

---

# The Best of Both Worlds: On the Dilemma of Out-of-distribution Detection

---

Qingyang Zhang<sup>1\*</sup>, Qiuxuan Feng<sup>1</sup>, Joey Tianyi Zhou<sup>2</sup>, Yatao Bian<sup>3</sup>,  
Qinghua Hu<sup>1</sup>, Changqing Zhang<sup>1†</sup>  
College of Intelligence and Computing, Tianjin University<sup>1</sup>  
A\*STAR<sup>2</sup>, Tencent AI Lab<sup>3</sup>

## Abstract

Out-of-distribution (OOD) detection is essential for model trustworthiness which aims to sensitively identify semantic OOD samples and robustly generalize for covariate-shifted OOD samples. However, we discover that the superior OOD detection performance of state-of-the-art methods is achieved by secretly sacrificing the OOD generalization ability. Specifically, the classification accuracy of these models could deteriorate dramatically when they encounter even minor noise. This phenomenon contradicts the goal of model trustworthiness and severely restricts their applicability in real-world scenarios. What is the hidden reason behind such a limitation? In this work, we theoretically demystify the “*sensitive-robust*” dilemma that lies in many existing OOD detection methods. Consequently, a theory-inspired algorithm is induced to overcome such a dilemma. By decoupling the uncertainty learning objective from a Bayesian perspective, the conflict between OOD detection and OOD generalization is naturally harmonized and a dual-optimal performance could be expected. Empirical studies show that our method achieves superior performance on standard benchmarks. To our best knowledge, this work is the first principled OOD detection method that achieves state-of-the-art OOD detection performance without compromising OOD generalization ability. Our code is available at <https://github.com/QingyangZhang/DUL>.

## 1 Introduction

Endowing machine learning models with out-of-distribution (OOD) detection and OOD generalization ability are both essential for their deployment in the open world [1, 2, 3]. We borrow an example of autonomous driving from [4] to demonstrate the motivation of these two tasks. Given a machine learning model trained on in-distribution (ID) data (top image in Fig. 1 (a)), OOD detection aims to sensitively perceive uncertainty arising upon outliers that do not belong to any known classes of training data [5] (bottom right image in Fig. 1 (a)). While OOD generalization expects machine learning models to be robust in the presence of unexpected noise or corruption, e.g., rainy or snowy weather (bottom left image in Fig. 1 (a)). In this paper, we reveal that many previous methods pursue OOD detection performance at a secret cost of sacrificing OOD generalization ability. To make things worse, we observe that some SOTA OOD detection methods may result in a catastrophic collapse in classification performance ( $\sim 15\%$  accuracy degradation) when encountering even slight noise. One pioneering work [4] makes a trade-off between OOD detection and OOD generalization, but the relationship between these two tasks is still largely unexplored. The learning objectives of these two tasks are seemingly conflicting at first glance. OOD detection encourages sensitive uncertainty awareness (highly uncertain prediction) on unseen data, while generalization expects the prediction to be confident and robust under unforeseeable distributional shifts. Previous work in OOD detection

---

\*Work done during an internship at Tencent AI Lab.

†Correspondence to Changqing Zhang <[zhangchangqing@tju.edu.cn](mailto:zhangchangqing@tju.edu.cn)>

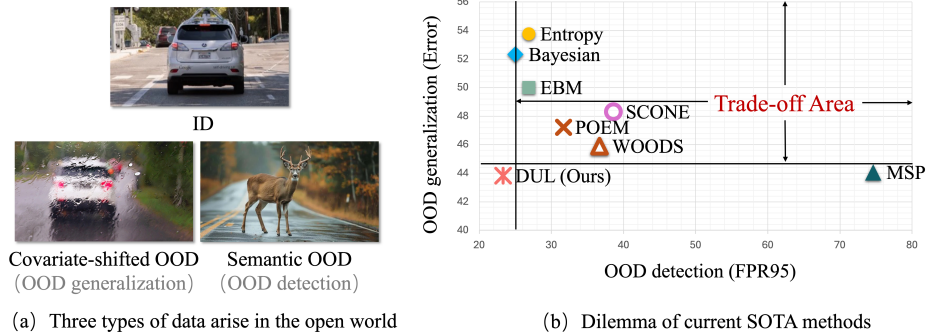


Figure 1: **(a)**: Models trained on in-distribution (ID) data inevitably encounter distributional shifts during their deployment. OOD generalization expects the model to correctly classify covariate-shifted data that undergoes noise or corruption due to environmental issues. OOD detection aims to identify samples that do not belong to any known classes for trustworthiness consideration. **(b)**: Limitations of current advanced OOD detection methods. We consider 8 representative OOD detection methods including the baseline method MSP [6] (without any OOD detection regularization), Entropy [7], EBM [8], Bayesian [9], SOTA OOD detection methods WOODS [10], POEM [11], recent advanced SCONE [4] which aims to seek for a good trade-off and the proposed DUL. All these methods exhibit a degraded generalization ability compared to baseline method MSP and lie in a trade-off area except our DUL. The goal of this paper is to understand and mitigate this phenomenon.

research area [4] characterizes the relationship between OOD detection and OOD generalization as a trade-off and thus striking for a balanced performance. However, this trade-off significantly limits the employment of current state-of-the-art OOD detection methods. Naturally, one might require the model to be aware of the OOD input for ensuring safety, but certainly does not expect to sacrifice the generalization ability, not to mention that the catastrophically collapsed classification performance under noise or corruption.

In this work, we first uncover the potential reason behind this limitation by characterizing the generalization error lower bound of previous OOD detection methods, which is referred to *sensitive-robust dilemma*. To overcome the dilemma, we devise a novel Decoupled Uncertainty Learning (DUL) framework for dual-optimal performance. The decoupled uncertainties are separately responsible for characterizing semantic OOD (detection) and covariate-shifted OOD (generalization). Thanks to the decoupled uncertainty learning objective, dual-optimal OOD detection and OOD generalization performance could be expected. Our emphasis lies on a particular category of OOD detection methods in the classification task, including max softmax probability (MSP) based model [6], energy-based model (EBM) [8] and Bayesian methods [9]. This selection offers two-fold advantages. First, MSP, EBM and Bayesian detectors encompass major OOD detection advances in classification task [5]. Second, numerous OOD detection works in diverse learning tasks (classification, object detection [12], time-series prediction and image segmentation) are all roughly related to classification [13]. The contributions of this paper are summarized as follows:

- This paper reveals that existing SOTA OOD detection methods may suffer from catastrophic degradation in terms of OOD generalization. That is, their superior OOD detection ability is achieved by (secretly) sacrificing OOD generalization ability. We theoretically demystify the sensitive-robust dilemma in learning objectives as the main reason behind such a limitation.
- In contrast to previous works that characterize OOD detection and generalization as conflictive learning tasks and thus implying an inevitable trade-off, we propose a novel learning framework termed Decoupled Uncertainty Learning (DUL) to successfully break through the limitation beyond a simple trade-off. Our DUL substantially harmonizes the conflict between OOD detection and OOD generalization, which achieves the best OOD detection performance without sacrificing the OOD generalization ability.
- We conduct extensive experiments on standard benchmarks to validate our findings. Our DUL achieves dual-optimal OOD detection and OOD generalization performance. To our best knowledge, DUL is the first method that gains state-of-the-art OOD detection performance without sacrificing OOD generalization ability.

## 2 Related works

**OOD detection** aims to indicate whether the input arises from unknown classes that are not present in training data, which is essential for model trustworthiness. In the classification task, the majority of advanced OOD detection methods include MSP detectors which characterize samples with lower max softmax probability as OOD [6, 14, 15, 7, 16]. EBM detectors identify high energy samples as OOD and frequently establish better performance than MSP detectors [8, 11, 17, 4], and various other types OOD detection methods such as distance-based detectors [18], non-parametric KNN-based detectors [19] which also show promises. According to the training paradigm, OOD detection methods can be split into auxiliary OOD-free and auxiliary OOD-required methods. Auxiliary OOD-free methods directly use the model pre-trained on ID data only for OOD detection. Another line of methods assumes that some OOD data is accessible during training and incorporates auxiliary outlier datasets (collected from websites or other datasets) for further enhancing OOD detection performance. By exposing the model to some semantic OOD during training, auxiliary OOD-required methods frequently outperform auxiliary OOD-free methods on commonly-used benchmarks [20, 5].

**OOD generalization** expects the model to be robust under unforeseeable noise or corruption [21, 22, 23, 24, 25]. Basically, OOD generalization expects invariant and confident prediction on OOD data. Examples include classic domain adaption (DA) methods which encourage the model’s behavior to be invariant across different distributions [21, 26, 27]. Besides, test-time adaption (TTA) directly encourages confident predictions on OOD data by minimizing predictive entropy [28, 29, 30]. However, as we will show later, confident prediction and invariance are seemingly conflictive to OOD detection purpose and further imply an unavoidable trade-off. The most related work to our paper is SCONE [4], which strikes to keep a balance between OOD detection and generalization performance. We argue that such a trade-off is not necessary and the conflict can be elegantly eliminated.

**Uncertainty estimation in Bayesian framework.** In the Bayesian framework, predictive uncertainty can be regarded as an indicator of whether the input sample is prone to be OOD. Since OOD samples are unseen during training and thus should be of higher uncertainty than ID. The overall predictive uncertainty of a classification model can be decomposed into three factors according to their source, including data (aleatoric) uncertainty (AU), distributional uncertainty (DU), and model (epistemic) uncertainty (EU) [31, 9]. AU measures the natural complexity of the data (e.g., class overlap, label noise) and EU results from the difficulty of estimating the model parameters with finite training data. DU arises due to the mismatch between the distributions of test and training data. A line of classic measurement can be used to capture various types of uncertainty including entropy, mutual information, and differential entropy [9].

## 3 Preliminaries

We consider  $K$ -class classification task with classifier  $f : \mathcal{X} \rightarrow \mathbb{R}^K$  parameterized by  $\theta$ , where  $\mathcal{X}$  is the input space and  $\mathcal{Y} = \{1, 2, \dots, K\}$  denotes the target space. The model output  $f_\theta(x)$  is considered as logits. The  $k$ -th element in logits is denoted as  $f_k(x)$  indicates the confidence of predicting  $x$  to class  $k$ . The predicted distribution  $F(x)$  is obtained by normalizing  $f(x)$  with the softmax function. We first formalize all possible distributions that the model might encounter.

- In-distribution  $P_{\mathcal{X}\mathcal{Y}}^{\text{ID}}$  which denotes the distribution of labeled training data.
- Covariate-shifted OOD  $P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}$  which is relevant to OOD generalization.  $P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}$  is of the same label space with ID. However, its marginal distribution  $P_{\mathcal{X}}^{\text{COV}}$  encounters shifts due to unexpected noise or corruption.
- Semantic OOD  $P_{\mathcal{X}\mathcal{Y}'}^{\text{SEM}}$  is the distribution of data that do not belong to any known class. Its label space has no overlap with the known ID label space, i.e.,  $\mathcal{Y}' \cap \mathcal{Y} = \emptyset$ .

In the following paper, we omit the subscript for simplicity. The goal of OOD detection is to build a detector  $G : \mathcal{X} \rightarrow [\text{IN}, \text{OUT}]$  to decide whether an input  $x$  is semantic OOD data or not through a thresholding function  $G$  deduced from classifier  $f$

$$G_\gamma(x) = \begin{cases} \text{IN} & g_f(x) \leq \gamma \\ \text{OUT} & g_f(x) > \gamma \end{cases}, \quad (1)$$

where  $\gamma$  is the threshold.  $g_f$  is an OOD scoring function deduced from  $f$ , which is expected to assign a higher value to OOD than ID. For example, in MSP detectors,  $g_f(x) = -\max_k F(x)$  where  $F(x)$  is the predicted softmax probability (negative max softmax probability). In EBM detectors,  $g_f$  is realized by the energy function  $E(x; f) := -\log \sum_{i=1}^K e^{f_i(x)}$  and the semantic input of OOD should be of high energy [8]. Since it is difficult to foresee  $P^{\text{SEM}}$  one will encounter, a board line of OOD detection works [7, 8, 9, 11, 17, 12, 32, 33] regularize the model on some auxiliary OOD data  $P_{\text{train}}^{\text{SEM}}$  during training (e.g., data from the web or other datasets), and expect the model can learn useful heuristic to handle unknown test-time OOD  $P_{\text{test}}^{\text{SEM}}$ . The learning objective is shown as follows

$$\min_{\theta} \mathbb{E}_{(x,y) \sim P^{\text{ID}}} [\mathcal{L}_{\text{CE}}(f(x), y)] + \lambda \mathbb{E}_{\tilde{x} \sim P_{\text{train}}^{\text{SEM}}} [\mathcal{L}_{\text{reg}} f(\tilde{x})], \quad (2)$$

where  $\mathcal{L}_{\text{CE}}$  is the standard cross entropy loss for the original classification task.  $\mathcal{L}_{\text{reg}}$  is the OOD detection regularization term depending on the detector used, which generally encourages a high uncertainty on  $P_{\text{train}}^{\text{SEM}}$ . For example,  $\mathcal{L}_{\text{reg}}$  is set to cross entropy between  $F(x)$  and the uniform distribution for MSP detector [7]. In EBM detectors [8],  $\mathcal{L}_{\text{reg}}$  is realized as a margin ranking loss to explicitly encourage a large energy gap between ID and semantic OOD. In this paper, we are interested in this setting for the following reasons:

- 1) In contrast to labeled data in supervised learning literature, auxiliary OOD data can be unlabeled and easy-to-collected in practice [11].
- 2) Most SOTA methods involve auxiliary outliers [5, 20] for superior performance.
- 3) Even under some strict assumptions that  $P_{\text{train}}^{\text{SEM}}$  is unavailable, recent works utilize GAN [15], diffusion model [34] or sampling strategy [12] to generate "virtual" outliers for training.

Thus we believe this setting is promising and the cost of auxiliary outliers is minor given the importance of ensuring model trustworthiness. At test-time, the model is evaluated in terms of

- ID accuracy (ID-Acc  $\uparrow$ ) which measures the model’s performance on  $P^{\text{ID}}$ ,
- OOD accuracy (OOD-Acc  $\uparrow$ ) measures the OOD generalization ability on  $P^{\text{COV}}$ ,
- False positive rate at 95% true positive rate (FPR95 $\downarrow$ ) :=  $\mathbb{E}_{x \sim P_{\text{test}}^{\text{SEM}}} (\mathbb{I}(G_{\gamma}(x) = \text{IN}))$  measures the OOD detection ability, where  $\gamma$  is chosen when true positive rate (TPR) is 95%.  $\mathbb{I}$  is the indicator function. In OOD detection, ID samples are considered as negative.

It is worth noting that in the standard OOD detection setting [11, 4], the test OOD data should not have any overlapped classes or samples with training-time auxiliary OOD data  $P_{\text{train}}^{\text{SEM}}$ . Let  $\mathcal{Y}_{\text{test}}^{\text{SEM}}$  and  $\mathcal{Y}_{\text{train}}^{\text{SEM}}$  be the label space of  $P_{\text{test}}^{\text{SEM}}$  and  $P_{\text{train}}^{\text{SEM}}$  respectively, we have  $\mathcal{Y}_{\text{test}}^{\text{SEM}} \cap \mathcal{Y}_{\text{train}}^{\text{SEM}} = \emptyset$ . Otherwise, OOD detection would be a trivial problem.

