

Gradient-Regularized Out-of-Distribution Detection

Sina Sharifi^{1*}, Taha Entesari^{1*}, Bardia Safaei^{1*},
Vishal M. Patel¹, and Mahyar Fazlyab¹

Johns Hopkins University, Baltimore MD 21218, USA
{sshari12, tentesa1, bsafaei1, vpatel36, mahyarfazlyab}@jhu.edu

Abstract. One of the challenges for neural networks in real-life applications is the overconfident errors these models make when the data is not from the original training distribution. Addressing this issue is known as Out-of-Distribution (OOD) detection. Many state-of-the-art OOD methods employ an auxiliary dataset as a surrogate for OOD data during training to achieve improved performance. However, these methods fail to fully exploit the local information embedded in the auxiliary dataset. In this work, we propose the idea of leveraging the information embedded in the gradient of the loss function during training to enable the network to not only learn a desired OOD score for each sample but also to exhibit similar behavior in a local neighborhood around each sample. We also develop a novel energy-based sampling method to allow the network to be exposed to more informative OOD samples during the training phase. This is especially important when the auxiliary dataset is large. We demonstrate the effectiveness of our method through extensive experiments on several OOD benchmarks, improving the existing state-of-the-art FPR95 by 4% on our ImageNet experiment. We further provide a theoretical analysis through the lens of certified robustness and Lipschitz analysis to showcase the theoretical foundation of our work. Our code is available at <https://github.com/o4lc/Greg-OOD>.

Keywords: Out-of-Distribution Detection, Gradient Regularization, Energy-based Sampling

1 Introduction

Neural networks are increasingly being utilized across a wide range of applications and fields, achieving unprecedented performance levels and surpassing traditional state-of-the-art approaches. However, concerns about robustness and safety, coupled with significant challenges in verifying robustness and ensuring safe performance impede their use in more sensitive applications [2, 32]. One concern that arises when deep models are deployed in the real world is their tendency to produce over-confident predictions upon encountering unfamiliar samples that are distant from the space in which they were trained [47, 51]. As

*Co-first authors, equally contributed to the work




Fig. 1: *Left:* Effect of the local structure of the score manifold on a two-dimensional toy example for OOD detection. The grey plane depicts the score function’s decision threshold. An equal amount of perturbation in these two scenarios results in different OOD detections, highlighting the importance of the local structure of the score manifold. *Right:* t-SNE plot showing the representation of ID and OOD datasets for CIFAR experiments.

a result, the field of Out-of-Distribution (OOD) detection has emerged to study and address this phenomenon.

Many works have been presented in recent years, trying different methods and ideas to regulate the network’s output to OOD samples in terms of softmax probability [22, 48, 76] or energy metrics [36]. One family of methods, known as post-hoc approaches [1, 34, 37, 41, 49, 53, 60, 66, 72], operate on a given trained network and employ certain statistics, such as the confidence level of the network, the number of active neurons, the contribution of important pathways in the network, etc., to perform OOD detection without using any samples from an auxiliary distribution. Other methods [7, 23, 27, 36, 43] use OOD data for training or finetuning a network. The latter family usually tends to achieve better performances in comparison to post-hoc methods, as their access to the auxiliary dataset provides them with additional information regarding the outlier distribution. OOD detection is closely related to anomaly [54, 69] and novelty detection [10, 39], each with slightly different problem formulations. Our work focuses on the general OOD detection problem.

In this paper, we propose *GReg*, a novel **G**radient **R**egularized OOD detection method, which aims to learn the local information of the score function to improve the OOD detection performance. Using information from the gradient of the score function, our method builds on top of current state-of-the-art OOD detection methods (that require training) to incorporate local information into the training process and lean the network towards a smoother manifold.

To motivate this idea, consider Figure 1a. This figure shows a two-dimensional example of a possible scenario in OOD detection. In this example, x_1 and x_2 are two OOD samples and the network has learned to assign the same OOD scores (denoted by red dots in the figure) to both samples and detect them correctly

as OOD. However, the difference in the local behavior of the network at these points, leaves samples in the vicinity of x_2 susceptible to misdetection, whereas, samples close to x_1 are always detected as OOD. To see this, consider the perturbations indicated by the red arrows, where the resulting points in each case represent two possible samples with equal distance from points x_1 and x_2 , respectively. However, as can be seen in the figure, the OOD detection algorithm cannot detect one of the samples correctly. This emphasizes the importance of regulating the local behavior of deep learning models, especially when OOD robustness is crucial. Furthermore, Figure 1b shows the t-SNE plot of the CIFAR and some of the well-known OOD benchmarks. This further confirms that even in high dimensions the ID data (CIFAR) are well-separated from the OOD data.

On top of that, several recent works have shown the effectiveness of some form of informative sampling or clustering on OOD detection [7, 27, 43]. Sampling gains more importance specifically when a large dataset of outlier samples is available and the network cannot be exposed to all of the data during training. Inspired by such works, we present *GReg+*, the coupling of *GReg* with an energy-based clustering method aimed at making better use of the auxiliary data during training. This is in line with our intuition of utilizing local information as clustering enables the use of samples from diverse regions of the feature space. The diversity put forward by the clustering mechanism allows *GReg+* to achieve its state-of-the-art performance in more regions of the space rather than only performing well on regions well-represented in the dataset.

In summary, our key contributions are:

- We propose the idea of regularizing the gradient of the OOD score function during the training (or fine-tuning). This allows the network to better learn the local information embedded in ID and OOD samples.
- We propose a novel energy-based sampling method that chooses samples –based on their energy levels– from the OOD dataset that represent more vulnerable regions of the space using clustering techniques.
- We provide empirical results and ablation studies on a wide range of architectures and different datasets to showcase the effectiveness of our method, along with an extensive comparison to the state-of-the-art. In addition, we provide a detailed theoretical analysis to justify our results.

To the best of our knowledge, this is the first work that utilizes the norm of the gradient of the score function during the training phase to learn the local behavior of the ID/OOD data aiming to improve the OOD detection performance.

2 Related Work

Several works have attempted to enable deep neural networks with OOD detection capabilities in recent years [4, 6, 16, 31, 42, 50, 55, 64, 75, 77]. These works can be divided into two main groups: methods that do not use auxiliary data (post-hoc) and those that utilize auxiliary data.

2.1 Post-hoc OOD Detection

In one of the earliest attempts to remedy the OOD detection problem, [22] uses the maximum softmax probability (MSP) score to recognize OOD samples during inference. ODIN [23] improves the MSP score’s separability by adding small noise to the input and utilizing temperature scaling. [25, 35–37] aim to formulate enhanced scoring functions to better distinguish OOD samples from ID samples. For example in [36], the authors propose using energy as a powerful OOD score. GradNorm [25] uses the vector norm of the gradient of the KL divergence between the softmax output and a uniform probability distribution to devise an OOD score. It was later observed by [58], that OOD data tend to have different activation patterns. As a solution, they propose activation truncation to improve OOD detection. Building on this observation, DICE [59] ranks weights based on the measure of contribution, and selectively uses the most contributing weights. LIne [1] further improves on DICE by using Shapley values [56] to more accurately detect important and contributing neurons.

