

Uncertainty-Aware Optimal Transport for Semantically Coherent Out-of-Distribution Detection

Fan Lu^{1,*} Kai Zhu^{1,*,\ddagger} Wei Zhai¹ Kecheng Zheng² Yang Cao^{1,3,\ddagger}

¹ University of Science and Technology of China ² Ant Group

³ Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

{lufan@mail., zkzy@mail., wzhai056@}ustc.edu.cn zkcloud@gmail.com forrest@ustc.edu.cn

Abstract

Semantically coherent out-of-distribution (SCOOD) detection aims to discern outliers from the intended data distribution with access to unlabeled extra set. The co-existence of in-distribution and out-of-distribution samples will exacerbate the model overfitting when no distinction is made. To address this problem, we propose a novel uncertainty-aware optimal transport scheme. Our scheme consists of an energy-based transport (ET) mechanism that estimates the fluctuating cost of uncertainty to promote the assignment of semantic-agnostic representation, and an inter-cluster extension strategy that enhances the discrimination of semantic property among different clusters by widening the corresponding margin distance. Furthermore, a T-energy score is presented to mitigate the magnitude gap between the parallel transport and classifier branches. Extensive experiments on two standard SCOOD benchmarks demonstrate the above-par OOD detection performance, outperforming the state-of-the-art methods by a margin of 27.69% and 34.4% on FPR@95, respectively. Code is available at <https://github.com/LuFan31/ET-ODD>.

1. Introduction

Deep learning models trained in close-set world often suffer from performance degradation in real-world scenarios due to the interference of out-of-distribution (OOD) inputs—unknown samples whose classes don’t overlap with those already seen during training [33]. Full-supervised models tend to make overconfident predictions on OOD inputs [20, 21], which severely limits the application in high-risk fields such as autonomous driving [7] and medical analysis [22]. To this end, OOD detection [10, 17, 18, 28, 31] aims to obtain a reliable model to identify these unknown

*Co-first Author. †Corresponding Author. ‡Work done during an internship at Ant Group.

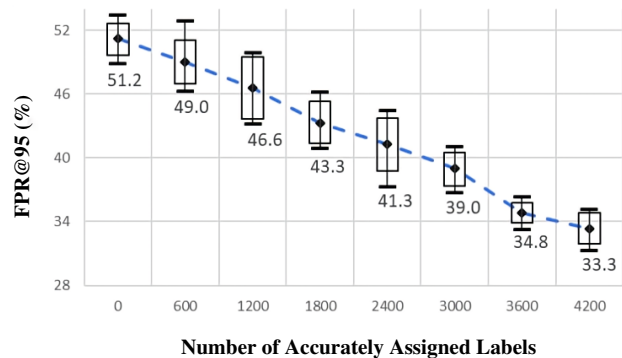


Figure 1. **The effect of the number of accurately assigned labels.** We gradually increase the ground-truth labels of unlabeled ID samples during the baseline [11] training and report the FPR@95% (lower value is better) results averaged by multiple tests. The trend shows that more accurately assigned labels enhance the OOD detection capability of the model.

samples as abnormal ones and reject them at test time.

To learn the discrepancy between ID/OOD samples better, a rich line of OOD detection methods like OE [11] and MCD [31] utilizing extra OOD data during training has been developed. OE flattens the prediction probability of OOD samples among all ID classes, and MCD maximizes the entropy discrepancy between ID and OOD samples outputted by two parallel classifiers. Although extra data greatly enhances the performance of models across traditional OOD benchmarks, these benchmarks still face two problems: 1) to ensure the sampling effectiveness, ID samples in extra unlabeled datasets are manually removed to obtain the pure OOD set, which does not correspond to reality since ID samples are inevitably mixed into the unlabeled set and expensive to be purified. 2) the great improvement on them is most likely due to the overfitting on the training OOD samples [28], so the model paying attention to the lower-level covariate shifts between datasets has difficulty generalizing to other OOD domains, especially when the scale of the dataset is not large.