## 4 Sensitive-robust Dilemma of Out-of-distribution Detection

In this section, we detail the limitation of current OOD detection methods: their OOD detection performance is achieved at the cost of generalization ability. This limitation implies the potential risk of SOTA OOD detection methods and underscores the urgent need for a better solution. Firstly, we re-examine representative OOD detection methods of six different types, including 1) baseline model MSP that is trained without any OOD detection regularization [6], 2) entropy-regularization (Entropy) that encourages high predictive entropy on OOD [7], which is devised for MSP detectors, 3) energy-regularization for EBM detectors that enforces the output with high energy score for OOD input [8], 4) Bayesian uncertainty learning that encourages high overall uncertainty on OOD [9], 5) state-of-the-art OOD detection methods WOODS [10] and POEM [11] 6) the most related SCONE [4] that seeks for a trade-off between OOD detection and generalization performance.

**Limitation of current OOD detection methods.** In Fig. 1 (b), we investigate current OOD detection methods in terms of OOD classification error and FPR95. The expected classifier should yield both low OOD classification error and FPR95. As it is observed, despite the superior OOD detection performance, all above methods significantly underperform the baseline MSP in terms of OOD generalization. By contrast, our method (DUL) successfully overcomes the limitation.

**Theoretical justification.** Toward understanding the limitation, we provide theoretical analysis for two types of most popular OOD detection methods, i.e., MSP and EBM detectors. Our analysis

identifies the “*sensitive-robust*” dilemma as the main reason behind such a limitation. The roadmap of our analysis is: (1) inspired by transfer learning theory, we first reveal that OOD detection regularization applied on semantic OOD may also affect the behavior of model on covariate-shifted OOD; (2) then we demonstrate why MSP detectors suffer from poor generalization by characterizing its generalization error bound; (3) we further identify that EBM methods [8] suffer from a similar drawback when incorporating with gradient-based optimization. First of all, we recap the definition of disparity discrepancy in transfer learning theory [35, 27].

**Definition 1** (Disparity with Total Variation Distance). *Given two hypotheses  $f', f \in \mathcal{F}$  and distribution  $P$ , we define the Disparity with Total Variation Distance between them as*

$$\text{disp}_P(f', f) = \mathbb{E}_P[TV(F_f || F_{f'})], \quad (3)$$

where  $F_{f'}, F_f$  are the class distributions predicted by  $f', f$  respectively.  $TV(\cdot || \cdot)$  is the total variation distance, i.e.,  $TV(F_f || F_{f'}) = \frac{1}{2} \sum_{k=1}^K ||F_{f,k} - F_{f',k}||$ .

**Definition 2** (Disparity Discrepancy with Total Variation Distance, DD with TVD). *Given a hypothesis space  $\mathcal{F}$  and two distributions  $P, Q$ , the Disparity Discrepancy with Total Variation Distance (DD with TVD) is defined as*

$$d_{\mathcal{F}}(P, Q) := \sup_{f', f \in \mathcal{F}} (\text{disp}_P(f', f) - \text{disp}_Q(f', f)). \quad (4)$$

Disparity discrepancy (DD) measures the “distance” between two distributions  $P, Q$  which considers the hypothesis space. DD is one of the most fundamental conceptions in transfer learning theory which constrains the behavior of hypothesis in  $\mathcal{F}$  should not be dissimilar substantially on different distributions  $P$  and  $Q$ .<sup>3</sup> If the DD between semantic OOD and covariate-shifted OOD is limited, one can suppose that OOD detection regularization applied to semantic OOD samples will also influence the model’s behavior on covariate-shifted OOD. Thus encouraging high uncertainty on semantic OOD may also result in highly uncertain prediction on covariate-shifted OOD, which is potentially harmful to generalization ability. We first formalize this intuition for MSP detectors.

**Theorem 1** (Sensitive-robust dilemma). *Let  $\mathcal{P}^{\text{COV}}, P_{\text{test}}^{\text{SEM}}$  be the covariate-shifted OOD and semantic OOD distribution.  $\text{GError}_{P^{\text{COV}}}(f)$  denotes standard cross entropy loss taking expectation on  $P^{\text{COV}}$ , i.e., generalization error. Then we have*

$$\underbrace{\text{GError}_{P^{\text{COV}}}(f)}_{\text{OOD generalization error}} \geq C - \sqrt{\frac{1}{8\kappa^2} \mathbb{E}_{P_{\text{test}}^{\text{SEM}}} [ \underbrace{\mathcal{L}_{\text{reg}}(f)}_{\text{OOD detection loss}} - \log K ]^{\frac{1}{2}}} - \frac{1}{2\kappa} d_{\mathcal{F}}(\mathcal{P}^{\text{COV}}, P_{\text{test}}^{\text{SEM}}), \quad (5)$$

where  $\mathcal{L}_{\text{reg}}$  is the OOD detection loss devised for MSP detectors defined in [7], i.e., cross-entropy between predicted distribution  $F(x)$  and uniform distribution.  $d_{\mathcal{F}}(P^{\text{COV}}, P_{\text{test}}^{\text{SEM}})$  is DD with TVD that measures the dissimilarity of covariate-shifted OOD and semantic OOD.  $C$  and  $\kappa$  are both some constants depending on hypothesis space  $\mathcal{F}$ ,  $P^{\text{COV}}$  and  $P_{\text{test}}^{\text{SEM}}$ .

The proof is deferred in Appendix A. Theorem 1 demonstrates that for MSP detectors, the OOD detection objective conflicts with OOD generalization. The model’s generalization error lower bound is negatively correlated with OOD detection loss that the model tries to minimize. Thus given a limited  $d_{\mathcal{F}}$ , pursuing low OOD detection loss on  $P_{\text{test}}^{\text{SEM}}$  will also inevitably result in highly uncertain prediction on  $P^{\text{COV}}$ . It is worth noting that such an interpretative theorem is applicable for all MSP-based OOD detectors no matter whether the model involves  $P_{\text{train}}^{\text{SEM}}$  during training or not. Since the inherent motivation of OOD detection methods lies in minimizing the OOD detection loss in  $P_{\text{test}}^{\text{SEM}}$ , regardless of the training strategies used.

**Why a limited  $d_{\mathcal{F}}(P^{\text{COV}}, P_{\text{test}}^{\text{SEM}})$  is practical?** In Theorem 1,  $d_{\mathcal{F}}(P^{\text{COV}}, P_{\text{test}}^{\text{SEM}})$  measures the dissimilarity between  $P^{\text{COV}}$  and  $P_{\text{test}}^{\text{SEM}}$ . It seems that this lower bound will be very small and trivial when  $d_{\mathcal{F}}(P^{\text{COV}}, P_{\text{test}}^{\text{SEM}})$  is large enough. However, since the semantic OOD samples can be any samples that do not belong to ID classes, one can suppose that semantic OOD samples are extremely diverse and some are of high similarity with ID and covariate-shifted OOD [37]. Detecting these “ID-like” OOD samples is inherently the core challenge of OOD detection [11, 17, 38]. Thus, it is reasonable to assume a limited  $d_{\mathcal{F}}(P^{\text{COV}}, P_{\text{test}}^{\text{SEM}})$ . We provide more discussions in the Appendix D.2.

<sup>3</sup>Encompassed by [36] as a special case, our definition is realized using TVD for theoretical convenience.

As presented before, the key limitation of MSP detectors is that they enforce high-entropy prediction on semantic OOD. We proceed to reveal that EBM detectors suffer from similar issues due to the natural property of gradient-based optimization. For EBM detectors,  $\mathcal{L}_{\text{reg}}$  is defined by [8] is

$$\mathcal{L}_{\text{reg}} = \mathbb{E}_{\tilde{x} \sim P_{\text{train}}^{\text{SEM}}} [\max(m_{\text{OUT}} - E(\tilde{x}), 0)]^2 + \mathbb{E}_{x \sim P^{\text{ID}}} [\max(0, E(x) - m_{\text{IN}})]^2, \quad (6)$$

which constrains the energy score  $E(x; f) := -\log \sum_{k=1}^K e^{f_k(x)}$  of ID sample  $x$  to be lower than that of OOD sample  $\tilde{x}$ .  $m_{\text{IN}}, m_{\text{OUT}}$  are manually selected margins. Although such regularization does not indicate high entropy prediction at first glance, unfortunately, we demonstrate that EBM detectors also tend to uncertain prediction when equipped with gradient-based optimization. Here we focus on Gradient Descent (GD) as a showcase. In each training epoch  $t$ , model  $f_\theta$  is updated with GD as  $\theta_{t+1} = \theta_t - \eta \nabla_\theta \mathcal{L}_{\text{reg}}$ , where  $\eta$  is the learning rate. For any sample  $\tilde{x}$  drawn from  $P^{\text{SEM}}$ , the gradient  $\nabla_\theta \mathcal{L}_{\text{reg}}$  can be written as

$$2 \sum_{k=1}^K \nabla f_k(\tilde{x}) \underbrace{\left[ e^{f_k(\tilde{x})} \left( \sum_{k=1}^K e^{f_k(\tilde{x})} \right)^{-1} \right]}_{\text{higher for larger } f_k(\tilde{x})} (m_{\text{OUT}} - E(\tilde{x})), \quad (7)$$

where  $m_{\text{OUT}} - E(\tilde{x}) > 0$  (otherwise the gradient is zero) and  $f_k(\tilde{x})$  is the predicted logits on the  $k$ -th class. For sample  $\tilde{x}$ , when class  $k$  has larger predicted logit, it contributes more to the overall gradient  $\nabla_\theta \mathcal{L}_{\text{reg}}$  and thus could obtain more optimization efforts during backpropagation. Eventually, one can infer that when an EBM detector is about to converge, it tends to high-entropy prediction on  $P_{\text{train}}^{\text{SEM}}$  accordingly. Incorporating this into the established Theorem 1, this is likely to harm the generalization ability of  $P^{\text{COV}}$ . Empirical evidence can support this supposition (see Table 14 in Appendix C). Therefore both MSP and EBM detectors face the “*sensitive-robust*” dilemma.

## 5 Decoupled Uncertainty Learning

We demonstrate how to handle the dilemma between OOD detection and generalization by decoupled uncertainty learning in the Bayesian framework. Unlike the most related work [4] which aims to seek a good trade-off, our method successfully gets out of the aforementioned sensitive-robust dilemma.

**Uncertainty Estimation in Bayesian Framework.** We first revisit the theoretical properties of different types of uncertainty in a Bayesian framework. Non-Bayesian classifiers consider the model’s output  $f(x)$  as logits, which is then normalized with softmax to directly model predictive categoricals  $p(\hat{y}|x)$ . While in Bayesian framework [9],  $f(x)$  is considered as parameters of a Dirichlet distribution  $p(\mu|x)$  firstly, which is used to model the prior of predictive categoricals  $p(\hat{y}|x)$  by

$$p(\mu|x) = \text{Dir}(\mu|\alpha) = \frac{\Gamma(\alpha_0)}{\prod_{k=1}^K \Gamma(\alpha_k)} \prod_{k=1}^K \mu_k^{\alpha_k - 1}, \quad \alpha = f(x), \quad (8)$$

where  $\text{Dir}(\mu|\alpha)$  is Dirichlet distribution and  $\alpha$  is the concentration parameters of Dirichlet. The sum of all  $\alpha_k \in \alpha$  (noted as  $\alpha_0$ ) is so called the strength of the Dirichlet distribution, i.e.,  $\alpha_k > 0, \alpha_0 = \sum_k \alpha_k$ . After obtaining prior  $p(\mu|x, \theta)$ , the final predicted posterior  $p(\hat{y}|x)$  over class labels is given by calculating the mean of the Dirichlet prior

$$\underbrace{p(\hat{y}|x, \theta)}_{\text{overall uncertainty}} = \int \underbrace{p(\hat{y}|\mu)}_{\text{data uncertainty}} \underbrace{p(\mu|x, \theta)}_{\text{distributional uncertainty}} d\mu. \quad (9)$$

From a Bayesian perspective, given deterministic parameters  $\theta$ , the overall uncertainty of final prediction  $p(\hat{y}|x)$  can be decomposed into two factors, including data (aleatoric) uncertainty (AU) and distributional uncertainty (DU). DU lies in  $p(\mu|x, \theta)$  which is defined as uncertainty due to the mismatch between the distributions of test and train data. AU is described by  $p(\hat{y}|\mu)$  which captures the natural complexity of the data (e.g., class-overlap) [9]. By definition, OOD detection is primarily associated with DU which is only a part of the overall uncertainty. While generalization is related to the overall uncertainty of  $p(\hat{y}|x)$  as we mentioned in related works (both AU and DU). One essential property of DU is that it can be high even if the expected categorical  $p(\hat{y}|x, \theta)$  expresses low overall

uncertainty. Such a property is well suited to achieve OOD detection and generalization jointly since high DU no longer necessarily indicates high overall uncertainty.

**Decoupled Uncertainty Learning.** While the aforementioned Bayesian framework enjoys theoretical potentiality, its learning object [9] lacks consideration of OOD generalization. Similar to other OOD detection methods, it also directly enforces high overall uncertainty on OOD

$$\min_{\theta} \mathbb{E}_{P^{\text{ID}}} \text{KL}(p(y|x)||p(\hat{y}|x)) + \mathbb{E}_{P_{\text{train}}^{\text{SEM}}} \text{KL}(p(\mu|\tilde{x})||\text{Dir}(\mu|\alpha = \mathbf{1})), \quad (10)$$

where  $p(y|x), p(\hat{y}|x)$  are the ground-truth distribution and predicted distribution on ID. The model’s prediction on OOD is enforced to be close to a rather flat Dirichlet distribution. It is worth noting that  $\text{Dir}(\mu|\alpha = \mathbf{1})$  means all classes are equiprobable, and the entropy of the final prediction is maximized. As shown in Fig. 1 (b), the vanilla Bayesian method [9] also suffers from degraded OOD generalization performance. To this end, we propose **Decoupled Uncertainty Learning (DUL)**, a novel OOD detection regularization method that explicitly encourages high DU on OOD samples without affecting the overall uncertainty. Similarly to previous OOD detection methods [8], our DUL is also devised in a finetune manner for effectiveness. Given a classifier  $f_{\theta_0}$  well pre-trained on  $P^{\text{ID}}$ , the goal of DUL lies in enhancing its OOD detection performance without sacrificing any generalization ability. Specifically, we finetune the model by encouraging higher DU but non-increased overall uncertainty on  $P_{\text{train}}^{\text{SEM}}$ . The learning objective of DUL is

$$\begin{aligned} \min_{\theta} & \underbrace{\mathbb{E}_{(x,y) \sim P^{\text{ID}}} [\mathcal{L}_{\text{CE}}(f(x), y)]}_{\text{ID classification}} + \lambda \underbrace{\mathbb{E}_{\tilde{x} \sim P_{\text{train}}^{\text{SEM}}} \|\max(0, (h_0 + m_{\text{OUT}}) - h)\|_{\tau}}_{\text{high distributional uncertainty (detection)}} \\ \text{s.t.} & \underbrace{H(p(\hat{y}|\tilde{x})) = H(p_0(\hat{y}|\tilde{x}))}_{\text{non-increased overall uncertainty (generalization)}} \quad \forall \tilde{x} \sim P_{\text{train}}^{\text{SEM}}, \end{aligned} \quad (11)$$

where  $H(\cdot)$  is the entropy.  $p(\hat{y}|\tilde{x})$  and  $p_0(\hat{y}|\tilde{x})$  are the predicted distribution on semantic OOD data  $\tilde{x}$  after and before finetuning. The first term is the original ID classification loss. The second term is OOD detection loss, which encourages high DU on outlier  $\tilde{x}$ .  $m_{\text{OUT}}$  and  $\tau > 0$  are hyperparameters.  $h_0, h$  are DU on  $\tilde{x}$  before and after finetuning. Here we measure DU with the differential entropy ( $h[p(\mu|\tilde{x}, \theta)] = -\int_{S^{K-1}} p(\mu|\tilde{x}) \ln(p(\mu|\tilde{x})) d\mu$ ,  $S$  is a  $K$ -simplex). We refer interested readers to the Appendix D.1 for mathematical details. The third term constraining on  $H(p(\hat{y}|\tilde{x}))$  avoids increment of overall uncertainty during finetuning and thus the generalization ability can be retained. Considering the difficulty of constrained optimization, we convert Eq. 11 into an unconstrained form and get our final minimizing objective

$$\underbrace{\mathbb{E}_{P^{\text{ID}}} [\mathcal{L}_{\text{CE}}(f(x), y)]}_{\text{ID classification}} + \underbrace{\mathbb{E}_{P_{\text{train}}^{\text{SEM}}} \{\lambda \|\max(0, (h_0 + m_{\text{OUT}}) - h)\|_{\tau}\}}_{\text{high distributional uncertainty}} + \underbrace{\gamma \text{KL}(p(\hat{y}|\tilde{x})||p_0(\hat{y}|\tilde{x}))}_{\text{unchanged overall uncertainty}}, \quad (12)$$

where  $\gamma$  is hyperparameter. In contrast to previous Bayesian method [9], DUL only encourages high DU rather than overall uncertainty on OOD and explicitly discourages high entropy in the final prediction. The implementation details are in Appendix D.1.

