to another distribution space while we need to measure the similarity between the class-prior distribution and the model-predicted distribution, we need to use the same mapping method to deal with the class-prior distribution $\mathbb{P}_{Y^{\text{in}}} = [p_1, p_2, ..., p_C]$, as follows:

$$\mathbb{P}'_{Y^{\text{in}}} = \left[\frac{\exp\left(p_1/T\right)}{\sum_{j=1}^{C} \exp\left(p_j/T\right)}, \frac{\exp\left(p_2/T\right)}{\sum_{j=1}^{C} \exp\left(p_j/T\right)}, \ldots, \frac{\exp\left(p_C/T\right)}{\sum_{j=1}^{C} \exp\left(p_j/T\right)}\right] \tag{13}$$

Then, we use this new class-prior distribution $\mathbb{P}'_{Y^{\text{in}}}$ to modify ODIN with RP strategy as Eq. (15).

$$\mathbf{h}_{\mathbf{\Theta}}(\mathbf{x}) = \left[\frac{\exp\left(\mathbf{f}_1(\mathbf{x})/T\right)}{\sum_{j=1}^{C} \exp\left(\mathbf{f}_j(\mathbf{x})/T\right)}, \frac{\exp\left(\mathbf{f}_2(\mathbf{x})/T\right)}{\sum_{j=1}^{C} \exp\left(\mathbf{f}_j(\mathbf{x})/T\right)}, \ldots, \frac{\exp\left(\mathbf{f}_C(\mathbf{x})/T\right)}{\sum_{j=1}^{C} \exp\left(\mathbf{f}_j(\mathbf{x})/T\right)}\right] \tag{14}$$

$$S_{\text{RP+ODIN}}(\mathbf{f}_{\mathbf{\Theta}}, \mathbf{x}) = \max\left(\mathbf{h}_{\mathbf{\Theta}}(\mathbf{x}) - \mathbb{P}'_{Y^{\text{in}}}\right) \tag{15}$$

When we follow the default setting $T = 1000$ in ODIN, we notice that $\mathbb{P}'_{Y^{\text{in}}}$ will be quite close to the uniform distribution, where each element is close to $1/K$. Thus,Eq. (15) can be regarded as $\mathbf{h}_{\mathbf{\Theta}}(\mathbf{x})$ minus a constant.

Table 15: Performance comparison of two different ID-class-prior distribution acquisition methods.

Method	ID-class-prior Distribution	iNaturalist		SUN		Places		Texture		Average	
		AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
RP+MSP	Model Prediction	63.36	96.05	68.19	90.40	67.44	90.58	42.93	98.17	60.48	93.80
	Data Label	63.34	96.09	68.19	90.44	67.43	90.64	42.93	98.17	60.48	93.84
RW+ODIN	Model Prediction	84.02	76.52	68.62	95.56	65.85	96.28	53.21	90.22	67.93	89.65
	Data Label	84.04	76.49	68.57	95.59	65.81	96.29	53.23	91.99	67.91	90.09
RW+Energy	Model Prediction	90.64	42.30	80.51	75.14	76.39	80.29	64.64	85.78	78.05	70.88
	Data Label	90.56	42.69	80.46	75.34	76.27	80.48	64.79	85.90	78.02	71.10
RP+GradNorm	Model Prediction	89.89	49.89	80.75	64.33	74.70	78.11	63.36	77.70	77.17	67.51
	Data Label	89.85	50.03	80.73	64.52	74.69	78.18	63.31	77.73	77.14	67.60

B.5. Discussion about m in Pareto Distribution

For each class x_i, the sample number is

$$y_i = N \times p(x_i) = N \times \frac{am^a}{x_i^{a+1}}, \tag{16}$$

where a is tail index, m is a constant and N is the sample number of the ImageNet-1K dataset.

After sampling, the new data distribution for each class is

$$p(y_i) = \frac{y_i}{\sum_{i=1}^{K} y_i} = \frac{N \times \frac{am^a}{x_i^{a+1}}}{\sum_{i=1}^{K}(N \times \frac{am^a}{x_i^{a+1}})} = \frac{1}{x_i^{a+1}\sum_{i=1}^{K}\frac{1}{x_i^{a+1}}}. \tag{17}$$

Obviously, the value of m do not affect the imbalance degree of sampled datasets. Thus, we keep $m = 1$ unchanged.

B.6. Discussion about Feature-based Methods

Feature-based methods, like KNN (Sun et al., 2022), need a training set to generate class prototypes, i.e., an average feature vector for each category. Under class-imbalanced situations, prototypes of tailed classes would be more unreliable than the majority due to the limitation of training samples. We think using ID-class-prior distribution to reweight features may be an effective way to solve the imbalanced problem in feature space.

B.7. Discussion about possibility for RP+MSP.

In order to discuss about the possibility for aligning the minimizer of the score function with the class priors, we conduct experiments for $\max_i(\text{softmax}_i(\mathbf{f}(\mathbf{x}))/\mathbb{P}_{Y^{\text{in}}}(i))$ and show the corresponding results in the below table. Our experiments show that $\max_i(\text{softmax}_i(\mathbf{f}(\mathbf{x}))/\mathbb{P}_{Y^{\text{in}}}(i))$ also performs very well.

Table 16: OOD detection performances on ImageNet-LT-a8 dataset.

Method	iNaturalist		SUN		Places		Textures		Average	
	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓	AUROC↑	FPR95↓
$\max_i(\text{softmax}_i(\mathbf{f}(\mathbf{x}))/\mathbb{P}_{Y^{\text{in}}}(i))$	81.76	69.75	57.80	94.80	54.90	94.68	52.01	88.97	61.62	87.05