2.2 OOD Detection with Auxiliary Outlier Dataset

When an auxiliary dataset is available, this group of methods [13, 14, 22, 36, 63, 68, 71, 74, 80] use this data to train the model to improve the network’s robustness against OOD data. OE [22] and Energy [8, 36] propose their respective loss functions to be used during the training of the network. Recently, [9] improved [36] by proposing balanced Energy loss that balances the number of samples of the auxiliary OOD data across classes. Furthermore, [52] increases the expected Frobenius Jacobian norm difference between ID and OOD. This is in sharp contrast with our intuition as we aim to decrease the norm of the score function in all regions. Another line of work [15, 19, 29, 44, 45] focuses on generating artificial samples that resemble the real OOD data, using various generative models such as GANs. Most recently, OpenMix [79] was proposed as a misclassification detection method, teaching the model to reject pseudo-samples generated from the available outlier samples.

Outlier Sampling: When large amounts of auxiliary data are available, choosing an informative and diverse set of samples to perform training becomes a priority. Many of the state-of-the-art OOD methods use their custom sampling techniques. NTOM [7] uses greedy outlier sampling where outliers are selected based on the estimated confidence. POEM [43] proposes a posterior sampling-based outlier mining framework for OOD detection. Most recently DOS [27] used K-Means [38] to perform clustering on the auxiliary dataset to provide the network with diverse informative OOD samples.

3 Preliminaries

3.1 Notation

We consider a classification setup where $\mathcal{X} \in \mathbb{R}^d$ denotes the input space and $\mathcal{Y} \in \{1, 2, \dots, K\}$ denotes the labels. We define a neural network $f: \mathbb{R}^d \rightarrow \mathbb{R}^K$

where $f_i(x)$ denotes the i -th logit of the neural network to an input x . We denote the feature extractor of the neural network as h , which is followed by a classification layer. We use $f(x) = Wh(x) + b$ for all the models in this paper. We denote the norm of a vector x as $\|x\|$. The log-sum-exp function is defined as $\text{LSE}(x) = \log \sum_{i=1}^n \exp(x_i)$. For simplicity, for a given score function S , we use $S(x)$ as a shorthand for $S(f(x))$.

The In-Distribution (ID), Out-of-Distribution (OOD) and Auxiliary distribution over \mathcal{X} are denoted by \mathcal{D}_{in} , \mathcal{D}_{out} and \mathcal{D}_{aux} , respectively. The corresponding datasets are assumed to be sampled i.i.d from their respective distributions. We denote the indicator function of an event $e \leq 0$ as $\mathbb{I}_{e \leq 0}$, i.e., if $e \leq 0$ then $\mathbb{I}_{e \leq 0} = 1$, and otherwise, $\mathbb{I}_{e \leq 0} = 0$. We denote $\mathcal{B}(x, \varepsilon) = \{y \mid \|y - x\| \leq \varepsilon\}$. Finally, a function f is said to be locally Lipschitz continuous on \mathcal{X} if there exists a positive constant L such that $\|f(x) - f(y)\| \leq L\|x - y\| \forall x, y \in \mathcal{X}$. The smallest such constant is called the local Lipschitz constant of f .

3.2 Out-of-Distribution Detection

It has been observed that when neural networks are faced with samples from a different distribution compared to their training distribution, they produce overconfident errors [21]. The goal of OOD detection is to differentiate the ID data from the OOD data. To achieve this goal, the main approach in the literature is to define a scoring function S and use a threshold γ to distinguish different samples. That is, for a sample x , if $S(x) \leq \gamma$, we label it as ID, and label it as OOD otherwise.

Most relevant to our setup, [36] uses the scoring function $S_{\text{En}}(x) = -\text{LSE}(f(x))$, and defines the following loss function to train the network to better differentiate the ID and OOD samples,

$$\begin{aligned} \mathcal{L}_{S_{\text{En}}} = & \mathbb{E}_{(x_{\text{in}}, y_{\text{in}}) \sim \mathcal{D}_{\text{in}}} [\mathbb{I}_{S_{\text{En}}(x_{\text{in}}) \geq m_{\text{in}}} (S_{\text{En}}(x_{\text{in}}) - m_{\text{in}})^2] \\ & + \mathbb{E}_{(x_{\text{aux}}, y_{\text{aux}}) \sim \mathcal{D}_{\text{aux}}} [\mathbb{I}_{S_{\text{En}}(x_{\text{aux}}) \leq m_{\text{aux}}} (m_{\text{aux}} - S_{\text{En}}(x_{\text{aux}}))^2], \end{aligned} \quad (1)$$

where m_{in} and m_{aux} define two thresholds to filter out points that already have an acceptable level of energy, i.e., if the energy score of an ID (OOD) sample is low (high) enough, it will be excluded from the energy loss.

4 Method

In this section, we propose our method which consists of regularizing the gradient of the score function by adding a new term to the loss, coupled with a novel energy-based sampling method to allow us to choose more informative samples during training. Figure 2 provides an overview of our method *GReg+* with a simple illustration of how the sampling algorithm utilizes clustering to provide more informative samples.

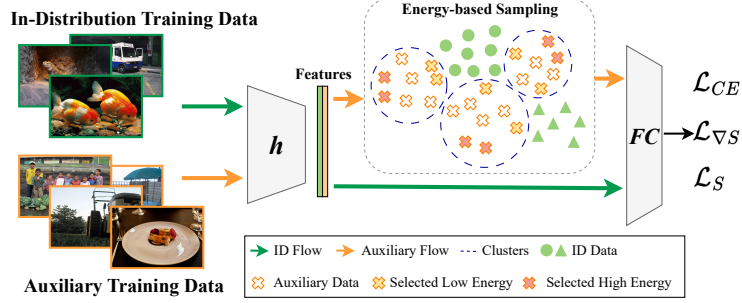


Fig. 2: Overview of *GReg+*. The gradient loss allows the method to obtain more local information from the training data. The sampling algorithm uses the normalized features of the OOD samples to perform clustering and choose the sampled data based on the energy score (see Algorithm 1) which will be used to calculate the gradient loss to expose the network to samples that can improve the performance of the model.