## 6 Experiment

We conduct experiments to validate our analysis and the superiority of DUL. The questions to be verified are Q1 Motivation. To what extent does OOD detection conflict with OOD generalization in previous methods? Q2 Effectiveness. Does DUL achieve better OOD detection and generalization performance compared to its counterparts? Q3 Interpretability. Does the proposed method well decouple uncertainty as expected?

### 6.1 Experimental Setup

Our settings follow the common practice [8, 11, 20, 5] in OOD detection. Here we present a brief description and more details about datasets, metrics, and implementation are in Appendix B.1 and B.2.

**Datasets.**  $\circ$  **ID datasets**  $P^{\text{ID}}$ . We train the model on different ID datasets including CIFAR-10, CIFAR-100 and ImageNet-200 (a subset of ImageNet-1K [39] with 200 classes).  $\circ$  **Auxiliary OOD datasets**  $P_{\text{train}}^{\text{SEM}}$ . In CIFAR experiments, we use ImageNet-RC as  $P_{\text{train}}^{\text{SEM}}$ . ImageNet-RC is a down-sampled variant of the original ImageNet-1K which is widely adopted in previous OOD detection

Table 1: OOD detection and generalization performance comparison. Substantial ( $\geq 0.5$ ) **improvement** and **degradation** compared to the baseline MSP [6] (training without any OOD detection regularization) are highlighted in blue or red, respectively. The **best** and **second-best** results are in bold or underlined. DUL is the only method that achieves SOTA OOD detection performance (mostly the best or second best) without sacrificing generalization i.e., the value of the entire row is either blue or black. Full results with standard deviation and diverse types of corruption are in Appendix C.

$P^{\text{ID}}/P^{\text{SEM}}_{\text{train}}$	Method	Model generalization		FPR $\downarrow$	OOD detection	
		ID-Acc $\uparrow$	OOD-Acc $\uparrow$		AUROC $\uparrow$	AUPR $\uparrow$
CIFAR-10 / None	MSP	96.11	87.35	41.96	89.28	68.00
	EBM (pretrain)	96.11	87.35	32.45	89.34	75.22
	Maxlogits	96.11	87.35	32.90	89.26	74.47
	Mahalanobis	96.11	87.35	32.53	93.93	74.96
CIFAR-10 / ImageNet-RC	Entropy	<u>96.04</u>	<u>72.57</u>	6.63	<u>98.72</u>	94.00
	EBM (finetune)	<b>96.10</b>	<u>79.03</u>	<u>3.61</u>	98.39	94.88
	POEM	<u>94.32</u>	<u>78.89</u>	<b>3.32</b>	<b>98.99</b>	<b>99.38</b>
	DPN	95.69	<u>85.52</u>	4.28	98.53	94.93
	WOODS	96.01	<u>80.14</u>	7.12	98.45	92.46
	SCONE	95.96	<u>78.80</u>	7.02	98.45	92.46
	DUL (ours)	<u>96.02</u> $\pm 0.00$	<b>88.01</b> $\pm 0.29$	<u>5.89</u> $\pm 0.12$	<u>98.47</u> $\pm 0.02$	<u>92.44</u> $\pm 1.29$
	DUL <sup>†</sup> (ours)	<u>96.04</u> $\pm 0.00$	<b>87.53</b> $\pm 0.49$	<u>5.99</u> $\pm 0.06$	<u>98.28</u> $\pm 0.01$	<u>98.40</u> $\pm 0.13$
CIFAR-10 / TIN-597	Entropy	<u>95.94</u>	<u>80.51</u>	11.60	97.93	92.16
	EBM (finetune)	<u>95.38</u>	<u>83.67</u>	19.36	<u>87.51</u>	83.63
	POEM	<u>95.44</u>	<u>83.17</u>	24.34	<u>86.83</u>	<u>94.25</u>
	DPN	<u>94.39</u>	<u>79.23</u>	17.27	94.92	87.67
	WOODS	<u>95.57</u>	<u>84.68</u>	<u>7.58</u>	<b>98.29</b>	93.08
	SCONE	<u>95.19</u>	<u>84.68</u>	8.02	<u>98.21</u>	93.08
	DUL (ours)	<b>96.06</b> $\pm 0.01$	<u>87.93</u> $\pm 0.39$	<b>6.87</b> $\pm 0.67$	<u>98.21</u> $\pm 0.01$	<u>91.29</u> $\pm 1.39$
	DUL <sup>†</sup> (ours)	<u>95.94</u> $\pm 0.01$	<b>88.10</b> $\pm 0.07$	<u>10.34</u> $\pm 0.11$	<u>97.67</u> $\pm 0.01$	<b>98.59</b> $\pm 0.06$
CIFAR-100 / None	MSP	80.99	55.95	74.63	80.19	42.59
	EBM (pretrain)	80.99	55.95	67.42	82.67	49.35
	Maxlogits	80.99	55.95	69.32	82.30	47.60
	Mahalanobis	80.99	55.95	61.51	85.97	56.10
CIFAR-100 / ImageNet-RC	Entropy	<u>80.21</u>	<u>45.48</u>	22.29	95.33	82.34
	EBM (finetune)	80.53	<u>48.14</u>	13.47	<u>96.78</u>	87.84
	POEM	<u>78.15</u>	<u>42.18</u>	<b>9.89</b>	<b>97.79</b>	<b>98.40</b>
	DPN	<u>78.90</u>	<u>50.14</u>	18.36	95.42	74.45
	WOODS	80.69	<u>54.38</u>	38.15	92.01	71.79
	SCONE	80.80	<b>56.73</b>	47.60	89.61	65.29
	DUL (ours)	<b>81.30</b> $\pm 0.04$	<u>56.27</u> $\pm 3.29$	<u>12.49</u> $\pm 0.05$	<u>95.24</u> $\pm 0.01$	<u>86.72</u> $\pm 0.58$
	DUL <sup>†</sup> (ours)	<u>81.23</u> $\pm 0.05$	<u>55.41</u> $\pm 0.54$	<u>11.12</u> $\pm 0.62$	<u>95.46</u> $\pm 0.36$	<u>96.49</u> $\pm 0.13$
CIFAR-100 / TIN-597	Entropy	<u>80.15</u>	<u>46.25</u>	26.88	93.50	79.81
	EBM (finetune)	<u>79.94</u>	<u>50.00</u>	26.87	91.68	80.08
	POEM	<u>78.68</u>	<u>52.53</u>	32.71	91.30	<u>94.65</u>
	DPN	<u>78.44</u>	<u>47.67</u>	24.99	<u>93.55</u>	81.63
	WOODS	<u>79.26</u>	<u>53.13</u>	36.71	92.15	73.42
	SCONE	<u>79.53</u>	<u>52.70</u>	35.60	92.47	73.58
	DUL (ours)	<b>80.85</b> $\pm 0.06$	<u>56.19</u> $\pm 2.33$	<u>23.32</u> $\pm 1.22$	<b>94.48</b> $\pm 0.12$	<u>80.82</u> $\pm 2.63$
	DUL <sup>†</sup> (ours)	<u>80.50</u> $\pm 0.06$	<b>56.22</b> $\pm 1.66$	<b>22.75</b> $\pm 0.78$	<u>90.88</u> $\pm 0.08$	<b>96.33</b>
ImageNet-200 / None	MSP	<u>85.15</u>	74.84	58.23	86.98	82.27
	EBM (pretrain)	<u>85.15</u>	74.84	51.94	88.18	84.75
	Maxlogits	<u>85.15</u>	74.84	<u>51.62</u>	88.30	84.71
ImageNet-200 / ImageNet-800	Entropy	84.92	74.75	53.62	<u>89.05</u>	<u>85.02</u>
	EBM (finetune)	<u>84.14</u>	<u>73.31</u>	<u>59.73</u>	87.54	82.81
	DPN	84.87	74.40	<u>63.84</u>	87.18	<u>80.69</u>
ImageNet-800	WOODS	84.99	<u>74.98</u>	51.71	88.30	84.80
	SCONE	84.93	74.91	52.52	88.19	84.50
	DUL (ours)	<b>85.65</b> $\pm 0.07$	<b>75.59</b> $\pm 0.12$	<u>49.14</u> $\pm 0.13$	<b>89.27</b> $\pm 0.03$	<b>85.62</b> $\pm 0.03$

works [8, 11, 17]. We also conduct experiments on the recent TIN-597 [20] as an alternative. When ImageNet-200 is ID, the remaining 800 classes termed ImageNet-800 are considered as  $P^{\text{SEM}}_{\text{train}}$ . **OOD detection test sets**  $P^{\text{SEM}}_{\text{test}}$  are a suite of diverse datasets introduced by commonly used benchmark [5]. In CIFAR experiments, we use SVHN [40], Places365 [41], Textures [42], LSUN-R, LSUN-C [43] and iSUN [44] as  $P^{\text{SEM}}_{\text{test}}$ . When  $P^{\text{ID}}$  is ImageNet-200,  $P^{\text{SEM}}_{\text{test}}$  consists of iNaturalist [45], Open-Image [46], NINCO [47] and SSB-Hard [48]. It is worth noting that in standard OOD detection settings, there should be no overlapped classes between  $P^{\text{ID}}$ ,  $P^{\text{SEM}}_{\text{train}}$  and  $P^{\text{SEM}}_{\text{test}}$ , otherwise OOD detection is a trivial problem. **OOD generalization test sets**  $P^{\text{COV}}$  is the original ID test set corrupted with additive Gaussian noise of  $\mathcal{N}(0, 5)$ , following [4]. Besides, we also conduct experiments on CIFAR10-C, CIFAR100-C and ImageNet-C which involve 15 diverse types of different noise or corruption (e.g., snow, rain, frost, fog...) in Appendix C.



**Metrics.** For OOD detection performance evaluation, we report the average FPR95, AUROC and AUPR to be consistent with [11]. OOD generalization ability is compared in terms of classification accuracy (OOD-Acc). Besides, we also report classification accuracy on ID test sets (ID-Acc).

**Compared methods.** We compare DUL with a board line of OOD detection methods, including auxiliary OOD required and auxiliary OOD free methods.

- **Auxiliary OOD-free methods** do not require  $P_{\text{train}}^{\text{SEM}}$  during training, including MSP [6], Maxlogits [49], pretrained EBM [8] and Mahalaobis [18].
- **Auxiliary OOD-required methods** explicitly regularize the model on  $P_{\text{train}}^{\text{SEM}}$ , including entropy-regularization (Entropy) [7], finetuned EBM [8], DPN of Bayesian framework [9], POEM [11] and WOODS [10]. We also compare our DUL to recent advanced SCONE [4] which aims to keep a balance between OOD detection and generalization.

## 6.2 Experimental Results

**Dilemma between OOD detection and generalization (Q1).** We validate the dilemma mentioned before in Fig. 1. As shown in Tab. 1, though many advanced methods establish superior OOD detection performance, their OOD generalization degrades a lot. For example, recent SOTA POEM achieves nearly perfect OOD detection performance on CIFAR10 when ImageNet-RC serves as  $P_{\text{train}}^{\text{SEM}}$  with 3.32% false positive error rate (FPR95). However, its OOD-Acc drops a lot (about 10%) compared to baseline MSP. This phenomenon is also observed in other advanced methods. To further detail this phenomenon, we reduce the weight of OOD detection regularization terms in Entropy and finetuned EBM and show the performance on both OOD detection and generalization. As shown in Table 3, when the regularization strength increases, OOD detection performance improves (lower FPR.), while the OOD generalization performance degrades (higher error rate).

**OOD detection and generalization ability (Q2).** As shown in Tab. 1, DUL establishes strong overall performance in terms of both OOD detection and generalization. We highlight a few essential observations: 1) **Compared to auxiliary OOD free methods**, DUL establishes substantial improvement due to additional regularization on auxiliary outliers. 2) **Compared to auxiliary OOD required methods**, our method achieves superior OOD detection performance without sacrificing generalization ability. Meanwhile, previous OOD detection methods commonly exhibit severely degraded classification accuracy, with many cases increasing by more than 10% error rate. 3) **Comparison to the most related work SCONE [4].** Despite recent advanced SCONE simultaneously considering both two targets, we observe that it can be hard to find a good trade-off. In contrast, dual-optimal OOD detection and generalization performance is achieved by our DUL. Noted that DUL is the only method that achieves *state-of-the-art* detection performance (mostly the best or second best) without degraded generalization ability (no red values in the entire row). The sensitive-robust dilemma is no longer observed in our method. These observations justify our expectation of DUL. 4) **Combining with existing methods.** Besides, to further demonstrate the effectiveness of the proposed DUL, we also add the unchanged overall uncertainty term in Eq.12 to the original Entropy and finetuned EBM. The results in Table 2 show that DUL regularization can also benefit EBM. However, combining Entropy with our regularization can not improve the accuracy substantially. This is not surprising, since the target of Entropy (high entropy prediction) and our DUL (non-increased entropy) directly conflict according to Theorem 1. 5) **Comparison to methods with an extra OOD detect branch.** Different from aforementioned methods, a line of recent OOD detectors [50, 51, 52, 17] employ extra output branches aside from the classification logits (with a shared backbone for feature extraction). For these OOD detectors, our theoretical analysis is not directly applicable and further analysis from a feature learning perspective may be needed in future work. However, the proposed DUL is devised in a finetune manner. Compared to OOD detectors with extra output branches that requires re-training the classifier from scratch, DUL can be applied to any pre-trained model (e.g., from torchvision, huggingface), with modest computation overhead.