To overcome these challenges, the SCOOD benchmarks




Figure 2. **Motivation of the proposed method.** When an extra unlabeled dataset is introduced into training, the resulting covariate shifts will easily confuse the minimum feature distance (e.g., *Euclidean distance*) from the correct class mean with other wrong classes. The left histogram reflects the minimum Euclidean distance distribution from the class mean for the ‘dog’ images in both CIFAR-10 (blue) and tiny-ImageNet (green). Unfortunately, although belonging to the same class, their feature distance distributions differ greatly, which limits the ability of cluster-based methods to mine unlabeled ID semantic knowledge in the SCOOD setting. In contrast, as shown in the right histogram, the distributions of the ‘dog’ images in the two different datasets maintain good consistency on the energy metric, facilitating a more stable semantic assignment process.

[28] are proposed to urge OOD detection models to focus on distinguishing semantic discrepancy between ID/OOD samples and avoid overfitting low-level covariate shifts on different data sources. K-means clustering is adopted to dynamically filter ID samples from the unlabeled set during training, enhancing the optimization for ID classification [28]. Unfortunately, we statistically observe that this strategy can only assign correct semantic labels to less than 1% unlabeled ID samples at each epoch, which means a large number of samples are added to the ID classification training with wrong labels and undoubtedly limits the ability of the model to fully understand the semantic discrepancy between ID/OOD samples.

To investigate the impact of accurately assigned semantic labels on SCOOD tasks, we gradually increase the ground-truth labels of unlabeled ID samples used in training with the baseline in [28], and then observe the variation trend of model performance shown in Fig. 1. The model performance consistently improves as the correctly assigned labels increase, leading the key question of how to allocate as many ID-related samples in the unlabeled set correctly as possible. Our further analysis indicates that compared to the traditional feature distance (e.g., *Euclidean distance*), the class-specific energy score remains consistent even when the assigned cluster contains shifted unlabeled samples with an ID class. As shown in Fig. 2, due to the covariate shifts of the same class samples from another dataset with a black circle, the minimum distance with the correct class mean is easily confused with other clusters. In contrast, the energy-based uncertainty prior can effectively mitigate the corresponding interference owing to the intra-class interaction.

To effectively involve energy-based guidance in the assignment process, we propose an uncertainty-aware optimal transport scheme. This core energy-based transport

(ET) mechanism estimates the energy transport cost to encourage lower-cost samples to be assigned to the same cluster and the higher-cost ones to be split uniformly among all clusters, leading to a semantic-consistent distribution. Specifically, the energy score is obtained by performing the aggregation operation in logit space.

To further enhance the assignment efficiency of the aforementioned transport mechanism, an inter-cluster extension strategy is proposed to narrow the margin of the same ID class while widening the margin between ID and OOD samples. An enhanced global feature representation will be mapped to the lower-dimensional logit space, from where we can obtain a more discriminative prediction distribution for each cluster and class-specific energy score better reflecting the semantic information of samples. By utilizing both of them, the semantics among samples assigned to different clusters by ET will be more discriminative. During inference, we also choose the energy score with temperature scaling produced by the classifier as the detection metric, aligning with the transport optimization more closely.

Our contributions are summarized as follows: **1)** We analyze the limitations of the clustering strategy for the SCOOD task and further explore how to address the fundamental challenge of OOD detection. **2)** We propose a novel uncertainty-aware optimal transport scheme to fully utilize the semantic consistency hidden in the unlabeled set, resulting in the covariate-invariant assignment. **3)** Extensive experiments demonstrate that the proposed method achieves state-of-the-art performance on SCOOD benchmarks.

2. Related Work

2.1. OOD Detection with Extra Data

The mainstream OOD detection works primarily focus on detecting samples with semantic shifts, i.e. OOD samples drawn from classes that do not overlap with known classes.