4.1 Gradient Regularization

Many of the current state-of-the-art methods focus only on the value of their scoring function S and aim at minimizing or maximizing it on different regions of the input space [9, 23, 27, 36, 43, 79]. We argue that by optimizing only based on the value of the score function, we are not utilizing all of the information that is embedded in the data. As such, the generalization of the model requires the auxiliary dataset to be a good representation of the actual OOD distribution and its corresponding support in the image space (which is extremely high-dimensional). In response, we aim to use the local information embedded in the data and propose to do so by leveraging the gradient of the score function S . In other words, we focus not only on the value of the score function S on the auxiliary samples but also on the local behavior of the score function around these data points. The rationale behind this approach is that, based on the definition of the problem, the distribution of ID and the OOD tends to be well separated. Therefore, the area close to an ID sample most likely does not include any OOD samples, and vice versa. Driven by this insight, we propose to add a regularization term $\mathcal{L}_{\nabla S}$ that promotes the smoothness of the score function around the training samples by penalizing the norm of its gradient.

Suppose x is a training sample. Using the first-order Taylor approximation of the score function around x , we have

$$S(x') \simeq S(x) + \nabla S(x)^\top (x' - x), \quad (2)$$

for x' sufficiently close to x . As a result, if $\|\nabla S(x)\|$ is small, then equation (2) hints that the difference $|S(x') - S(x)|$ would likely remain small, implying stability of the score manifold around the point x . To this end, we propose the following regularizer. The specific definition of $\mathcal{L}_{\nabla S}$ is a design choice, however, for the case of the energy loss, which is the main focus of this paper, we use the

following formulation.

$$\begin{aligned} \mathcal{L}_{\nabla S_{\text{En}}} = & \mathbb{E}_{(x_{\text{in}}, y_{\text{in}}) \sim \mathcal{D}_{\text{in}}} \|\mathbb{I}_{S_{\text{En}}(x_{\text{in}}) \leq m_{\text{in}}} \nabla_{x_{\text{in}}} S_{\text{En}}(x_{\text{in}})\| \\ & + \mathbb{E}_{(x_{\text{aux}}, y_{\text{aux}}) \sim \mathcal{D}_{\text{aux}}} \|\mathbb{I}_{S_{\text{En}}(x_{\text{aux}}) \geq m_{\text{aux}}} \nabla_{x_{\text{aux}}} S_{\text{En}}(x_{\text{aux}})\| \end{aligned} \quad (3)$$

To elaborate on the intuition behind this choice, consider the following cases:

1. If an ID (OOD) sample is correctly detected, we would want the local area around it to be detected the same.
2. If an ID (OOD) sample is mis-detected as OOD (ID), we do not want its local area to behave the same.

Hence, we select thresholds such that the loss only penalizes the gradient of the correctly detected samples. This is in sharp contrast to (1), where the aim is to penalize the energy score of the samples that have undesired energy levels. Therefore, by designing the gradient loss as (3), we allow the energy loss to train the samples that are not yet well learned and only use the samples whose score values are properly learned. Having specified $\mathcal{L}_{\nabla S}$, we train our model using the following loss $\mathcal{L} = \mathcal{L}_{\text{CE}} + \lambda_S \mathcal{L}_S + \lambda_{\nabla S} \mathcal{L}_{\nabla S}$, where λ_S and $\lambda_{\nabla S}$ are regularization constants and \mathcal{L}_{CE} is the well-known cross-entropy loss.

4.2 OOD Sampling

One concern faced by OOD training schemes is that the OOD data may outnumber the ID data by orders of magnitude. Having access to such a large auxiliary dataset begs the question of which samples should the network see during the training phase. This creates two possible scenarios. Either the training algorithm exposes the model to an equal number of both ID and OOD data points, which means that the model cannot fully utilize the OOD dataset as most of the OOD dataset remains unseen. Alternatively, the algorithm may provide disproportionately more OOD samples, leading to bias in the model due to the imbalance in the samples [62]. These observations hint towards employing a more informative sampling strategy [7, 33] that utilizes as much of the dataset as possible while avoiding the pitfalls caused by imbalanced datasets. Whilst the first option that comes to mind is to perform a greedy sampling based on the score function S , it has been observed [27] that such greedy samples could bias the model to some specific regions of the image space, resulting in a model unable to generalize well.

To address these issues, we present Algorithm 1, which uses clustering (to discourage greedy behavior) alongside energy-based scoring to sample more informative images. Next, we explain the sampling algorithm in more detail.

Clustering: Given samples $\{x_i\}_{i=1}^{n_{\text{OOD}}}$, we first calculate the features $z_i = h(x_i)$ and the scores $s_i = -\text{LSE}(f(x_i))$. The next step is to perform clustering in the *feature space*. To perform clustering, we use K-Means with a fixed number of clusters. As K-Means assigns clusters based on Euclidean distances, we normalize the features to $\hat{z}_i = \frac{z_i}{\|z_i\|_2}$ to avoid skewed clusters towards samples with disproportionate feature values.

As for the number of clusters, we use the number of samples in each mini-batch of training, i.e., if in a training iteration the model is presented with n_{ID} ID samples, we cluster the n_{OOD} OOD samples into $k = n_{\text{ID}}$ clusters. We study the choice of the number of clusters in the Supplementary Material. Finally, given a clustering $\{C_j\}_{j=1}^k$, we can acquire cluster labels l_i for each data point x_i , $i = 1, \dots, n_{\text{OOD}}$.

Energy-Based Sampling: The final step in our sampling algorithm is to choose the *best* samples from each cluster with respect to some criterion. As we use the energy scores s_i to perform OOD detection, we use the same scores for sample selection. Given our choice of loss functions in (1) and (3), we use the samples with the smallest energy scores s_i of each cluster to provide samples useful for the loss term (1), and we use the samples with the largest energy scores to provide samples that are useful for the loss term (3). Intuitively, $\mathcal{L}_{\mathcal{S}_{\text{En}}}$ aims to increase the energy score of the OOD samples, and so, we need to provide samples whose energy scores are small. On the other hand, $\mathcal{L}_{\nabla \mathcal{S}_{\text{En}}}$ focuses on samples that already have high energy scores to make the model locally behave similarly, and so, we need to also provide samples with high energy scores.

Overall, by clustering, we ensure that the model is exposed to *diverse* regions of the feature space, and by sampling using the scores we ensure that from each region more *informative* samples are used.