**Visualization of estimated uncertainty (Q3).** To evaluate the uncertainty estimation, we visualize the distribution of ID (CIFAR-10) and OOD (SVHN) samples in terms of uncertainty. As we can see in Fig. 2 (b), our DUL establishes a distinguishable (distributional) uncertainty gap between test-time ID and OOD data, which indicates a good sensitiveness for OOD detection. By contrast, the baseline method MSP (Fig.2 (a)) can not effectively discriminate ID and OOD. Besides, we visualize the predictive entropy (overall uncertainty) on covariate-shifted OOD (CIFAR-10 with Gaussian noise) in Fig. 2 (c), our DUL yields much lower entropy compared to other methods. Besides, we visualize the data uncertainty on semantic OOD test data (Textures) when CIFAR-10 is ID in Fig. 6.2. The

investigated methods are 1) pretrained model training on ID dataset only, 2) finetuned model with OOD detection regularization (ablating the last term in Eq.12), and 3) finetuned model with the full DUL method described by Eq.12. As shown in Fig. 6.2, to keep the overall uncertainty and enlarge the distributional uncertainty (for OOD detection), the data uncertainty must be reduced. We use Eq.17 from [9] to calculate data uncertainty. The distributional uncertainty is shifted by subtracting that on ID dataset. These results meet our expectation.




Figure 2: Visualization of different types of uncertainty estimated by DUL.




Figure 3: Visualization of uncertainty on semantic OOD test dataset when CIFAR-10 is ID dataset. Without DUL (orange), all three types of uncertainty will increase altogether. In contrast, DUL (green) increases the DU but decreases the AU, which further lead to unchanged overall uncertainty.

Table 2: Additional results when equip DUL to existing methods i.e., Entropy and finetuned EBM. ID dataset is CIFAR-10.  $P_{train}^{SEM}$  is ImageNet-RC.  $P_{test}^{Cov}$  is the original CIFAR-10 testset corrupted by Gaussian noise  $\mathcal{N}(0, 5)$ .

Method	Model generalization		OOD detection	
	ID-Acc $\uparrow$	OOD-Acc $\uparrow$	FPR $\downarrow$	AUC $\uparrow$
Entropy	96.04	72.57	6.63	98.72
EBM (finetune)	96.10	79.03	3.61	98.39
POEM	94.32	78.89	3.32	98.99
EBM w. DUL	95.19	87.45	6.17	98.28
Entropy w. DUL	96.10	87.41	29.56	95.92
DUL	96.02	88.01	5.89	98.47
DUL <sup>†</sup>	96.04	87.53	5.99	98.28

Table 3: We tune the weight of OOD detection regularization term for EBM as well as Entropy and report the FPR (OOD detection metric) and error rate (Err, OOD generalization metric). The experimental settings are the same with Table 2.

$\lambda$	Entropy		EBM		
	OOD-Err $\downarrow$	FPR $\downarrow$	OOD-Err $\downarrow$	FPR $\downarrow$	
0	9.55	35.15	0	9.55	20.57
$5 \times 10^{-4}$	13.58	8.36	$1 \times 10^{-4}$	9.46	14.69
$5 \times 10^{-3}$	15.48	6.37	$1 \times 10^{-3}$	10.32	13.54
$5 \times 10^{-2}$	17.97	5.71	$1 \times 10^{-2}$	16.43	8.15
$5 \times 10^{-1}$	18.53	5.60	$1 \times 10^{-1}$	24.38	6.11

## 7 Conclusion

This paper provides both theoretical and empirical analysis towards understanding the dilemma between OOD detection and generalization. We demonstrate that the superior OOD detection performance of current advances are achieved at the cost of generalization ability. The theory-inspired algorithm successfully removes the conflict between previous OOD detection and generalization methods. For SOTA OOD detection performance, our implementation assumes that auxiliary outliers are available during training. This limitation is noteworthy for our DUL as well as the most existing SOTA OOD detection methods. We argue that this added cost is minor and reasonable given the significance of ensuring model trustworthiness in open-environments. Reducing the dependency on auxiliary OOD data can be an interesting research direction for the future exploration.

## Acknowledgement

This work was supported by the National Natural Science Foundation of China (62376193 and 61925602). The authors thank NeurIPS anonymous peer reviewers for their helpful suggestions.

## References

- [1] Chunjong Park, Anas Awadalla, Tadayoshi Kohno, and Shwetak Patel. Reliable and trustworthy machine learning for health using dataset shift detection. *Advances in Neural Information Processing Systems*, 34:3043–3056, 2021.
- [2] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [3] Jiashuo Liu, Zheyang Shen, Yue He, Xingxuan Zhang, Renzhe Xu, Han Yu, and Peng Cui. Towards out-of-distribution generalization: A survey. *arXiv preprint arXiv:2108.13624*, 2021.
- [4] Haoyue Bai, Gregory Canal, Xuefeng Du, Jeongyeol Kwon, Robert D Nowak, and Yixuan Li. Feed two birds with one scone: Exploiting wild data for both out-of-distribution generalization and detection. In *International Conference on Machine Learning*, 2023.
- [5] Jingyang Zhang, Jingkang Yang, Pengyun Wang, Haoqi Wang, Yueqian Lin, Haoran Zhang, Yiyun Sun, Xuefeng Du, Kaiyang Zhou, Wayne Zhang, Yixuan Li, Ziwei Liu, Yiran Chen, and Hai Li. Openood v1.5: Enhanced benchmark for out-of-distribution detection. *arXiv preprint arXiv:2306.09301*, 2023.
- [6] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *International Conference on Learning Representations*, 2017.
- [7] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. *International Conference on Learning Representations*, 2019.
- [8] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. *Advances in neural information processing systems*, 2020.
- [9] Andrey Malinin and Mark Gales. Predictive uncertainty estimation via prior networks. *Advances in neural information processing systems*, 2018.
- [10] Julian Katz-Samuels, Julia B Nakhleh, Robert Nowak, and Yixuan Li. Training ood detectors in their natural habitats. In *International Conference on Machine Learning*, 2022.
- [11] Yifei Ming, Ying Fan, and Yixuan Li. Poem: Out-of-distribution detection with posterior sampling. In *International Conference on Machine Learning*, 2022.
- [12] Xuefeng Du, Zhaoning Wang, Mu Cai, and Yixuan Li. Vos: Learning what you don’t know by virtual outlier synthesis. *International Conference on Learning Representations*, 2022.
- [13] Jingkang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*, 2021.
- [14] Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. *International Conference on Learning Representations*, 2018.
- [15] Kimin Lee, Honglak Lee, Kibok Lee, and Jinwoo Shin. Training confidence-calibrated classifiers for detecting out-of-distribution samples. *International Conference on Learning Representations*, 2018.
- [16] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. *Advances in neural information processing systems*, 2019.
- [17] Jiefeng Chen, Yixuan Li, Xi Wu, Yingyu Liang, and Somesh Jha. Atom: Robustifying out-of-distribution detection using outlier mining. In *Machine Learning and Knowledge Discovery in Databases*, 2021.
- [18] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 2018.

- [19] Yiyou Sun, Yifei Ming, Xiaojin Zhu, and Yixuan Li. Out-of-distribution detection with deep nearest neighbors. In *International Conference on Machine Learning*, pages 20827–20840. PMLR, 2022.
- [20] Jingkang Yang, Pengyun Wang, Dejian Zou, Zitang Zhou, Kunyuan Ding, Wenxuan Peng, Haoqi Wang, Guangyao Chen, Bo Li, Yiyou Sun, et al. Openood: Benchmarking generalized out-of-distribution detection. *Advances in Neural Information Processing Systems*, 35:32598–32611, 2022.
- [21] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- [22] Zhaolin Hu and L Jeff Hong. Kullback-leibler divergence constrained distributionally robust optimization. *Available at Optimization Online*, 2013.
- [23] Paul Michel, Tatsunori Hashimoto, and Graham Neubig. Modeling the second player in distributionally robust optimization. *International Conference on Learning Representations*, 2021.
- [24] Aman Sinha, Hongseok Namkoong, Riccardo Volpi, and John Duchi. Certifying some distributional robustness with principled adversarial training. *International Conference on Learning Representations*, 2018.
- [25] Zongbo Han, Zhipeng Liang, Fan Yang, Liu Liu, Lanqing Li, Yatao Bian, Peilin Zhao, Bingzhe Wu, Changqing Zhang, and Jianhua Yao. Umix: Improving importance weighting for subpopulation shift via uncertainty-aware mixup. *Advances in Neural Information Processing Systems*, 2022.
- [26] Han Zhao, Remi Tachet Des Combes, Kun Zhang, and Geoffrey Gordon. On learning invariant representations for domain adaptation. In *International conference on machine learning*, pages 7523–7532. PMLR, 2019.
- [27] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. In *International Conference on Machine Learning*, 2019.
- [28] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. *International Conference on Learning Representations*, 2021.
- [29] Marvin Zhang, Sergey Levine, and Chelsea Finn. Memo: Test time robustness via adaptation and augmentation. *Advances in neural information processing systems*, 35:38629–38642, 2022.
- [30] Shuaicheng Niu, Jiaxiang Wu, Yifan Zhang, Zhiquan Wen, Yafo Chen, Peilin Zhao, and Mingkui Tan. Towards stable test-time adaptation in dynamic wild world. *arXiv preprint arXiv:2302.12400*, 2023.
- [31] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 2017.
- [32] Qizhou Wang, Zhen Fang, Yonggang Zhang, Feng Liu, Yixuan Li, and Bo Han. Learning to augment distributions for out-of-distribution detection. *Advances in Neural Information Processing Systems*, 36, 2024.
- [33] Hyunjun Choi, Hawook Jeong, and Jin Young Choi. Balanced energy regularization loss for out-of-distribution detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15691–15700, 2023.
- [34] Xuefeng Du, Yiyou Sun, Jerry Zhu, and Yixuan Li. Dream the impossible: Outlier imagination with diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024.
- [35] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning*, 2015.
- [36] Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation with multiple sources. *Advances in neural information processing systems*, 21, 2008.

- [37] William Yang, Byron Zhang, and Olga Russakovsky. Imagenet-ood: Deciphering modern out-of-distribution detection algorithms. *International Conference on Learning Representations*, 2024.
- [38] Yichen Bai, Zongbo Han, Changqing Zhang, Bing Cao, Xiaoheng Jiang, and Qinghua Hu. Id-like prompt learning for few-shot out-of-distribution detection. *arXiv preprint arXiv:2311.15243*, 2023.
- [39] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015.
- [40] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Baolin Wu, Andrew Y Ng, et al. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*. Granada, Spain, 2011.
- [41] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [42] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014.
- [43] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
- [44] Junting Pan and Xavier Giró-i Nieto. End-to-end convolutional network for saliency prediction. *arXiv preprint arXiv:1507.01422*, 2015.
- [45] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018.
- [46] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International journal of computer vision*, 128(7):1956–1981, 2020.
- [47] Julian Bitterwolf, Maximilian Mueller, and Matthias Hein. In or out? fixing imagenet out-of-distribution detection evaluation. In *International Conference on Machine Learning*, 2023.
- [48] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. The semantic shift benchmark. In *ICML 2022 Shift Happens Workshop*, 2022.
- [49] Dan Hendrycks, Steven Basart, Mantas Mazeika, Andy Zou, Joe Kwon, Mohammadreza Mostajabi, Jacob Steinhardt, and Dawn Song. Scaling out-of-distribution detection for real-world settings. *International conference on machine learning*, 2022.
- [50] Wenjun Miao, Guansong Pang, Xiao Bai, Tianqi Li, and Jin Zheng. Out-of-distribution detection in long-tailed recognition with calibrated outlier class learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 4216–4224, 2024.
- [51] Julian Bitterwolf, Alexander Meinke, Maximilian Augustin, and Matthias Hein. Breaking down out-of-distribution detection: Many methods based on ood training data estimate a combination of the same core quantities. In *International Conference on Machine Learning*, pages 2041–2074. PMLR, 2022.
- [52] Kai Liu, Zhihang Fu, Chao Chen, Sheng Jin, Ze Chen, Mingyuan Tao, Rongxin Jiang, and Jieping Ye. Category-extensible out-of-distribution detection via hierarchical context descriptions. *Advances in Neural Information Processing Systems*, 36, 2024.

- [53] Clément L Canonne. A short note on an inequality between kl and tv. *arXiv preprint arXiv:2202.07198*, 2022.
- [54] Haoqi Wang, Zhizhong Li, Litong Feng, and Wayne Zhang. Vim: Out-of-distribution with virtual-logit matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4921–4930, 2022.
- [55] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Open-set recognition: A good closed-set classifier is all you need? *International Conference on Learning Representations*, 2022.
- [56] Dan Hendrycks and Thomas G Dietterich. Benchmarking neural network robustness to common corruptions and surface variations. *International Conference on Learning Representations*, 2019.
- [57] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *British Machine Vision Conference*, 2016.
- [58] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.
- [59] Yarin Gal et al. Uncertainty in deep learning. 2016.

## Appendices

<b>A Proofs</b>	<b>15</b>
<b>B Experimental Details</b>	<b>17</b>
B.1 Datasets details . . . . .	17
B.2 Implementation details . . . . .	18
<b>C Additional Results</b>	<b>20</b>
C.1 Uncertainty estimation. . . . .	20
C.2 Time-consuming comparison. . . . .	20
C.3 Full results with standard deviation. . . . .	20
C.4 Results of different types of corruption. . . . .	21
C.5 OOD detection results on individual datasets. . . . .	21
C.6 Empirical evidence . . . . .	24
<b>D Discussions</b>	<b>24</b>
D.1 Math derivation . . . . .	24
D.2 Discussion about Disparity Discrepancy . . . . .	24
D.3 Social Impact . . . . .	25

## A Proofs

First, we recap the definitions of Disparity with Total Variation Distance and Disparity Discrepancy.

**Definition 3** (Disparity with Total Variation Distance). *Given two hypotheses  $f', f \in \mathcal{F}$  and distribution  $P$ , we define the Disparity with Total Variation Distance between them as*

$$\text{disp}_P(f', f) = \mathbb{E}_P[TV(F_f || F_{f'})], \quad (13)$$

where  $F_f, F_{f'}$  are the class distributions predicted by  $f', f$  respectively.  $TV(\cdot || \cdot)$  is the total variation distance (TVD), i.e.,  $TV(F_f || F_{f'}) = \frac{1}{2} \sum_{k=1}^K \|F_{f,k} - F_{f',k}\|_1$ .

**Definition 4** (Disparity Discrepancy with Total Variation Distance, DD with TVD). *Given a hypothesis space  $\mathcal{F}$  and two distributions  $P, Q$ , the Disparity Discrepancy with Total Variation Distance (DD with TVD) is defined as*

$$d_{\mathcal{F}}(P, Q) := \sup_{f', f \in \mathcal{F}} (\text{disp}_P(f', f) - \text{disp}_Q(f', f)). \quad (14)$$

Since TVD is a distance measurement of two distribution. It yields the triangle equality. That is, for any distribution  $P_{\mathcal{X}}$  support on  $\mathcal{X}$  and hypotheses  $f^1, f^2$  and  $f^3 \in \mathcal{F}$ , we have

$$\begin{aligned} \text{disp}_{P_{\mathcal{X}}}(f^1, f^2) &\leq \mathbb{E}_{x \sim P_{\mathcal{X}}}[TV(F_{f^1}(x) || F_{f^3}(x))] + \mathbb{E}_{x \sim P_{\mathcal{X}}}[TV(F_{f^2}(x) || F_{f^3}(x))], \\ \text{disp}_{P_{\mathcal{X}}}(f^1, f^2) &\geq \mathbb{E}_{x \sim P_{\mathcal{X}}}[TV(F_{f^1}(x) || F_{f^3}(x))] - \mathbb{E}_{x \sim P_{\mathcal{X}}}[TV(F_{f^2}(x) || F_{f^3}(x))]. \end{aligned} \quad (15)$$

To prove Theorem 1, we need the following lemmas.