To enhance the ability to understand ID/OOD discrepancy, some methods synthesize OOD training data such as [6, 15, 23, 26], and another series of OOD detection methods introduce extra OOD data into training. OE [11] uses purified unlabeled OOD data, which is encouraged to produce a uniform softmax distribution during training. MCD [31] maximizes the entropy discrepancy between ID and extra OOD samples outputted by two parallel classifiers. The survey [29] indicates that OOD detection with external data generally performs better. However, massive unlabeled samples could be very easily obtained considering their accessibility, and ID data would inevitably exist among them. Hence, ‘purified unlabeled OOD data’ means the mixed ID data need to be manually removed, which brings a high cost and is impractical.

Against the above issues, UDG [28] proposes a realistic benchmark that retains ID samples in the unlabeled set, namely SCOOD (Semantically Coherent Out-of-Distribution Detection) benchmarks. UDG separates the ID/OOD samples in the unlabeled dataset based on K-means clustering, and adds the ID samples to the labeled samples according to semantics for training together. During training, UDG focuses on the semantic shifts between samples, which is a fundamental problem to be solved in OOD detection tasks, and well explores the semantic knowledge in the unlabeled set. UDG underlines that the OOD detection task should concentrate on semantic shifts, but in experiments we found that the clustering method cannot sufficiently collect unlabeled ID samples and can only assign less than 1% correct semantic labels to unlabeled ID samples at each epoch. In this paper, we propose an uncertainty-aware optimal transport scheme to achieve more accurate label assignment and learn a more significant difference between ID and OOD semantics.

2.2. Optimal Transport

wGAN [1] applies optimal transport (OT) theory to the field of computer vision and has received extensive attention. Its basic idea is to minimize the Wasserstein distance between the distribution of sampled data and the image distribution synthesized by deep generative models. In summary, OT is to transform one distribution into another with minimal cost (e.g., Wasserstein distance), so it is applied to distribution matching tasks such as few-shot classification [8] and domain adaptation [16, 25]. In close-set unsupervised classification tasks, OT can assign pseudo-labels for unlabeled samples. SeLa [2] implements equipartition constraint on the label assignment matrix \mathbf{Q} , and introduces KL divergence as a regularization term into the optimization objective, then uses a fast version of the Sinkhorn-Knopp [5] algorithm to assign labels. SwAV [3] keeps the soft-label assignment produced by the Sinkhorn-Knopp algorithm without approximating it into a one-hot

assignment like SeLa. For the SCOOD task, we adopt an individual OT strategy called energy-based transport (ET) mechanism, which introduces the energy score as the transport cost to guide the cluster distribution of samples.

3. Method

3.1. Problem Setting

In this section, we will introduce the setting of SCOOD task. Firstly, we denote in-distribution and out-of-distribution set as I and O , respectively. The training set D on SCOOD benchmarks consists of labeled training set D_L and unlabeled training set D_U . While the samples contained in D_L are all from I , the D_U is a mixture of partial ID set D_U^I from I and OOD set D_U^O from O , which is unlike previous OOD benchmarks. So the training set can be represented as $D = (D_L \subset I) \cup (D_U^I \subset I) \cup (D_U^O \subset O)$. Likewise, the test set T comprises $T^I \subset I$ and $T^O \subset O$. We desire to assign correct ID labels to as many samples in D_U^I as possible during training, and identify data from T^O as negative (OOD) samples and classify samples from T^I correctly at test time.

3.2. Overall Pipeline

As shown in Fig. 3, we design a novel uncertainty-aware optimal transport (OT) scheme, which consists of an energy-based transport (ET) mechanism and an inter-cluster extension strategy L_{rep} to assign correct ID class label to partial unlabeled ID samples and then involve them into the joint training. The ET introduces the energy score as uncertainty and estimates the uncertainty-based transport cost to guide the cluster distribution of all samples. To further promote the discrimination in logit space, from which the uncertainty can significantly reflect the difference between ID/OOD samples, the inter-cluster extension strategy enhances the global feature representation mixed with ID and OOD samples and then the enhanced representation will be mapped into more discriminate logits.