5 Theoretical Analysis

In this section, we focus on the theoretical implications of gradient regularization. Traditionally, during the training phase, there is no control over the smoothness of the network and the score function, which can lead to overly sensitive OOD detection. The smoothness of networks has been studied in depth in certified robustness [17, 26, 40, 57] and has been deemed as a desired or even necessary property for the robustness of neural networks. The desirable properties sought by smoothness consequently follow over to OOD detection; simply because OOD detection can be cast as a simple classification task that the aforementioned literature studies. Intuitively, the lack of smoothness decreases the network’s OOD robustness as the score values can abruptly change with small perturbations in

Algorithm 1 Energy-Based Sampling

Input: Auxiliary dataset $\mathcal{X} \sim \mathcal{D}_{\text{aux}}$ with n_{OOD} samples, number of clusters k , feature extractor h , score function S

Output: Selected samples: $\bar{\mathcal{X}} \subset \mathcal{X}$;

- 1: $z_i \leftarrow h(x_i)$, $s_i \leftarrow S(x_i)$, $i = 1, \dots, n_{\text{OOD}}$.
 - 2: $(\hat{z}_1, \dots, \hat{z}_{n_{\text{OOD}}}) \leftarrow h_N(z_1, \dots, z_{n_{\text{OOD}}})$
 - 3: Labels $l_i \leftarrow \text{K-Means}(\{\hat{z}_i\}_{i=1}^{n_{\text{OOD}}}, k)$.
 - 4: **for** $j = 1 \dots k$ **do**
 - 5: $\mathcal{I} \leftarrow \{i : l_i = j\}$
 - 6: $\bar{\mathcal{X}}_j \leftarrow \text{Sample } \{x_i\}_{i \in \mathcal{I}} \text{ based on } \{s_i\}_{i \in \mathcal{I}}$
 - 7: **end for**
 - 8: return $\bar{\mathcal{X}} \leftarrow \cup_{j=1}^k \bar{\mathcal{X}}_j$
-




Fig. 3: The evolution of the Energy loss and the norm of the gradient with respect to the input of the score function during training with and without gradient regularization. Gradient regularization reduces the slope of the increase of the gradient norm, without negatively affecting the reduction in energy loss. The sudden change in the behavior of the plots towards the final iterations is due to learning rate scheduling.

the image space. To support our intuition, consider Figure 3 which portrays the evolution of the energy loss and the norm of its gradient with respect to the input in two cases, with and without *GReg*. Although the energy loss decreases in a similar fashion in both scenarios, the norm of the gradient increases much more rapidly in the absence of gradient regularization. The excess increase in the average norm of the gradient potentially points to a more sensitive score manifold. Note that the general increasing trend of the norm of the gradient in Figure 3 is not unexpected as the network is initialized randomly.

We further motivate our idea using the concept of penalizing the norm of the gradient [5] or upper bounds on the norm of the gradient, like the Lipschitz constant [17, 26] for achieving robustness. In other words, if the gradient of the score function S is bounded around some arbitrary point x , there exists an ε -neighborhood in which our prediction of ID/OOD would not change. This means that if the point x is detected as ID (OOD), then all the points $x' \in \mathcal{B}(x, \varepsilon)$ are also provably ID (OOD). This is desirable for us as it is usually assumed that the ID and OOD data are well separated. Formally, suppose S satisfies the following local Lipschitz property on $\mathcal{B}(x, \varepsilon)$ where $x \sim \mathcal{D}_{\text{in}}$

$$|S(x') - S(x)| \leq L_S \|x' - x\|, \quad \forall x' \in \mathcal{B}(x, \varepsilon),$$

where $L_S > 0$. Suppose $x \sim \mathcal{D}_{\text{in}}$ is correctly classified by the score function as ID, i.e., $S(x) < \gamma$. If we want x' to be also labeled as ID, it is sufficient to satisfy $S(x) + L_S \|x' - x\| \leq \gamma$, which in turn gives the following certificate

$$\|x' - x\| \leq \varepsilon^* = \min\left\{\varepsilon, \frac{\gamma - S(x)}{L_S}\right\}, \quad (4)$$

i.e., all points x' within ε^* distance to x will also be labeled correctly as ID. A similar argument follows for OOD samples. This bound suggests that we can increase the certified radius by controlling L_S during training, e.g., through

penalizing L_S . There exists a vast literature on Lipschitz estimations for neural networks [17, 18, 28]. However, using these methods on deep models increases the computational load of training. Instead, we penalize the norm of the gradient, which can be considered as a proxy to the local Lipschitz constant.

Furthermore, for piecewise-linear networks like ReLU and LeakyReLU networks, the local Lipschitz constant is equal to the norm of the gradient. More specifically, for every point x^* , there exists an $\hat{\varepsilon} > 0$ such that the network remains linear in this region, i.e.,

$$f(x') = Ax' + b, \quad \forall x' \in \mathcal{B}(x, \hat{\varepsilon}).$$

In such a region, the local Lipschitz constant is equal to the norm of the gradient of the points [3], i.e., $L_f = \|A\|$. This further motivates directly penalizing the norm of the gradient of S . Thus, for such networks, minimizing $\|\nabla S(x)\|$ decreases the local Lipschitz constant in the neighborhood of x , which in turn could increase the maximum certified perturbation radius in (4). Note that our argument for using the norm of the gradient in practice does not require the knowledge of $\hat{\varepsilon}$. We just use the fact that such an $\hat{\varepsilon}$ exists and that in the corresponding region, the network will be linear, and as a result, minimizing the norm of the gradient is equivalent to minimizing the local Lipschitz constant for this local neighborhood.

Post-hoc Methods and Gradient Regularization: So far, we have provided a certified robustness radius against OOD misdetection using the Lipschitz constant of the network. This radius (4) is proportional to the inverse of the Lipschitz constant of the score function, and by regularizing the gradient we aim to reduce the Lipschitz constant to increase this radius. Here, we note that many of the state-of-the-art OOD detection methods can also be analyzed in this framework. For example, DICE [59] and LINE [1] use activation pruning and sparsification to remove noisy neurons. It has been shown [26] that for a given matrix, removing a row or column of it results in a reduction of the Lipschitz constant. As activation pruning is essentially a form of masking the weights of the network, methods utilizing activation pruning reduce the upper bound on the Lipschitz constant, acquiring a network with a smaller Lipschitz upper bound.

6 Experiments

We perform extensive OOD detection experiments on CIFAR [30] and ImageNet [12] benchmarks. In our comparisons, we use a wide range of CNN-based architectures. We also provide ablation analysis to demonstrate the effect of each component within our framework.

6.1 Setup

Datasets. For CIFAR experiments, we use CIFAR-10 and CIFAR-100 as ID datasets. We evaluate the models on six different OOD datasets including Tex-

*We discard the negligible case of points on the boundaries of multiple piecewise-linear spaces.

Table 1: Comparison on the CIFAR-10 benchmark. Mean AUROC and FPR95 percentages on various benchmarks. The experimental results are reported over three trials. The best mean results are bolded and the runner-up is underlined. AUROC and FPR95 are percentages.