**Lemma 1.** *For any  $f \in \mathcal{F}$ , we have*

$$\mathbb{E}_{P^{\text{COV}}} TV(F_f || U) \leq \mathbb{E}_{P_{\text{test}}^{\text{SEM}}} TV(F_f || U) + d_{\mathcal{F}}(P^{\text{COV}}, P_{\text{test}}^{\text{SEM}}) + \lambda \quad (16)$$

where  $\lambda$  is a constant independent of  $f$ .  $U$  is the  $K$ -classes uniform distribution.  $P^{\text{COV}}$  is the covariate-shifted OOD distribution.  $P_{\text{test}}^{\text{SEM}}$  is the semantic OOD distribution.

*Proof.* Let  $f^*$  be the hypothesis which jointly minimizes the total variance distance between the predicted distribution  $F_f$  with uniform distribution  $U$  taking expectation on  $P^{\text{COV}}$  and  $P_{\text{test}}^{\text{SEM}}$ , which is to say

$$f^* := \operatorname{argmin}_{f \in \mathcal{F}} \{ \mathbb{E}_{x \sim P^{\text{COV}}} [TV(F_f(x)||U)] + \mathbb{E}_{x \sim P_{\text{test}}^{\text{SEM}}} [TV(F_f(x)||U)] \}. \quad (17)$$

Set  $\lambda = \mathbb{E}_{x \sim P^{\text{COV}}} [TV(F_{f^*}(x)||U)] + \mathbb{E}_{x \sim P_{\text{test}}^{\text{SEM}}} [TV(F_{f^*}(x)||U)]$ , then by the triangle equality we have

$$\begin{aligned} \mathbb{E}_{P^{\text{COV}}} TV(F_f||U) &\leq \operatorname{disp}_{P^{\text{COV}}}(f, f^*) + \mathbb{E}_{P^{\text{COV}}} TV(F_{f^*}||U) \\ &\leq \mathbb{E}_{P_{\text{test}}^{\text{SEM}}} TV(F_f||U) - \mathbb{E}_{P_{\text{test}}^{\text{SEM}}} TV(F_f||U) + \operatorname{disp}_{P^{\text{COV}}}(f, f^*) + \mathbb{E}_{P^{\text{COV}}} TV(F_{f^*}||U) \\ &\leq \mathbb{E}_{P_{\text{test}}^{\text{SEM}}} TV(F_f||U) + \mathbb{E}_{P_{\text{test}}^{\text{SEM}}} TV(F_{f^*}||U) \\ &\quad - \operatorname{disp}_{P_{\text{test}}^{\text{SEM}}}(f, f^*) + \operatorname{disp}_{P^{\text{COV}}}(f, f^*) + \mathbb{E}_{P^{\text{COV}}} TV(F_{f^*}||U) \\ &\leq \mathbb{E}_{P_{\text{test}}^{\text{SEM}}} TV(F_f||U) + d_{\mathcal{F}}(P_{\text{test}}^{\text{SEM}}, P^{\text{COV}}) + \lambda. \end{aligned} \quad (18)$$

□

Intuitively speaking, Lemma 1 demonstrates that if the classifier  $f$  express high overall uncertainty on  $P_{\text{test}}^{\text{SEM}}$  (i.e., the predicted distribution  $F_f$  is close to uniform distribution), it will also tend to high uncertain prediction on  $P^{\text{COV}}$  given a limited  $d_{\mathcal{F}}(P_{\text{test}}^{\text{SEM}}, P^{\text{COV}})$ .

**Lemma 2.** [Inequality between KL and TV] For any  $K$ -class distribution  $P$  and  $Q$  on  $\{1, \dots, K\}$  and  $\kappa > 0$ , the following inequality holds

$$\frac{1}{2\kappa} KL(P||Q) + \frac{\kappa}{4} \geq TV(P||Q). \quad (19)$$

*Proof.* The proof can be found in [53] page 8. □

**Lemma 3.** Denote the OOD detection loss used for MSP detectors as  $\mathcal{L}_{\text{reg}}$ , then we have

$$\mathbb{E}_{P_{\text{test}}^{\text{SEM}}} TV(F_f||U) \leq \mathbb{E}_{P_{\text{test}}^{\text{SEM}}} \sqrt{\frac{1}{2}(\mathcal{L}_{\text{reg}}(f) - \log K)} \quad (20)$$

where  $TV(\cdot||\cdot)$  is total variance distance (TVD).  $U$  denotes uniform distribution support on  $\mathcal{Y} = \{1, 2 \dots K\}$ .  $\mathcal{L}_{\text{reg}}$  is defined in [7] is the cross-entropy between predicted distribution  $F_f(x)$  and uniform distribution  $U$ .

Lemma 2 means that minimizing the OOD detection loss will constrains the predicted distribution to be close to uniform distribution, which is a intuitive and straightforward result.

*Proof.* In  $K$ -classes classification task, for any sample  $\tilde{x}$  drawn from  $P_{\mathcal{X}}^{\text{SEM}}$ , we have

$$\mathcal{L}_{\text{reg}}(f(\tilde{x})) = KL(U||F_f) + H(U). \quad (21)$$

Applying Pinsker's Inequality, the following inequality holds

$$\mathcal{L}_{\text{reg}}(f(\tilde{x})) = KL(F_f||U) + H(U) \geq 2TV(F_f(\tilde{x})||U)^2 + H(U). \quad (22)$$

Noted that  $H(U) = \log K$ , we can re-write above inequality as

$$TV(F_f(\tilde{x})||U) \leq \sqrt{\frac{1}{2}(\mathcal{L}_{\text{reg}}(f(\tilde{x})) - \log K)}. \quad (23)$$

Then, by taking expectation on  $P_{\mathcal{X}}^{\text{SEM}}$  we can get the result. □

Now we are ready to present the proof of Theorem 1.



*Proof.* By the definition, the generalization error can be written as

$$\begin{aligned} \text{GError}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}}(f) &:= \mathbb{E}_{(x,y) \sim P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} \mathcal{L}_{\text{CE}}(f(x), y) \\ &= \mathbb{E}_{(x,y) \sim P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} [KL(P_{\text{true}} \| F_f(x)) + H(P_{\text{true}})] \end{aligned} \quad (24)$$

where  $P_{\text{true}}$  is the target distribution given input  $x$  (i.e., the true class distribution) and  $\mathcal{L}_{\text{CE}}(\cdot)$  denotes the cross-entropy loss.

Applying Lemma 2, for any  $x$  we have

$$KL(P_{\text{true}} \| F_f) \geq \frac{1}{2\kappa} (TV(P_{\text{true}} \| F_f)) + \frac{\kappa}{4}. \quad (25)$$

By the sub-additivity of TVD, we have

$$\begin{aligned} KL(P_{\text{true}} \| F_f(x)) &\geq \frac{1}{2\kappa} (TV(P_{\text{true}} \| F_f)) + \frac{\kappa}{4} \\ &\geq \frac{1}{2\kappa} [TV(P_{\text{true}} \| U) - TV(F_f(x) \| U)] + \frac{\kappa}{4}. \end{aligned} \quad (26)$$

Taking expectation on  $P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}$ , we have

$$\begin{aligned} \text{GError}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}}(f) &= \mathbb{E}_{(x,y) \sim P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} KL(P_{\text{true}} \| F_f(x)) + H(P_{\text{true}}) \\ &\geq \frac{1}{2\kappa} \mathbb{E}_{(x,y) \sim P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} [TV(P_{\text{true}} \| U) - TV(F_f(x) \| U)] + \frac{\kappa}{4} + \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} H(P_{\text{true}}) \end{aligned} \quad (27)$$

Applying Lemma 1 and Lemma 2,

$$\begin{aligned} \text{GError}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}}(f) &= \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} KL[P_{\text{true}} \| F_f(x)] + \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} H(P_{\text{true}}) \\ &\geq \frac{1}{2\kappa} \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} [TV(P_{\text{true}} \| U) - TV(F_f(x) \| U)] + \frac{\kappa}{4} + \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} H(P_{\text{true}}) \\ &\geq \frac{1}{2\kappa} \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} TV(P_{\text{true}} \| U) - \frac{1}{2\kappa} \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} TV(F_f \| U) + \frac{\kappa}{4} + \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} H(P_{\text{true}}) \\ &\geq \frac{1}{2\kappa} \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} TV(P_{\text{true}} \| U) - \frac{1}{2\kappa} \mathbb{E}_{P_{\mathcal{X}}^{\text{SEM}}} TV(F_f \| U) \\ &\quad - \frac{1}{2\kappa} d_{\mathcal{F}}(P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}, P_{\mathcal{X}}^{\text{SEM}}) - \frac{\lambda}{2\kappa} + \frac{\kappa}{4} + \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} H(P_{\text{true}}) \\ &\geq \frac{1}{2\kappa} \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} TV(P_{\text{true}} \| U) - \frac{1}{2\kappa} \mathbb{E}_{P_{\mathcal{X}}^{\text{SEM}}} \sqrt{\frac{1}{2} (\mathcal{L}_{\text{reg}}(f) - \log K)} \\ &\quad - \frac{1}{2\kappa} d_{\mathcal{F}}(P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}, P_{\mathcal{X}}^{\text{SEM}}) - \frac{1}{2\kappa} \lambda + \frac{4}{\kappa} + \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} H(P_{\text{true}}). \end{aligned} \quad (28)$$

Given the fact that  $P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}$ ,  $P_{\text{true}}$  are both fixed,  $H(P_{\text{true}})$  and  $TV(P_{\text{true}} \| U)$  are constants for each  $x$ . Finally, we get

$$\text{GError}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}}(f) \geq C - \frac{1}{2\kappa} \mathbb{E}_{P_{\mathcal{X}}^{\text{SEM}}} \sqrt{\frac{1}{2} (\mathcal{L}_{\text{reg}}(f) - \log K)} - \frac{1}{2\kappa} d_{\mathcal{F}}(P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}, P_{\mathcal{X}}^{\text{SEM}}), \quad (29)$$

where  $C = \frac{1}{2\kappa} (\mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} TV(P_{\text{true}} \| U) - \lambda + 8) + \mathbb{E}_{P_{\mathcal{X}\mathcal{Y}}^{\text{COV}}} H(P_{\text{true}})$  is a constant.  $\square$

## B Experimental Details

### B.1 Datasets details

**ID datasets  $P^{\text{ID}}$ .** ID datasets are chosen following common practice in OOD detection. We use CIFAR-10, CIFAR-100 and ImageNet-200 as  $P^{\text{ID}}$ . ImageNet-200 is a subset of the original ImageNet-1K introduced by [20, 5].

**Auxiliary OOD datasets  $P_{\text{train}}^{\text{SEM}}$ .** For CIFAR experiments, we use ImageNet-RC and TIN-597 as auxiliary datasets. ImageNet-RC is a down-sampled variant of the ImageNet-1K, which consists

of 1000 classes and 1,281,167 images. We also conduct experiments on TIN-597 as an alternative for ImageNet-RC. TIN-597 is introduced by recent work [5]. The resolutions of ID and auxiliary samples are both  $64 \times 64$ . For ImageNet experiments, we use a subset of ImageNet-1K consisting of 200 classes as ID datasets. The remaining images belong to other 800 classes are utilized as auxiliary datasets. The resolutions of ID and auxiliary images are both  $224 \times 224$ .

**OOD detection test datasets**  $P_{\text{test}}^{\text{SEM}}$ . In CIFAR experiments, following standard practice [8], we use SVHN [40], Textures [42], Places365 [41], iSUN [44], LSUN-C and LSUN-R [43] to evaluate the OOD detection performance.

- The SVHN test set comprises 26,032 color images of house numbers.
- Textures (Describable Textures Dataset, DTD) consists of 5,640 images depicting natural textures.
- Places365 dataset consists scenic images of 365 different categories. Each class consists of 900 images.
- The iSUN dataset is a subset of the SUN database with 8,925 images.
- The Large-scale Scene Understanding dataset (LSUN) comprises a testing set with 10,000 images of 10 different scenes. LSUN offers two datasets, LSUN-C and LSUN-R. In LSUN-C, the original high-resolution images are randomly cropped into  $32 \times 32$ . Meanwhile, in LSUN-R, the images are resized to  $32 \times 32$ . In ImageNet experiments, we follow the settings of [5], where OpenImage-O [54], SSB-hard [55], Textures [42], iNaturalist [45] and NINCO [47] are selected as OOD detection test datasets.
- OpenImage-O contains 17632 manually filtered images and is  $7.8 \times$  larger than the ImageNet dataset.
- SSB-hard is selected from ImageNet-21K. It consists of 49K images and 980 categories.
- Textures (Describable Textures Dataset, DTD) consists of 5,640 images depicting natural textures.
- iNaturalist consists of 859000 images from over 5000 different species of plants and animals.
- NINCO consists with a total of 5879 samples of 64 classes which are non-overlapped with ImageNet-1K.

**OOD generalization test datasets**  $P^{\text{COV}}$ . Following previous work [4], we corrupt the original test data with Gaussian noise of zero mean and variance of 5 in the main paper. In appendix, we conduct additional experiments involving CIFAR10-C, CIFAR100-C and ImageNet-C [56] with 15 diverse types of noise.

## B.2 Implementation details

### B.2.1 CIFAR experiments.

We use WideResNet-40-10 [57] as the backbone network, which comprises 40 layers. The widen factor is set to 10. We use SGD optimizer to train all methods with dropout strategy. The dropout rate is 0.3. The momentum is set to 0.9 and weight decay is set to 0.0005.

**Pretraining details.** The pretrained model is obtained by training WideResNet-40-10 for 200 epochs with an initial learning rate of 0.1. We decay the learning rate by a factor of 0.2 at the 60-th, 120-th, and 160-th epochs. Batch size is set to 128.