Dataset	Method	ResNet		WRN		DenseNet	
		FPR95 ↓	AUROC ↑	FPR95 ↓	AUROC ↑	FPR95 ↓	AUROC ↑
CIFAR-10	MSP [ICLR17] [22]	58.04	90.49	51.52	90.8	52.56	91.7
	ODIN [ICLR18] [34]	34.46	91.59	35.23	89.82	24.88	94.42
	Energy Score [NeurIPS20] [36]	39.32	92.24	33.4	91.76	28.99	94.09
	ReAct [NeurIPS21] [58]	39.44	92.30	37.54	91.49	25.83	95.27
	Dice [ECCV22] [59]	42.40	91.25	34.12	91.74	26.37	94.61
	LiNe [CVPR23] [1]	45.25	90.59	36.27	90.2	14.84	96.95
	OE [ICLR18] [23]	19.92	95.71	21.12	95.55	21.76	95.8
	Energy Loss [NeurIPS20] [36]	<u>11.14</u>	<u>97.53</u>	<u>13.11</u>	<u>97.14</u>	<u>11.26</u>	<u>97.43</u>
	OpenMix [CVPR23] [79]	22.24	96.26	21.92	96	22.86	95.65
	<i>GReg</i>	7.9	97.95	7.95	98.1	7.93	98.12

Table 2: Comparison on the CIFAR-100 benchmark. Mean AUROC and FPR95 percentages on various benchmarks. The experimental results are reported over three trials. The best mean results are bolded and the runner-up is underlined.

Dataset	Method	ResNet		WRN		DenseNet	
		FPR95 ↓	AUROC ↑	FPR95 ↓	AUROC ↑	FPR95 ↓	AUROC ↑
CIFAR-100	MSP [ICLR17] [22]	74.66	80.10	80.39	75.37	85.87	68.78
	ODIN [ICLR18] [34]	64.00	84.33	67.95	79.86	69.97	81.55
	Energy Score [NeurIPS20] [36]	68.46	84.15	74.15	79.26	78.91	76.12
	ReAct [NeurIPS21] [58]	59.38	87.59	70.65	80.2	76.27	80.1
	Dice [ECCV22] [59]	77.29	80.35	72.32	79.62	69.80	80.50
	LiNe [CVPR23] [1]	72.09	82.93	67.5	80.9	35.16	88.72
	OE [ICLR18] [23]	77.32	63.89	74.22	65.70	60.52	84.21
	Energy Loss [NeurIPS20] [36]	62.77	84.05	65.62	80.18	64.37	83.86
	POEM [ICML22] [43]	61.11	82.43	62.58	79.74	59.33	79.81
	OpenMix [CVPR23] [79]	59.64	88.09	73.12	79.23	67.29	84.47
	DOS [ICLR24] [27]	<u>54.62</u>	<u>88.30</u>	45.26	<u>90.76</u>	<u>34.92</u>	<u>93.22</u>
	<i>GReg</i>	59.6	82.92	58.26	86.56	56.29	87.35
	<i>GReg+</i>	50.78	88.75	<u>48.12</u>	91.02	30.55	93.38

tures [11], SVHN [46], Places365 [78], LSUN-cropped, LSUN-resized [70], and iSUN [67]. Following [59], at test time, we use 10000 randomly selected samples from each of these 6 datasets. Following [79], instead of the recently retracted TinyImages, we use the 300K RandomImages [23] as the auxiliary outlier dataset.

For the ImageNet experiments, following [27], we consider a set of 10 *random* classes from ImageNet-1K as the ID dataset and the rest of the dataset as the auxiliary outlier dataset. For evaluation, we use 10000 randomly selected samples from Textures, Places, SUN [65], and iNaturalist [61] as OOD datasets.

Experimental Setup: For the experiments of *GReg* on CIFAR, we use a pre-trained model of the corresponding architecture and finetune for 20 epochs using SGD with cosine annealing. For the experiments of *GReg+* on CIFAR, we train a model from scratch for 50 epochs with a learning rate of 0.1 and then further train for another 10 epochs with a learning rate of 0.01, using the same optimizer. Following [36], we set $\lambda_S = 0.1$ and set $\lambda_{\nabla S} = 1$. For the ImageNet experiment, we only perform finetuning on a pre-trained DenseNet-121. See the Supplementary Material for more details.

Table 3: Comparison on the ImageNet benchmark. The experimental results are reported over three trials. The best mean results are bolded and the runner-up is underlined. AUROC and FPR95 are percentages.

Method	OOD Datasets								Average	
	iNaturalist		SUN		Places		Textures			
	FPR95 ↓ AUROC ↑	FPR95 ↓ AUROC ↑	FPR95 ↓ AUROC ↑	FPR95 ↓ AUROC ↑	FPR95 ↓ AUROC ↑	FPR95 ↓ AUROC ↑	FPR95 ↓ AUROC ↑	FPR95 ↓ AUROC ↑	FPR95 ↓ AUROC ↑	FPR95 ↓ AUROC ↑
MSP [ICLR17] [22]	62.6	87.09	64.3	86.64	62.3	86.72	78.3	75.03	66.87	83.87
ODIN [ICLR18] [34]	48.5	92.07	53.9	90.77	49.3	91.06	69.6	80.77	55.32	88.66
Energy Score [Neurips20] [36]	54	90.77	52.5	90.7	46.6	91.37	73.1	77.54	56.55	87.59
ReAct [NeurIPS21] [58]	51.88	91.52	52.48	90.81	49.56	90.92	64.61	83.77	54.63	89.25
DICE [ECCV22] [59]	32.86	93.6	41.35	92.9	45.82	91.58	65.2	78.61	46.3	89.17
LiNe [CVPR23] [1]	29.47	93.77	37.29	92.85	38.85	92.16	52.32	86.38	39.48	91.29
OE [ICLR18] [23]	34.56	92.75	54.63	89.75	54.86	89.05	76.2	75.55	55.06	86.78
Energy Loss [Neurips20] [36]	18.44	95.91	50.19	90.37	49.32	90.47	62.96	77.78	45.23	88.63
DOS [ICLR24] [27]	61.63	88.68	42.69	93.23	45.49	93.37	47.45	92.63	49.31	91.97
<i>GReg</i>	29.54	94.94	45.86	92.53	45.35	92.06	68.29	80.98	47.26	90.13
<i>GReg+</i>	24.11	95	33.98	92.31	32.89	92.7	49.32	88.22	35.08	92.06

For the CIFAR experiments, we evaluate our results on ResNet-18 [20], WRN-40 [73], and DenseNet-101 [24] to span a wide range of architectures and parameter numbers.

Sampling Details: Following the explanations in Section 4.2, we cluster the auxiliary samples into n_{ID} clusters and choose the lowest and highest scoring samples of each cluster, i.e., in each iteration, we present $2n_{ID}$ auxiliary samples for the training of that iteration. For the CIFAR experiments, we use the whole auxiliary dataset in each epoch for clustering. In other words, at each epoch, the 300K images are split into mini-batches, and clustering is performed on each mini-batch at each iteration so that no auxiliary sample is missed. For the ImageNet experiment, in each epoch, we sample 600K randomly selected images from the remaining (990) classes as the auxiliary dataset and perform the clustering on each mini-batch.

Evaluation Metrics: We report the following metrics: 1) *FPR95*: the false positive rate of the OOD samples when the true positive rate of ID samples is at 95%. 2) *AUROC*: the area under the receiver operating characteristic curve.

Comparison Methods: For comparison with post-hoc OOD methods, we compare to MSP [22], ODIN [23], Energy (score) [36], ReAct [58], DICE [59] and LiNe [1]. For the methods that use auxiliary data, we compare with OE [23], Energy (loss) [36], POEM [43], OpenMix [79] and DOS [27]. MSP uses the maximum softmax probability to detect the OOD samples. ODIN improves upon MSP by introducing noise and temperature scaling. Energy proposes a score function for OOD detection. ReAct uses activation clipping. DICE and LiNe build up on that idea by also detecting the important neurons and activation to improve OOD detection. OE and Energy train based on their respective loss functions. POEM uses posterior sampling and DOS proposes diverse sampling to improve OOD detection. OpenMix uses pseudo-samples to improve misclassification detection.

Table 4: Ablation Study of $GReg$. AUROC and FPR95 percentages on CIFAR benchmarks averaged over all OOD datasets and three runs on DenseNet.

Method	CIFAR-10		CIFAR-100	
	FPR95 ↓	AUROC ↑	FPR95 ↓	AUROC ↑
OE	21.76	95.8	60.52	84.21
OE + Grad	18.79	96.32	62.54	84.71
Energy	11.26	97.43	64.37	83.86
Energy + Grad	7.93	98.12	56.29	87.35
POEM	29.15	94.18	59.33	79.81
POEM + Grad	23.87	95.41	47.94	86.63

6.2 Results

Evaluation on CIFAR-10: We first evaluate the performance of our method on the CIFAR-10 dataset. Following the literature [1], Table 1 shows the average results over the six commonly used OOD datasets. The detailed results on individual OOD datasets can be found in the Supplementary Material. As the results on CIFAR-10 are saturated, in Table 1 we focus on the effectiveness of gradient regularization and report the rest of the methods in the Supplementary Material. It can be seen that $GReg$ outperforms all the methods in both of the reported metrics. Specifically, $GReg$ decreases the FPR95 of [36] by 3.9%.

Evaluation on CIFAR-100: For the results of this part, we use the same experimental setup and only change the ID dataset to CIFAR-100. The average results of this experiment set are presented in Table 2 and the individual results can be found in the Supplementary Material. In this setup, $GReg$ provides competitive results on the ResNet and WRN architectures with state-of-the-art methods that do not employ sampling. Moreover, coupled with sampling, $GReg+$ outperforms all current state-of-the-art methods on almost all metrics of the different architectures. On average, $GReg+$ outperforms the FPR95 metric of DOS [27] by 1.8%, which showcases the effectiveness of energy-based sampling. Results on each OOD dataset can be found in the Supplementary Material.

Evaluation on ImageNet: For this experiment, we randomly select 10 classes of Imagenet-1K as ID and use the rest as the auxiliary data. The results of this large-scale experiment are presented in Table 3. We can see that in terms of FPR95, $GReg+$ outperforms DOS and LINE by large margins of 14.23% and 3.68%, respectively, while achieving state-of-the-art AUROC performance. The energy-based sampling approach enables our method to acquire more informative samples compared to the sampling method in [27], resulting in improved FPR95.

6.3 Ablation Study

In this section, we perform extensive experiments to show the effectiveness of components within our framework, individually. That is, we show the effect of regularizing the gradient on other methods and then compare our proposed sampling method with the state-of-the-art. We also study the effect of our OOD schemes on the ID accuracy of the models.

Table 5: Ablation Study on Sampling. Comparison among sampling methods on CIFAR-100 on DenseNet averaged over six OOD datasets.

Method	CIFAR-100	
	FPR95 ↓	AUROC ↑
POEM	59.33	79.81
DOS	34.92	93.22
Energy only clustering	31.77	93.20
<i>GReg+</i>	30.55	93.38

Table 6: Ablation Study on ID Accuracy. Accuracy comparison on CIFAR-10 on two architectures.

Method	CIFAR-10	
	WRN	DenseNet
Clean	94.84	94.03
Energy loss	88.15	91.82
<i>GReg</i>	89.55	92.14
DOS	78.90	87.06
<i>GReg+</i>	83.06	89.89

Gradient Regularization: Due to the generality of the idea of gradient regularization, our method can be used in conjunction with many existing methods. To leverage the gradient regularization idea, one needs to define a suitable loss function based on the norm of the gradient of the score, to promote a smoother score manifold. To show the effectiveness of *GReg* on different OOD training schemes, we choose three well-known OOD detection methods, Energy, OE, and POEM, and train a DenseNet with and without gradient regularization on both CIFAR benchmarks. Table 4 shows the overall improvement of all these methods with gradient regularization, in both the FPR95 and AUROC, supporting our claim that coupling *GReg* with other methods could improve their performance.

Sampling Methods: To study the effect of various sampling methods, we compare POEM and DOS to our method with and without gradient regularization. We denote *GReg+* without gradient regularization as Energy-only-clustering. Table 5 compares different methods that utilize some form of sampling and shows the impact of energy-based clustering. On top of that, the comparison between Energy-only-clustering and *GReg+* further encourages gradient regularization.

ID Accuracy: To study the effect of gradient regularization on accuracy, we compare the ID accuracy of our methods to the most relevant state-of-the-art on WRN and DenseNet architectures. Table 6 shows that gradient regularization (*GReg*) improves the ID accuracy of energy on both architectures and *GReg+* also has superior ID accuracy in comparison with DOS.

7 Conclusions

In this paper, we presented the idea of gradient regularization of the score function in OOD detection. We also developed an energy-based sampling algorithm to improve sampling quality when the auxiliary dataset is large. We demonstrate that *GReg* exhibits superior performance compared to methods that do not require clustering, and *GReg+* outperforms all state-of-the-art methods. Furthermore, our experiments show that regularizing the gradient could result in robust models capable of extracting local information from the ID and OOD data. We also provide detailed analytical analysis and ablation studies to support our method.