**Finetuning details.**

- For Entropy and EBM, we finetune the pretrained model for 20 epochs with an initial learning rate of 0.001, utilizing a cosine annealing strategy to adjust the learning rate. Following the official implementation, the weight of OOD detection regularization term is set to 0.5 and 0.1 for Entropy and EBM (finetune) respectively. The hyperparameters  $m_{\text{ID}}$  and  $m_{\text{OOD}}$  in EBM regularization learning are set to -25 and -7 respectively. The ID batch size is 128 and the OOD batch size is set to 256.
- For SCONE, we finetune the pretrained model for 10 epochs with an initial learning rate of 0.0002, utilizing a cosine annealing strategy to adjust the learning rate. The batch size is 32, the OOD batch size is 64. The margin of the OOD detection boundary is set to 1. To be aligned with most previous works in OOD detection and generalization, we assume  $P_{\mathcal{X}}^{\text{COV}}$  is unavailable during finetuning.
- For WOODS, we finetune the pretrained model for 10 epochs with an initial learning rate of 0.0002, utilizing a cosine annealing strategy to adjust the learning rate. The ID batch size is 32, the OOD batch size is 64. Other hyperparameters settings are consistent with SCONE.
- For DUL,  $\alpha_0$  is set to 12. While finetuning on CIFAR10, the  $m_{\text{ID}}$  and  $m_{\text{OOD}}$  are set to 10 and 30 respectively. The weight  $\lambda, \gamma$  are set to 0.3 and 2. We train for 20 epochs with an initial learning rate of 0.00005, utilizing a cosine annealing strategy to adjust the learning rate. While finetuning on CIFAR100/TIN-597, the  $m_{\text{ID}}$  and  $m_{\text{OOD}}$  are set to 10 and 30 respectively. The weights  $\lambda, \gamma$  are set to 0.05 and 2 respectively. We finetune for 20 epochs with an initial learning rate of 0.00005, utilizing a cosine annealing strategy to adjust the learning rate. While finetuning on CIFAR100/ImageNet-RC, we set  $h_0 = 0$ . The  $m_{\text{ID}}$  and  $m_{\text{OOD}}$  are set to -430 and -370 respectively. We train for 30 epochs with an initial learning rate of 0.0001, utilizing a cosine annealing strategy to adjust the learning rate.

The weights  $\lambda, \gamma$  are set to 0.1 and 1 respectively. For CIFAR-100/ImageNet-RC, we set  $\tau = 2$  and otherwise  $\tau = 1$ . For DUL<sup>†</sup>, we use Thompson sampling strategy [11] for OOD informativeness mining. The sampling hyperparameters are consistent with that of POEM.

**Training from scratch details.** For POEM, we train from scratch for 200 epochs with an initial learning rate of 0.1, and decay the learning rate by a factor of 0.2 at the 60-th, 120-th, and 160-th epochs following [57]. The ID and OOD batch size are set to 128 and 256 respectively. Following the official implementation, the pool of outliers consists of randomly selected 400,000 samples from auxiliary datasets, and only 50,000 samples (same size as the ID training set) are selected for training based on the boundary score. For DPN, we train for 200 epochs with an initial learning rate of 0.1, and decay the learning rate by a factor of 0.2 at the 60-th, 120-th, and 160-th epochs. The Dirichlet parameters  $\alpha$  are calculated by performing ReLU plus one on the model’s outputs, i.e.,  $\alpha = \text{ReLU}(f(x)) + 1$ .  $\alpha_0$  is set to 15 and 12 respectively when training on CIFAR10 and CIFAR100. The auxiliary datasets are ImageNet and TIN-597. The ID and OOD batch size are set to 128 and 256 respectively. When training on CIFAR100/TIN-597, the OOD regularization weight  $\lambda$  is set to 0.05. In other cases,  $\lambda$  is set to 0.5.

## B.2.2 ImageNet experiments.

We use ResNet18 [58] as the backbone network. We use SGD optimizer to train all the models. The momentum is set to 0.9.

**Pretraining details.** The pretrained model is obtained by training ResNet18 for 100 epochs with an initial learning rate of 0.1, utilizing a cosine annealing strategy to adjust the learning rate. The weight decay is set to 0.0001. Batch size is set to 64.

**Finetuning details.** For Energy regularized learning, we finetune the pretrained model for 10 epochs with an initial learning rate of 0.001, utilizing a cosine annealing strategy to adjust the learning rate. The weight decay is set to 0.0001. Following the official implementation, the weights of OOD detection regularization term are set to 0.1. Specifically, the  $m_{\text{ID}}$  and  $m_{\text{OOD}}$  in energy regularization method are set to -25 and -7 respectively. The ID batch size is 64 and the OOD batch size is set to 128. For Entropy, we finetune the pretrained model for 10 epochs with an initial learning rate of 0.001, utilizing a cosine annealing strategy to adjust the learning rate. The weight decay is set to 0.0001. The ID and OOD batch size are set to 64 and 128 respectively. Following the official implementation, the weights of OOD detection regularization term are set to 0.5. For SCONE, we finetune the pretrained model for 10 epochs with an initial learning rate of 0.0002, utilizing a cosine annealing strategy to adjust the learning rate. The weight decay is set to 0.0005. The batch size is 32, the OOD batch size is 64. The margin of the OOD detection boundary is set to 1. To be aligned with most previous works in OOD detection and generalization, we assume  $P_x^{\text{COV}}$  is unavailable during finetuning. For WOODS, we finetune the pretrained model for 10 epochs with an initial learning rate of 0.0002, utilizing a cosine annealing strategy to adjust the learning rate. The batch size is 32, the OOD batch size is 64. Other hyperparameters of WOODS are consistent with SCONE. For our DUL, the Dirichlet parameters  $\alpha$  are calculated by performing ReLU and exp operation on the model’s outputs, i.e.,  $\alpha = \exp(\text{ReLU}(f(x)))$ . For numerical stability on large scale benchmark, we measure the distributional uncertainty by the strength of Dirichlet distribution.  $\lambda, \gamma$  are set to 0.1 and 4 respectively. We set  $\tau = 1$  in large-scale ImageNet experiments.

**Training from scratch details.** For DPN, we train ResNet18 for 100 epochs with an initial learning rate of 0.1, utilizing a cosine annealing strategy to adjust the learning rate. The weight decay is set to 0.0001. Batch size is set to 64. The Dirichlet parameters  $\alpha$  are calculated by performing ReLU plus one on the model’s outputs, i.e.,  $\alpha = \text{ReLU}(f(x)) + 1$ . The ID classification loss is set to KL-divergence between predicted class distribution under Dirichlet prior and target distribution because of the inconvenience of directly setting  $\alpha_0$ . The target distribution is obtained by label smoothing strategy with parameter of 0.01 [9]. The weight of regularization term applied on OOD auxiliary samples is 1.

---

**Algorithm 1:** Pseudo Code of Decoupled Uncertainty Learning (DUL)

---

**Input :** ID data  $P^{\text{ID}}$ , auxiliary outliers  $P_{\text{train}}^{\text{SEM}}$ , classifier  $f_{\theta_0}$  pretrained on  $P^{\text{ID}}$ .**Output :** finetuned classifier  $f_{\theta}$ 

- 1 Initialize  $\theta = \theta_0$ ;
  - 2 **for each iteration do**
  - 3     Obtain ID sample  $(x, y)$  from  $P^{\text{ID}}$  and auxiliary outlier  $\tilde{x}$  from  $P_{\text{train}}^{\text{SEM}}$ ;
  - 4     Update model parameters  $\theta$  by minimizing objective defined in Eq. 12;
  - 5 **end**
- 

## C Additional Results

### C.1 Uncertainty estimation.

We add Gaussian noise with zero mean and varying variance  $\epsilon$  on CIFAR-10 and investigate the estimated distributional uncertainty and overall uncertainty. Distributional uncertainty is measured by differential entropy. It clear that with DUL regularization, the prediction yields a low overall uncertainty and high distributional uncertainty on covariate-shifted data. We conduct experiments on CIFAR-10/ImageNet-RC and CIFAR-10/TIN-597, tabular results are shown in Tab. 4.

Table 4: Mean value of estimated uncertainty on CIFAR-10-C with varying severity of Gaussian noise with zero mean and variance of  $\epsilon$ .

$\mathcal{P}_x^{\text{in}}/\mathcal{P}_x^{\text{aux}}$	Uncertainty type	DUL	$\epsilon = 0.0$	$\epsilon = 2.0$	$\epsilon = 4.0$	$\epsilon = 6.0$	$\epsilon = 8.0$	$\epsilon = 10.0$
CIFAR-10	Distributional uncertainty	$\times$	-21.33	-18.94	-15.62	-13.71	-12.91	-12.85
		$\checkmark$	-21.23	-19.42	-16.96	-15.47	-14.85	-14.58
ImageNet-RC	Total uncertainty	$\times$	0.04	0.47	1.31	1.93	2.20	2.28
		$\checkmark$	0.03	0.17	0.48	0.77	0.98	1.14
CIFAR-10	Distributional uncertainty	$\times$	-20.94	-20.14	-18.29	-16.23	-14.71	-13.72
		$\checkmark$	-21.48	-20.68	-19.21	-17.70	-16.57	-15.78
TIN-597	Total uncertainty	$\times$	0.06	0.15	0.51	1.06	1.55	1.92
		$\checkmark$	0.04	0.08	0.18	0.34	0.51	0.68

### C.2 Time-consuming comparison.

We compare the time-cost of proposed DUL to other training-required OOD detection methods in Tab. 5. We run all the experiments on one single NVIDIA GeForce RTX-3090 GPU. Compared with other OOD detection methods in a finetune manner, DUL does not introduce noticeably extra cost of computation.

Table 5: Average execution times (s) per epoch of training required OOD detection methods. Compare to other OOD detection methods, DUL does not introduce noticeable computational cost.

Method	CIFAR-10/ImageNet-RC	CIFAR-10/TIN-597	CIFAR-100/ImageNet-RC	CIFAR-100/TIN-597
EBM (finetune)	354.87	229.12	355.70	112.47
Entropy	355.98	140.88	1108.43	120.85
DPN	842.05	75.62	841.69	80.98
POEM	615.51	483.04	825.87	500.01
WOODS	808.77	291.37	906.67	282.50
SCONE	911.83	183.57	1169.33	160.00
DUL	329.58	101.08	925.62	100.43
DUL*	598.68	597.82	544.26	595.25

### C.3 Full results with standard deviation.

Full results with standard deviation are presented this section. In CIFAR experiments, we report the mean and standard deviation in 5 random runs. In ImageNet experiments, we report the mean and

standard deviation in 3 random runs to be consist with [5]. CIFAR experimental results are shown in Tab. 6. Large-scale ImageNet results are shown in Tab. 7.

Table 6: OOD detection and generalization performance comparison with standard variance. Marginal **improvement** and **degradation** ( $\geq 0.5$ ) compare to the baseline method MSP are highlighted in blue or red respectively. The **best** and second best results are in bold or underlined. DUL is the only method achieves state-of-art OOD detection performance (mostly the best or second best) without trade-offs on generalization i.e., the value of entire row is either blue or black.

$\mathcal{P}^{\text{ID}}/\mathcal{P}^{\text{SEM}}_{\text{train}}$	Method	ID/OOD generalization		FPR ↓	OOD detection	
		ID-Acc. ↑	OOD-Acc. ↑		AUROC ↑	AUPR ↑
CIFAR-10 Only	MSP	96.11 $\pm$ 0.09	87.35 $\pm$ 0.58	41.96 $\pm$ 3.85	89.28 $\pm$ 1.12	68.00 $\pm$ 2.19
	EBM (pretrain)	96.11 $\pm$ 0.09	87.35 $\pm$ 0.58	32.45 $\pm$ 3.45	89.34 $\pm$ 1.21	75.22 $\pm$ 2.67
	Maxlogits	96.11 $\pm$ 0.09	87.35 $\pm$ 0.58	32.90 $\pm$ 3.51	89.26 $\pm$ 1.21	74.47 $\pm$ 2.55
	Mahalanobis	96.11 $\pm$ 0.09	87.35 $\pm$ 0.58	32.53 $\pm$ 9.61	93.93 $\pm$ 2.68	74.96 $\pm$ 7.47
CIFAR-10 ImageNet-RC	Entropy	96.04 $\pm$ 0.14	72.57 $\pm$ 3.87	6.63 $\pm$ 0.80	98.72 $\pm$ 0.14	94.00 $\pm$ 1.00
	EBM (pretrain)	<b>96.10</b> $\pm$ 0.23	79.03 $\pm$ 2.53	3.61 $\pm$ 0.71	98.39 $\pm$ 0.39	94.88 $\pm$ 0.91
	POEM	94.32 $\pm$ 0.14	78.89 $\pm$ 2.25	3.32 $\pm$ 0.41	<b>98.99</b> $\pm$ 0.17	<b>99.38</b> $\pm$ 0.12
	DPN	95.69 $\pm$ 0.17	85.52 $\pm$ 0.51	4.28 $\pm$ 0.60	98.53 $\pm$ 0.17	94.93 $\pm$ 0.60
	WOODS	96.01 $\pm$ 0.16	80.14 $\pm$ 1.69	7.12 $\pm$ 1.54	98.40 $\pm$ 0.21	92.92 $\pm$ 0.96
	SCONE	95.96 $\pm$ 0.08	78.80 $\pm$ 1.57	7.02 $\pm$ 1.06	98.45 $\pm$ 0.12	92.46 $\pm$ 0.93
	DUL (ours)	96.02 $\pm$ 0.07	<b>88.01</b> $\pm$ 0.54	5.89 $\pm$ 0.35	98.47 $\pm$ 0.12	92.44 $\pm$ 1.14
	DUL <sup>†</sup> (ours)	<u>96.04</u> $\pm$ 0.03	<u>87.52</u> $\pm$ 0.70	5.99 $\pm$ 0.25	98.28 $\pm$ 0.11	<u>98.40</u> $\pm$ 0.36
CIFAR-10 TIN-597	Entropy	95.94 $\pm$ 0.00	80.51 $\pm$ 0.68	11.60 $\pm$ 0.82	97.93 $\pm$ 0.15	92.16 $\pm$ 0.50
	EBM (pretrain)	95.38 $\pm$ 0.13	83.67 $\pm$ 1.41	19.36 $\pm$ 1.92	87.51 $\pm$ 1.53	83.63 $\pm$ 1.73
	POEM	95.44 $\pm$ 0.18	83.17 $\pm$ 1.39	24.34 $\pm$ 2.48	86.83 $\pm$ 1.13	<u>94.25</u> $\pm$ 0.53
	DPN	94.39 $\pm$ 0.38	79.23 $\pm$ 2.95	17.27 $\pm$ 1.07	94.92 $\pm$ 0.65	87.67 $\pm$ 0.88
	WOODS	95.57 $\pm$ 0.64	83.12 $\pm$ 1.71	7.58 $\pm$ 0.52	<b>98.29</b> $\pm$ 0.04	93.39 $\pm$ 0.39
	SCONE	95.19 $\pm$ 0.77	84.68 $\pm$ 1.44	8.02 $\pm$ 0.92	98.22 $\pm$ 0.08	93.08 $\pm$ 0.30
	DUL (ours)	<b>96.06</b> $\pm$ 0.08	<u>87.93</u> $\pm$ 0.62	6.87 $\pm$ 0.82	98.21 $\pm$ 0.12	91.29 $\pm$ 1.18
	DUL <sup>†</sup> (ours)	<u>95.94</u> $\pm$ 0.09	<b>88.10</b> $\pm$ 0.27	10.34 $\pm$ 0.34	97.67 $\pm$ 0.09	<b>98.59</b> $\pm$ 0.24
CIFAR-100 Only	MSP	80.99 $\pm$ 0.16	55.95 $\pm$ 1.38	74.63 $\pm$ 2.43	80.19 $\pm$ 1.65	42.59 $\pm$ 2.79
	EBM (pretrain)	80.99 $\pm$ 0.16	55.95 $\pm$ 1.38	67.42 $\pm$ 4.35	82.67 $\pm$ 1.82	49.35 $\pm$ 4.00
	Maxlogits	80.99 $\pm$ 0.16	55.95 $\pm$ 1.38	69.32 $\pm$ 3.97	82.30 $\pm$ 1.79	47.60 $\pm$ 3.68
	Mahalanobis	80.99 $\pm$ 0.16	55.95 $\pm$ 1.38	61.51 $\pm$ 3.62	85.97 $\pm$ 1.22	56.10 $\pm$ 3.22
CIFAR-100 ImageNet-RC	Entropy	80.21 $\pm$ 0.09	45.48 $\pm$ 0.78	22.29 $\pm$ 1.32	95.33 $\pm$ 0.28	82.34 $\pm$ 1.11
	EBM (finetune)	80.53 $\pm$ 0.22	48.14 $\pm$ 0.33	13.47 $\pm$ 0.43	96.78 $\pm$ 0.13	87.84 $\pm$ 0.86
	POEM	78.15 $\pm$ 0.18	42.18 $\pm$ 2.34	9.89 $\pm$ 0.36	<b>97.79</b> $\pm$ 0.12	<b>98.40</b> $\pm$ 0.08
	DPN	78.90 $\pm$ 0.25	50.14 $\pm$ 0.36	18.36 $\pm$ 0.82	95.42 $\pm$ 0.17	74.45 $\pm$ 18.40
	WOODS	80.69 $\pm$ 0.30	54.38 $\pm$ 4.42	38.15 $\pm$ 12.91	92.01 $\pm$ 3.23	71.79 $\pm$ 7.98
	SCONE	80.80 $\pm$ 0.30	<b>56.73</b> $\pm$ 4.66	47.60 $\pm$ 14.73	89.61 $\pm$ 3.75	65.29 $\pm$ 9.66
	DUL (ours)	<b>81.30</b> $\pm$ 0.19	<u>56.27</u> $\pm$ 1.82	12.49 $\pm$ 0.22	95.24 $\pm$ 0.09	86.72 $\pm$ 0.76
	DUL <sup>†</sup> (ours)	<u>81.22</u> $\pm$ 0.23	<u>56.07</u> $\pm$ 0.54	<u>11.75</u> $\pm$ 1.69	95.33 $\pm$ 0.79	<u>96.45</u> $\pm$ 0.42
CIFAR-100 TIN-597	Entropy	80.15 $\pm$ 0.17	46.25 $\pm$ 1.42	26.88 $\pm$ 2.06	93.50 $\pm$ 0.36	79.81 $\pm$ 1.31
	EBM (finetune)	79.94 $\pm$ 0.27	50.00 $\pm$ 0.93	26.87 $\pm$ 1.15	91.68 $\pm$ 0.45	80.08 $\pm$ 0.76
	POEM	78.68 $\pm$ 0.13	52.53 $\pm$ 1.06	32.71 $\pm$ 0.96	91.30 $\pm$ 0.68	<u>94.65</u> $\pm$ 0.49
	DPN	78.44 $\pm$ 0.22	47.67 $\pm$ 0.28	25.02 $\pm$ 2.19	<u>93.55</u> $\pm$ 0.19	81.63 $\pm$ 1.27
	WOODS	79.26 $\pm$ 2.28	53.13 $\pm$ 2.97	36.71 $\pm$ 7.92	92.15 $\pm$ 2.35	73.42 $\pm$ 4.72
	SCONE	79.53 $\pm$ 1.82	52.70 $\pm$ 0.96	35.60 $\pm$ 9.50	92.47 $\pm$ 2.19	73.58 $\pm$ 5.21
	DUL (ours)	<b>80.85</b> $\pm$ 0.24	<u>56.19</u> $\pm$ 1.53	<u>23.32</u> $\pm$ 1.11	<b>94.48</b> $\pm$ 0.35	80.82 $\pm$ 1.62
	DUL <sup>†</sup> (ours)	<u>80.50</u> $\pm$ 0.25	<b>56.22</b> $\pm$ 1.29	<b>22.75</b> $\pm$ 0.88	90.88 $\pm$ 0.28	<b>96.33</b> $\pm$ 0.09