Acknowledgements

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-23-2-0008. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

1. Ahn, Y.H., Park, G.M., Kim, S.T.: Line: Out-of-distribution detection by leveraging important neurons. arXiv preprint arXiv:2303.13995 (2023)
2. Bafghi, R.A., Gurari, D.: A new dataset based on images taken by blind people for testing the robustness of image classification models trained for imagenet categories. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16261–16270 (June 2023)
3. Bhowmick, A., D’Souza, M., Raghavan, G.S.: Lipbab: Computing exact lipschitz constant of relu networks. In: Artificial Neural Networks and Machine Learning–ICANN 2021: 30th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 14–17, 2021, Proceedings, Part IV 30. pp. 151–162. Springer (2021)
4. Cao, C., Zhong, Z., Zhou, Z., Liu, Y., Liu, T., Han, B.: Envisioning outlier exposure by large language models for out-of-distribution detection. arXiv preprint arXiv:2406.00806 (2024)
5. Chan, A., Tay, Y., Ong, Y.S., Fu, J.: Jacobian adversarially regularized networks for robustness. arXiv preprint arXiv:1912.10185 (2019)
6. Chen, G., Qiao, L., Shi, Y., Peng, P., Li, J., Huang, T., Pu, S., Tian, Y.: Learning open set network with discriminative reciprocal points. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16. pp. 507–522. Springer (2020)
7. Chen, J., Li, Y., Wu, X., Liang, Y., Jha, S.: Atom: Robustifying out-of-distribution detection using outlier mining. In: Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part III 21. pp. 430–445. Springer (2021)
8. Chen, S., Huang, L.K., Schwarz, J.R., Du, Y., Wei, Y.: Secure out-of-distribution task generalization with energy-based models. *Advances in Neural Information Processing Systems* **36** (2024)
9. Choi, H., Jeong, H., Choi, J.Y.: Balanced energy regularization loss for out-of-distribution detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15691–15700 (2023)
10. Choi, S., Lee, H., Lee, H., Lee, M.: Projection regret: Reducing background bias for novelty detection via diffusion models. *Advances in Neural Information Processing Systems* **36**, 19230–19245 (2023)
11. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3606–3613 (2014)

12. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
13. Dhamija, A.R., Günther, M., Boulton, T.: Reducing network agnostophobia. *Advances in Neural Information Processing Systems* **31** (2018)
14. Du, X., Fang, Z., Diakonikolas, I., Li, Y.: How does unlabeled data provably help out-of-distribution detection? *arXiv preprint arXiv:2402.03502* (2024)
15. Du, X., Wang, Z., Cai, M., Li, Y.: Vos: Learning what you don’t know by virtual outlier synthesis. *arXiv preprint arXiv:2202.01197* (2022)
16. Fan, K., Liu, T., Qiu, X., Wang, Y., Huai, L., Shanguan, Z., Gou, S., Liu, F., Fu, Y., Fu, Y., Jiang, X.: Test-time linear out-of-distribution detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 23752–23761 (June 2024)
17. Fazlyab, M., Entesari, T., Roy, A., Chellappa, R.: Certified robustness via dynamic margin maximization and improved lipschitz regularization (2023)
18. Fazlyab, M., Robey, A., Hassani, H., Morari, M., Pappas, G.: Efficient and accurate estimation of lipschitz constants for deep neural networks. In: *Advances in Neural Information Processing Systems*. pp. 11427–11438 (2019)
19. Ge, Z., Demyanov, S., Chen, Z., Garnavi, R.: Generative openmax for multi-class open set classification. *arXiv preprint arXiv:1707.07418* (2017)
20. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
21. Hein, M., Andriushchenko, M., Bitterwolf, J.: Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 41–50 (2019)
22. Hendrycks, D., Gimpel, K.: A baseline for detecting misclassified and out-of-distribution examples in neural networks. In: *International Conference on Learning Representations* (2017), <https://openreview.net/forum?id=Hkg4TI9x1>
23. Hendrycks, D., Mazeika, M., Dietterich, T.: Deep anomaly detection with outlier exposure. In: *International Conference on Learning Representations* (2019), <https://openreview.net/forum?id=HyxCxhRcY7>
24. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4700–4708 (2017)
25. Huang, R., Geng, A., Li, Y.: On the importance of gradients for detecting distributional shifts in the wild. *Advances in Neural Information Processing Systems* **34**, 677–689 (2021)
26. Huang, Y., Zhang, H., Shi, Y., Kolter, J.Z., Anandkumar, A.: Training certifiably robust neural networks with efficient local lipschitz bounds. *Advances in Neural Information Processing Systems* **34**, 22745–22757 (2021)
27. Jiang, W., Cheng, H., Chen, M., Wang, C., Wei, H.: Dos: Diverse outlier sampling for out-of-distribution detection. *arXiv preprint arXiv:2306.02031* (2023)
28. Jordan, M., Dimakis, A.G.: Exactly computing the local lipschitz constant of relu networks. *Advances in Neural Information Processing Systems* **33**, 7344–7353 (2020)
29. Kong, S., Ramanan, D.: Opengan: Open-set recognition via open data generation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 813–822 (2021)

30. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
31. Li, J., Chen, P., He, Z., Yu, S., Liu, S., Jia, J.: Rethinking out-of-distribution (ood) detection: Masked image modeling is all you need. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11578–11589 (2023)
32. Li, T., Qiao, F., Ma, M., Peng, X.: Are data-driven explanations robust against out-of-distribution data? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3821–3831 (2023)
33. Li, Y., Vasconcelos, N.: Background data resampling for outlier-aware classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13218–13227 (2020)
34. Liang, S., Li, Y., Srikant, R.: Enhancing the reliability of out-of-distribution image detection in neural networks. arXiv preprint arXiv:1706.02690 (2017)
35. Lin, Z., Roy, S.D., Li, Y.: Mood: Multi-level out-of-distribution detection. In: Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition. pp. 15313–15323 (2021)
36. Liu, W., Wang, X., Owens, J., Li, Y.: Energy-based out-of-distribution detection. *Advances in neural information processing systems* **33**, 21464–21475 (2020)
37. Liu, X., Lochman, Y., Zach, C.: Gen: Pushing the limits of softmax-based out-of-distribution detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 23946–23955 (2023)
38. Lloyd, S.: Least squares quantization in pcm. *IEEE transactions on information theory* **28**(2), 129–137 (1982)
39. Lo, S.Y., Oza, P., Patel, V.M.: Adversarially robust one-class novelty detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(4), 4167–4179 (2022)
40. Losch, M., Stutz, D., Schiele, B., Fritz, M.: Certified robust models with slack control and large lipschitz constants. arXiv preprint arXiv:2309.06166 (2023)
41. Lu, H., Gong, D., Wang, S., Xue, J., Yao, L., Moore, K.: Learning with mixture of prototypes for out-of-distribution detection. arXiv preprint arXiv:2402.02653 (2024)
42. Ming, Y., Cai, Z., Gu, J., Sun, Y., Li, W., Li, Y.: Delving into out-of-distribution detection with vision-language representations. *Advances in Neural Information Processing Systems* **35**, 35087–35102 (2022)
43. Ming, Y., Fan, Y., Li, Y.: Poem: Out-of-distribution detection with posterior sampling. In: International Conference on Machine Learning. pp. 15650–15665. PMLR (2022)
44. Moon, W., Park, J., Seong, H.S., Cho, C.H., Heo, J.P.: Difficulty-aware simulator for open set recognition. In: European Conference on Computer Vision. pp. 365–381. Springer (2022)
45. Neal, L., Olson, M., Fern, X., Wong, W.K., Li, F.: Open set learning with counterfactual images. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 613–628 (2018)
46. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning (2011)
47. Nguyen, A., Yosinski, J., Clune, J.: Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 427–436 (2015)

48. Noh, S., Jeong, D., Lee, J.H.: Simple and effective out-of-distribution detection via cosine-based softmax loss. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16560–16569 (2023)
49. Olber, B., Radlak, K., Popowicz, A., Szczepankiewicz, M., Chachula, K.: Detection of out-of-distribution samples using binary neuron activation patterns. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3378–3387 (2023)
50. Oza, P., Patel, V.M.: C2ae: Class conditioned auto-encoder for open-set recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2307–2316 (2019)
51. Park, J., Jung, Y.G., Teoh, A.B.J.: Nearest neighbor guidance for out-of-distribution detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1686–1695 (2023)
52. Park, J., Park, H., Jeong, E., Teoh, A.B.J.: Understanding open-set recognition by jacobian norm of representation. arXiv preprint arXiv:2209.11436 (2022)
53. Peng, B., Luo, Y., Zhang, Y., Li, Y., Fang, Z.: Conjnorn: Tractable density estimation for out-of-distribution detection. arXiv preprint arXiv:2402.17888 (2024)
54. Perini, L., Davis, J.: Unsupervised anomaly detection with rejection. *Advances in Neural Information Processing Systems* **36** (2024)
55. Safaei, B., Vibashan, V., de Melo, C.M., Hu, S., Patel, V.M.: Open-set automatic target recognition. In: ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 1–5. IEEE (2023)
56. Shapley, L.S., et al.: A value for n-person games (1953)
57. Srinivas, S., Matoba, K., Lakkaraju, H., Fleuret, F.: Efficient training of low-curvature neural networks. *Advances in Neural Information Processing Systems* **35**, 25951–25964 (2022)
58. Sun, Y., Guo, C., Li, Y.: React: Out-of-distribution detection with rectified activations. *Advances in Neural Information Processing Systems* **34**, 144–157 (2021)
59. Sun, Y., Li, Y.: Dice: Leveraging sparsification for out-of-distribution detection. In: European Conference on Computer Vision. pp. 691–708. Springer (2022)
60. Tang, K., Hou, C., Peng, W., Chen, R., Zhu, P., Wang, W., Tian, Z.: Cores: Convolutional response-based score for out-of-distribution detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10916–10925 (2024)
61. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8769–8778 (2018)
62. Wang, L., Han, M., Li, X., Zhang, N., Cheng, H.: Review of classification methods on unbalanced data sets. *IEEE Access* **9**, 64606–64628 (2021)
63. Wang, Q., Fang, Z., Zhang, Y., Liu, F., Li, Y., Han, B.: Learning to augment distributions for out-of-distribution detection. *Advances in Neural Information Processing Systems* **36** (2024)
64. Wang, Q., Ye, J., Liu, F., Dai, Q., Kalander, M., Liu, T., Hao, J., Han, B.: Out-of-distribution detection with implicit outlier transformation. arXiv preprint arXiv:2303.05033 (2023)
65. Xiao, J., Hays, J., Ehinger, K.A., Oliva, A., Torralba, A.: Sun database: Large-scale scene recognition from abbey to zoo. In: 2010 IEEE computer society conference on computer vision and pattern recognition. pp. 3485–3492. IEEE (2010)

66. Xu, M., Lian, Z., Liu, B., Tao, J.: Vra: variational rectified activation for out-of-distribution detection. *Advances in Neural Information Processing Systems* **36**, 28941–28959 (2023)
67. Xu, P., Ehinger, K.A., Zhang, Y., Finkelstein, A., Kulkarni, S.R., Xiao, J.: Turkergaze: Crowdsourcing saliency with webcam based eye tracking. *arXiv preprint arXiv:1504.06755* (2015)
68. Yang, J., Wang, H., Feng, L., Yan, X., Zheng, H., Zhang, W., Liu, Z.: Semantically coherent out-of-distribution detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8301–8309 (2021)
69. Yang, Y., Lee, K., Dariush, B., Cao, Y., Lo, S.Y.: Follow the rules: Reasoning for video anomaly detection with large language models. In: *European Conference on Computer Vision* (2024)
70. Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., Xiao, J.: Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365* (2015)
71. Yu, Q., Aizawa, K.: Unsupervised out-of-distribution detection by maximum classifier discrepancy. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 9518–9526 (2019)
72. Yuan, Y., He, R., Dong, Y., Han, Z., Yin, Y.: Discriminability-driven channel selection for out-of-distribution detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 26171–26180 (2024)
73. Zagoruyko, S., Komodakis, N.: Wide residual networks. *arXiv preprint arXiv:1605.07146* (2016)
74. Zhang, J., Inkawhich, N., Linderman, R., Chen, Y., Li, H.: Mixture outlier exposure: Towards out-of-distribution detection in fine-grained environments. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 5531–5540 (2023)
75. Zhang, J., Yang, J., Wang, P., Wang, H., Lin, Y., Zhang, H., Sun, Y., Du, X., Zhou, K., Zhang, W., Li, Y., Liu, Z., Chen, Y., Li, H.: Openood v1.5: Enhanced benchmark for out-of-distribution detection. *arXiv:2306.09301* (2023)
76. Zhang, Z., Xiang, X.: Decoupling maxlogit for out-of-distribution detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3388–3397 (June 2023)
77. Zheng, H., Wang, Q., Fang, Z., Xia, X., Liu, F., Liu, T., Han, B.: Out-of-distribution detection learning with unreliable out-of-distribution sources. *Advances in Neural Information Processing Systems* **36**, 72110–72123 (2023)
78. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence* **40**(6), 1452–1464 (2017)
79. Zhu, F., Cheng, Z., Zhang, X.Y., Liu, C.L.: Openmix: Exploring outlier samples for misclassification detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 12074–12083 (June 2023)
80. Zhu, J., Geng, Y., Yao, J., Liu, T., Niu, G., Sugiyama, M., Han, B.: Diversified outlier exposure for out-of-distribution detection via informative extrapolation. *Advances in Neural Information Processing Systems* **36** (2024)