#### C.4 Results of different types of corruption.

We conduct additional experiments on CIFAR10-C, CIFAR100-C and ImageNet-C with 15 different types of corruption. The results validate that the proposed method can improve the overall performance under different types of corruption.

#### C.5 OOD detection results on individual datasets.

We provide OOD detection results of DUL on each individual OOD detection test dataset in Tab. 8 and Tab. 9, based on the checkpoint with random seed 1.

Table 7: OOD detection and generalization performance comparison with standard variance. Substantially **improvement** and **degradation** ( $\geq 0.5$ ) compare to baseline method w.r.t. MSP are highlighted in blue or red respectively. The **best** and **second best** results are in bold or underlined. Similar with CIFAR experiments, DUL establishes strong OOD detection performance (always the best or second best) without degraded generalization i.e., the entire row is either blue or black.

$\mathcal{P}^{ID}/\mathcal{P}_{train}^{SEM}$	Method	ID/OOD Generalization		OOD Detection		
		ID-Acc. $\uparrow$	OOD-Acc. $\uparrow$	FPR $\downarrow$	AUROC $\uparrow$	AUPR $\uparrow$
ImageNet-200 ImageNet-800	MSP	<u>85.15</u> $\pm 0.33$	74.84 $\pm 0.47$	58.23 $\pm 1.54$	86.98 $\pm 0.24$	82.27 $\pm 0.32$
	EBM (pretrain)	<u>85.15</u> $\pm 0.33$	74.84 $\pm 0.47$	51.94 $\pm 0.82$	88.18 $\pm 0.11$	84.75 $\pm 0.08$
	Maxlogits	<u>85.15</u> $\pm 0.33$	74.84 $\pm 0.47$	<u>51.62</u> $\pm 0.20$	88.30 $\pm 0.09$	84.71 $\pm 0.07$
	Entropy	84.92 $\pm 0.30$	74.75 $\pm 0.52$	<u>53.62</u> $\pm 0.76$	<u>89.05</u> $\pm 0.01$	<u>85.02</u> $\pm 0.08$
	EBM (finetune)	<u>84.14</u> $\pm 0.11$	<u>73.31</u> $\pm 0.57$	<u>59.73</u> $\pm 0.83$	<u>87.54</u> $\pm 0.03$	<u>82.81</u> $\pm 0.16$
	DPN	84.87 $\pm 0.30$	74.40 $\pm 0.90$	<u>63.84</u> $\pm 0.70$	87.18 $\pm 0.18$	<u>80.69</u> $\pm 0.35$
	WOODS	84.99 $\pm 0.62$	<u>74.98</u> $\pm 0.46$	<u>51.71</u> $\pm 2.84$	<u>88.30</u> $\pm 0.56$	<u>84.80</u> $\pm 0.98$
	SCONE	84.93 $\pm 0.71$	74.91 $\pm 0.49$	<u>52.52</u> $\pm 3.54$	<u>88.19</u> $\pm 0.41$	<u>84.50</u> $\pm 1.08$
	DUL (ours)	<b>85.65</b> $\pm 0.07$	<b>75.59</b> $\pm 0.12$	<b>49.14</b> $\pm 0.13$	<b>89.27</b> $\pm 0.03$	<b>85.62</b> $\pm 0.03$

Table 8: OOD detection results of DUL on each individual OOD detection test dataset.

$\mathcal{P}^{ID}/\mathcal{P}_{train}^{SEM}$	Method	LSUN-crop		Places365		LSUN-resize		iSUN		Texture		SVHN	
		FPR $\downarrow$	AUROC $\uparrow$	FPR $\downarrow$	AUROC $\uparrow$	FPR $\downarrow$	AUROC $\uparrow$	FPR $\downarrow$	AUROC $\uparrow$	FPR $\downarrow$	AUROC $\uparrow$	FPR $\downarrow$	AUROC $\uparrow$
CIFAR-10	DUL	6.75	98.75	15.55	96.34	0.00	99.63	0.00	99.57	3.20	98.89	8.35	98.24
ImageNet-RC	DUL $\dagger$	12.79	98.08	14.99	96.11	0.00	99.49	0.00	99.45	1.72	98.79	4.84	98.47
CIFAR-10	DUL	0.90	99.47	22.40	95.21	0.10	99.52	0.30	99.45	8.00	97.94	5.90	98.52
TIN-597	DUL $\dagger$	3.48	99.25	32.81	91.48	0.00	99.78	0.00	99.78	13.97	97.15	9.69	98.28
CIFAR-100	DUL	44.65	83.73	21.05	79.30	0.00	99.63	0.00	99.61	3.05	98.53	5.00	98.15
ImageNet-RC	DUL $\dagger$	47.75	81.29	14.92	95.08	0.00	99.58	0.00	99.46	1.70	98.70	4.44	98.02
CIFAR-100	DUL	7.05	98.69	65.35	83.60	3.80	99.03	3.50	98.98	34.45	92.21	21.15	95.81
TIN-597	DUL $\dagger$	6.25	98.55	78.95	64.09	0.02	99.84	0.00	99.83	39.79	85.02	5.61	98.51

Table 9: OOD detection results of DUL on each OOD detection test dataset.

$\mathcal{P}^{ID}/\mathcal{P}_{train}^{SEM}$	Method	OpenImage-O		SSB-hard		Textures		iNaturalist		NINCO	
		FPR $\downarrow$	AUROC $\uparrow$	FPR $\downarrow$	AUROC $\uparrow$	FPR $\downarrow$	AUROC $\uparrow$	FPR $\downarrow$	AUROC $\uparrow$	FPR $\downarrow$	AUROC $\uparrow$
ImageNet-200/800	DUL	49.68	91.31	72.40	80.60	30.76	92.98	33.21	94.76	59.92	86.61

Table 10: Classification error rate comparison on CIFAR10-C. ID dataset is CIFAR10.

Method	$\mathcal{P}_{train}^{SEM}$	Noise			Blur				Weather				Digital				Avg.
		Gauss.	Shot	Impul.	Defoc.	Glass	Motion	Zoom	Snow	Frost	Fog	Brit.	Contr.	Elastic	Pixel	JPEG	
MSP	None	53.35	39.98	44.07	15.61	42.35	19.33	19.83	14.39	18.02	9.61	5.09	18.19	13.45	22.71	18.61	23.64
EBM (pretrain)		53.35	39.98	44.07	15.61	42.35	19.33	19.83	14.39	18.02	9.61	5.09	18.19	13.45	22.71	18.61	23.64
Maxlogits		53.35	39.98	44.07	15.61	42.35	19.33	19.83	14.39	18.02	9.61	5.09	18.19	13.45	22.71	18.61	23.64
Mahalanobis		53.35	39.98	44.07	15.61	42.35	19.33	19.83	14.39	18.02	9.61	5.09	18.19	13.45	22.71	18.61	23.64
Entropy	ImageNet-RC	67.20	53.99	70.93	17.09	81.99	20.00	22.17	22.41	31.78	11.43	5.82	16.99	14.28	28.92	21.37	32.42
EBM (finetune)		62.49	50.00	73.01	21.01	73.93	20.06	25.80	21.02	27.00	11.55	5.60	15.08	15.15	31.46	22.69	31.72
DPN		50.19	38.80	59.81	16.98	52.81	19.68	22.24	18.25	20.11	11.10	5.74	18.01	14.41	24.30	18.76	26.08
POEM		47.52	38.16	63.39	22.64	67.37	24.12	28.43	23.44	27.33	14.30	7.65	19.02	18.37	31.87	22.23	30.40
WOODS		62.36	49.09	60.77	15.56	74.37	19.05	20.16	18.79	25.41	9.84	5.45	16.76	14.16	24.07	20.45	29.09
SCONE		63.29	50.01	61.96	15.61	77.16	18.65	20.24	19.77	26.43	9.82	5.54	16.88	14.09	24.46	20.64	29.64
DUL (Ours)		53.46	40.11	43.71	15.62	42.77	19.49	19.89	14.27	18.07	9.82	5.06	18.49	13.56	22.65	18.69	23.71
Entropy		TIN-597	65.68	52.08	56.49	18.68	51.96	22.73	24.43	18.57	24.64	10.56	5.51	16.38	15.88	25.19	41.25
EBM (finetune)	58.95		44.89	52.91	18.30	47.25	23.99	24.33	17.70	22.18	12.23	6.07	19.78	17.11	26.85	51.31	29.59
DPN	48.07		38.05	42.56	22.54	47.39	27.29	30.41	20.46	26.17	14.53	7.22	24.54	18.33	27.40	26.23	28.08
POEM	51.06		39.56	47.20	16.08	48.97	20.23	21.40	17.36	22.15	11.29	5.94	18.08	14.87	22.08	31.74	25.87
WOODS	60.01		46.74	51.26	18.04	49.83	20.95	23.17	17.34	22.66	10.02	5.69	16.02	16.32	23.55	25.92	27.17
SCONE	58.20		44.92	50.36	17.14	48.47	20.64	22.80	17.05	21.15	9.87	6.07	17.24	16.17	24.89	25.53	26.70
DUL (Ours)	54.56		40.84	44.18	15.67	43.61	19.93	19.74	14.49	18.34	9.84	5.18	19.24	13.95	22.91	18.66	24.08

Table 11: Classification error rate comparison on CIFAR100-C. ID dataset is CIFAR100.

Method	$\mathcal{P}_{\text{train}}^{\text{SEM}}$	Noise				Blur				Weather				Digital			Avg.
		Gauss.	Shot	Impul.	Defoc.	Glass	Motion	Zoom	Snow	Frost	Fog	Brit.	Contr.	Elastic	Pixel	JPEG	
MSP	None	76.98	67.70	74.51	36.08	78.49	40.69	41.39	40.52	46.66	31.44	22.59	40.29	35.31	42.83	44.93	48.03
EBM (pretrain)		76.98	67.70	74.51	36.08	78.49	40.69	41.39	40.52	46.66	31.44	22.59	40.29	35.31	42.83	44.93	48.03
Maxlogits		76.98	67.70	74.51	36.08	78.49	40.69	41.39	40.52	46.66	31.44	22.59	40.29	35.31	42.83	44.93	48.03
Mahalanobis		76.98	67.70	74.51	36.08	78.49	40.69	41.39	40.52	46.66	31.44	22.59	40.29	35.31	42.83	44.93	48.03
Entropy	ImageNet-RC	81.69	72.89	86.99	38.19	88.94	42.20	44.62	45.50	53.90	33.14	24.40	39.77	37.19	47.62	48.82	52.39
EBM (finetune)		82.33	73.51	90.58	39.20	90.71	41.43	44.40	47.29	54.88	34.26	24.72	39.15	37.65	51.47	49.29	53.39
DPN		77.60	68.59	83.59	39.74	86.77	43.33	44.91	47.01	53.69	36.37	26.11	43.45	37.56	47.48	44.47	52.04
POEM		83.65	76.63	88.02	44.09	90.05	44.29	46.44	52.80	60.84	39.49	28.90	43.35	42.49	56.06	54.50	56.77
WOODS		77.51	68.48	77.47	36.38	80.62	40.92	41.93	41.27	47.08	31.36	23.30	39.74	36.21	42.87	46.16	48.75
SCONE		76.43	67.30	75.32	36.23	77.00	40.92	41.59	40.22	45.53	31.35	23.09	40.10	35.87	42.42	45.42	47.92
DUL (Ours)		77.23	67.95	75.13	35.83	78.52	39.74	40.64	39.76	46.13	30.96	22.43	39.17	34.81	43.03	44.64	47.73
Entropy	TIN-597	83.34	75.44	79.77	38.02	82.74	42.22	44.17	44.52	53.36	32.80	23.67	38.00	37.85	44.55	62.91	52.22
EBM (finetune)		81.24	72.78	78.14	37.55	80.00	42.94	43.64	43.70	51.11	33.52	23.92	40.19	38.13	45.05	73.27	52.34
DPN		81.81	72.87	79.09	39.61	81.25	44.66	46.21	45.42	52.40	34.28	25.64	40.89	39.15	47.66	76.40	53.82
POEM		78.88	70.32	74.68	38.59	77.95	43.41	44.27	43.00	50.22	34.64	25.37	43.09	37.28	45.43	81.55	52.58
WOODS		81.14	72.39	76.49	38.50	79.00	43.23	44.59	42.37	49.21	33.03	24.32	40.23	39.01	41.19	47.80	50.17
SCONE		80.90	72.30	77.19	37.91	78.70	42.36	43.99	42.81	49.49	32.19	24.00	39.26	37.96	41.80	48.34	49.95
DUL (Ours)		77.01	67.67	74.16	36.16	78.35	40.79	41.32	40.17	46.33	31.42	22.87	40.79	35.34	42.74	45.05	48.01

Table 12: Classification error rate comparison on ImageNet-C. Here we test compared methods on a subset of the original ImageNet-C consisting of 200 classes. ID dataset is ImageNet-200.

Method	$\mathcal{P}_{\text{train}}^{\text{SEM}}$	Noise				Blur				Weather				Digital			Avg.
		Gauss.	Shot	Impul.	Defoc.	Glass	Motion	Zoom	Snow	Frost	Fog	Brit.	Contr.	Elastic	Pixel	JPEG	
MSP	None	52.2	70.7	71.5	56.0	55.5	52.6	54.1	67.3	67.0	63.5	56.5	53.3	46.3	50.3	51.9	57.9
EBM (pretrain)		52.2	70.7	71.5	56.0	55.5	52.6	54.1	67.3	67.0	63.5	56.5	53.3	46.3	50.3	51.9	57.9
Maxlogits		52.2	70.7	71.5	56.0	55.5	52.6	54.1	67.3	67.0	63.5	56.5	53.3	46.3	50.3	51.9	57.9
Entropy	ImageNet-800	51.3	71.2	71.7	54.4	54.6	51.8	53.2	66.9	66.2	62.7	55.8	51.4	45.1	49.7	51.1	57.1
EBM (finetune)		52.9	72.2	72.8	56.0	56.2	53.5	54.1	67.8	67.4	64.0	57.0	52.2	46.5	51.4	52.9	58.5
DPN		51.9	69.2	69.6	56.5	55.0	52.0	53.1	65.5	65.3	62.3	55.5	54.5	46.0	49.7	51.4	57.2
WOODS		51.4	69.4	70.0	55.2	54.8	51.9	52.9	66.2	66.0	62.5	55.6	52.6	45.6	49.7	51.3	57.0
SCONE		51.6	69.4	70.0	55.4	55.0	52.1	53.1	66.3	66.0	62.6	55.7	53.0	45.8	49.9	51.4	57.1
DUL (Ours)		51.1	69.1	70.5	55.1	54.5	51.6	52.4	66.2	65.9	62.6	55.7	53.0	45.4	49.4	50.9	56.9

Table 13: Comprehensive comparison involving 15 different types of corruption from commonly-used domain adaption benchmark [52]. Substantial ( $\geq 0.5$ ) **improvement** and **degradation** compared to the baseline MSP [6] are highlighted in blue or red respectively. DUL is the only method that achieves SOTA OOD detection performance without sacrificing generalization i.e., the value of the entire row is almost black or blue. The **best** or **second best** results are highlighted in bold or underlined. MD is the shorthand of Mahalanobis.

Method	$\mathcal{P}_{\text{train}}^{\text{ID}}$	$\mathcal{P}_{\text{train}}^{\text{SEM}}$	OOD generalization (Error rate ↓)													OOD detection				
			Gauss.	Shot	Impul.	Defoc.	Glass	Motion	Zoom	Snow	Frost	Fog	Brit.	Contr.	Elast.	Pixel	JPEG	Avg.	FPR↓	AUC↑
MSP	CIFAR-10	None	77.0	67.7	74.5	36.1	78.5	40.7	41.4	40.5	46.7	31.4	22.6	40.3	35.3	42.8	44.9	48.0	42.0	89.3
EBM			77.0	67.7	74.5	36.1	78.5	40.7	41.4	40.5	46.7	31.4	22.6	40.3	35.3	42.8	44.9	48.0	42.0	89.3
Maxlogits			77.0	67.7	74.5	36.1	78.5	40.7	41.4	40.5	46.7	31.4	22.6	40.3	35.3	42.8	44.9	48.0	42.0	89.3
MD		77.0	67.7	74.5	36.1	78.5	40.7	41.4	40.5	46.7	31.4	22.6	40.3	35.3	42.8	44.9	48.0	42.0	89.3	
Entropy		ImageNet-RC	81.7	72.9	87.0	38.2	88.9	42.2	44.6	45.5	53.9	33.1	24.4	39.8	37.2	47.6	48.8	52.4	6.6	98.7
EBM (FT)			82.3	73.5	90.6	39.2	90.7	41.4	44.4	47.3	54.9	34.3	24.7	39.2	37.7	51.5	49.3	53.4	3.6	98.4
DPN			77.6	68.6	83.6	39.7	86.8	43.3	44.9	47.0	53.7	36.4	26.1	43.5	37.6	47.5	44.5	52.0	4.3	98.5
POEM			83.7	76.6	88.0	44.1	90.1	44.3	46.4	52.8	60.8	39.5	28.9	43.4	42.5	56.1	54.5	56.8	3.3	99.0
WOODS			77.5	68.5	77.5	36.4	80.6	40.9	41.9	41.3	47.1	31.4	23.3	39.7	36.2	42.9	46.2	48.8	7.1	98.5
SCONE			<b>76.4</b>	<b>67.3</b>	<b>75.3</b>	<b>36.2</b>	<b>77.0</b>	<b>40.9</b>	<b>41.6</b>	<b>40.2</b>	<b>45.5</b>	<b>31.4</b>	<b>23.1</b>	40.1	<b>35.9</b>	<b>42.4</b>	<b>45.4</b>	47.9	7.0	98.5
DUL (Ours)			<b>77.2</b>	<b>68.0</b>	<b>75.1</b>	<b>35.8</b>	<b>78.5</b>	<b>39.7</b>	<b>40.6</b>	<b>39.8</b>	<b>46.1</b>	<b>31.0</b>	<b>22.4</b>	<b>39.2</b>	<b>34.8</b>	43.0	<b>44.6</b>	<b>47.7</b>	<b>5.9</b>	<b>98.5</b>
Entropy		TIN-597	83.3	75.4	79.8	38.0	82.7	42.2	44.2	44.5	53.4	32.8	23.7	38.0	37.9	44.6	62.9	52.2	11.6	97.9
EBM (FT)			81.2	72.8	78.1	37.6	80.0	42.9	43.6	43.7	51.1	33.5	23.9	40.2	38.1	45.1	73.3	52.3	19.4	87.5
DPN			81.8	72.9	79.1	39.6	81.3	44.9	46.2	45.4	52.4	34.3	25.6	40.9	39.2	47.7	76.4	53.8	17.3	94.9
POEM			<b>78.9</b>	<b>70.3</b>	<b>74.7</b>	<b>38.6</b>	<b>78.0</b>	43.4	44.3	43.0	50.2	34.6	25.4	43.1	<b>37.3</b>	45.4	81.6	52.6	34.3	86.8
WOODS	81.1		72.4	76.5	38.6	79.0	43.2	44.6	42.4	49.2	33.0	24.3	40.2	39.0	<b>41.2</b>	47.8	50.2	<b>7.6</b>	<b>98.3</b>	
SCONE	80.9		72.3	77.2	37.9	78.7	42.4	43.4	42.8	49.5	32.2	24.0	<b>39.3</b>	38.0	<b>41.8</b>	48.3	<b>50.0</b>	<b>8.0</b>	<b>98.2</b>	
DUL (Ours)	<b>77.0</b>		<b>67.7</b>	<b>74.2</b>	<b>36.2</b>	<b>78.4</b>	<b>40.8</b>	<b>41.3</b>	<b>40.2</b>	<b>46.3</b>	<b>31.4</b>	<b>22.9</b>	<b>40.8</b>	<b>35.3</b>	42.7	<b>45.1</b>	<b>48.0</b>	<b>6.9</b>	<b>98.2</b>	
MSP	ImageNet-200	None	52.2	70.7	71.5	56.0	55.5	52.6	54.1	67.3	67.0	63.5	56.5	53.3	46.3	50.3	51.9	57.9	58.2	82.3
EBM			52.2	70.7	71.5	56.0	55.5	52.6	54.1	67.3	67.0	63.5	56.5	53.3	46.3	50.3	51.9	57.9	58.2	82.3
Maxlogits			52.2	70.7	71.5	56.0	55.5	52.6	54.1	67.3	67.0	63.5	56.5	53.3	46.3	50.3	51.9	57.9	58.2	82.3
Entropy		ImageNet-800	51.3	71.2	71.7	<b>54.4</b>	<b>54.6</b>	<b>51.8</b>	53.2	66.9	66.2	62.7	55.8	<b>51.4</b>	<b>45.1</b>	<b>49.7</b>	<b>51.1</b>	57.1	53.6	89.1
EBM (FT)			52.9	72.2	72.8	56.0	<b>56.2</b>	<b>53.5</b>	54.1	67.8	67.4	64.0	57.0	<b>52.2</b>	46.5	51.4	<b>52.9</b>	58.5	59.7	87.5
DPN			51.9	<b>69.2</b>	<b>69.6</b>	56.5	55.6	52.4	53.1	<b>65.5</b>	<b>65.0</b>	<b>62.2</b>	<b>55.5</b>	<b>54.5</b>	46.0	<b>49.7</b>	51.4	<b>57.2</b>	63.8	87.2
WOODS			51.4	69.4	<b>70.0</b>	55.2	54.8	51.9	<b>52.9</b>	<b>66.2</b>	66.0	<b>62.5</b>	<b>55.6</b>	52.6	45.6	<b>49.7</b>	51.3	<b>57.0</b>	<b>51.7</b>	88.3
SCONE			51.6	69.4	<b>70.0</b>	55.4	55.0	52.1	53.1	66.3	66.0	62.6	55.7	53.0	45.8	49.9	51.4	57.1	52.5	88.2
DUL (Ours)			<b>51.0</b>	<b>69.1</b>	<b>70.5</b>	<b>55.1</b>	<b>54.5</b>	<b>51.5</b>	<b>52.4</b>	<b>66.2</b>	<b>65.9</b>	62.6	55.7	53.0	<b>45.4</b>	<b>49.4</b>	<b>50.9</b>	<b>56.9</b>	<b>49.1</b>	<b>89.3</b>

## C.6 Empirical evidence

Here we provide empirical supports to our intuition about energy-based OOD detection regularization. We calculate the entropy of predicted distribution before and after finetuning with Energy regularization [8]. The results show can support our claim in Section 4.

Table 14: Predictive entropy of predicted distribution on covariate-shifted OOD dataset before and after finetuning with energy-based OOD detection regularization [8].

$\mathcal{P}^{\text{ID}}$	$\mathcal{P}_{\text{train}}^{\text{SEM}}$	Before	After
CIFAR10	ImageNet-RC	0.11	1.05
CIFAR10	TIN-597	0.11	0.14
CIFAR100	ImageNet-RC	0.92	3.39
CIFAR100	TIN-597	0.92	1.15

## D Discussions

### D.1 Math derivation

**Differential entropy.** The proposed DUL calculates differential entropy as OOD detection measurement. Here we detail how to calculate the differential entropy of a Dirichlet distribution. The following derivation of differential entropy is taken from [9]. The differential entropy of a Dirichlet parameterized by  $\alpha$  is calculated by

$$\begin{aligned}
 h[p(\mu|x)] &= - \int_S p(\mu|x) \ln(p(\mu|x)) d\mu \\
 &= \sum_k^K \ln \Gamma(\alpha_k) - \ln \Gamma(\alpha_0) - \sum_k^K (\alpha_k - 1) \cdot (\psi(\alpha_k) - \psi(\alpha_0))
 \end{aligned}
 \tag{30}$$

where  $\alpha_0$  is the strength of Dirichlet, i.e.,  $\alpha_0 = \sum_K \alpha_k$ .  $\alpha_k$  denotes the  $k$ -th element in  $\alpha$ .  $\Gamma$  is the Gamma function and  $\psi$  is the digamma function. Here we provide a PyTorch implementation on how to calculate distribution uncertainty measured by differential entropy.

```

def diff_entropy(alphas):
    alpha0 = torch.sum(alphas, dim=1)
    return torch.sum(
        torch.lgamma(alphas) - (alphas - 1) * (torch.digamma(alphas) -
                                                torch.digamma(alpha0)).
        unsqueeze(1)),
        dim=1) - torch.lgamma(alpha0)

logits = model(x)
alpha = torch.ReLU(logits)+1
diff_entropy = diff_entropy(alpha)

```

We refer interested readers to [9] and Gal’s PhD Thesis [59] for more detailed math derivations.

### D.2 Discussion about Disparity Discrepancy

In section 4, we claim that a limited disparity discrepancy between test-time semantic OOD and covariate-shifted OOD is practical. Here we provide some empirical evidence and discussion to support such a claim.

**The key challenge of OOD detection lies in identifying ID-like semantic OOD.** As mentioned in recent works [38], effectively distinguishing between the most challenging OOD samples that are much like in-distribution (ID) data is the core challenge of OOD detection. Recent works regularize models on ID-like OOD to enhance the OOD detection performance. Since the ID-like semantic OOD samples are more difficult to be detected and more informative. For example, NTOM [17] and



POEM [11] utilizes greedy and Thompson sampling strategies to find semantic OOD samples which are more closely to ID. [38] proposes to explicitly discover outliers near ID by prompt learning.

**Semantic OOD and covariate OOD can be very similar in practice.** As shown in Fig. 4. There exists many similar samples from semantic OOD and covariate OOD in large-scale commonly used benchmarks. We borrow some examples from recent works [47] to show case.



(a) Semantic OOD classes from **NINCO**



(b) ID classes from **ImageNet**

Figure 4: Semantic OOD samples can be very similar to ID.

### D.3 Social Impact

AI safety and trustworthiness are closely related to our work. This paper presents work to harmonize the conflicts between out-of-distribution detection methods and model generalization. The proposed method puts effort to enhance machine learning models for their safely deployment on out-of-distribution data, avoiding both undesirable behavior and degraded performance in challenging high-stake tasks. However, due to the bias from data used by current OOD detection benchmark, e.g., large-scale ImageNet, the ones using the proposed method need to carefully consider the selection of auxiliary outliers for safety-critical applications.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: In conclusion and Appendix.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The paper provide the full set of assumptions and a complete (and correct) proof in Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Section 6 and Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: See Appendix and supplemental material.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The paper provides above details in Section 6 and Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: In Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics [https://neurips.cc/public/EthicsGuidelines?](https://neurips.cc/public/EthicsGuidelines)

Answer: [Yes]

Justification: The authors have carefully reviewed the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: In Appendix.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to

generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

#### 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

#### 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The license and terms of use explicitly mentioned and properly respected.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

#### 13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

#### 14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

#### 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.