3.3. Energy-Based Transport Mechanism

Unlike assigning labels to unlabeled samples in the close-set classification tasks, the situation we face is more challenging because we have to overcome the interference of OOD samples when assigning labels. To address it, we propose an energy-based transport (ET) mechanism based on logit space to explore the semantic discrepancy and knowledge hidden in the unlabeled set more sufficiently.

Specifically, we add a K -dimensional OT head h_{ot} in parallel with the classification head h_{cls} to map training samples to K clusters. Given the set of all N samples $\mathcal{X} = \{x_1, x_2, \dots, x_N\}$ and the set of their feature extracted by encoder $\mathcal{Z} = \{z_1, z_2, \dots, z_N\}$, we denote the probability of the i -th sample over all clusters as

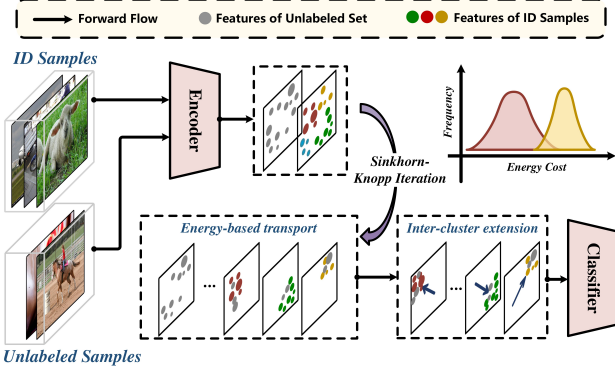


Figure 3. **Overall framework of our uncertainty-aware optimal transport scheme.** The energy scores for both ID and OOD samples are firstly calculated from output logit and then utilized to estimate the energy-based transport cost, guiding the assignment of unlabeled ID samples during the Sinkhorn-Knopp iteration. Finally, the energy scores will be further enhanced by the unlabeled samples obtaining correct ID labels and the inter-cluster extension strategy, facilitating subsequent assignment and classification.

$p(\mathbf{c}|x_i) = \text{softmax}(h_{ot}(z_i))$, where \mathbf{c} is a K -dimension vector representing K clusters. Then we can define the cost matrix in optimal transport problem as $\mathbf{P} \in \mathbb{R}^{K \times N}$ and $P_{ji} = p(c_j|x_i)$ denotes the probability of x_i belonging to c_j , where c_j is the j -th cluster. Likewise, we represent $\mathbf{Q} \in \mathbb{R}^{K \times N}$ as the assignment matrix and $Q_{ji} = q(c_j|x_i)$ denotes the posterior probability of x_i assigned to c_j . Note that the assignment matrix \mathbf{Q} just represents the assignment of a certain cluster to each sample, rather than directly assigning labels for classification training. When we represent the distributions of N samples and K clusters using an N -dimension vector β and a K -dimensional vector α , respectively, all feasible solutions of assignment matrix \mathbf{Q} in *transportation polytope* [5] can be formulated as:

$$U(\alpha, \beta) := \{\mathbf{Q} \in \mathbb{R}^{K \times N} \mid \mathbf{Q}\mathbf{1}_N = \alpha, \mathbf{Q}^\top \mathbf{1}_K = \beta\}, \quad (1)$$

where $\mathbf{1}$ is an all-ones vector of corresponding dimension, and α and β are the marginal projections of matrix \mathbf{Q} onto its rows and columns, respectively.

Energy score from logit space is an effective and readily available metric to distinguish ID/OOD samples [18], so we introduce the energy $\mathbf{e} \in \mathbb{R}^N$ as uncertainty guidance into our ET module. The energy of x_i distributed across all clusters is shown as:

$$e_i = \log \sum_{j=1}^K e^{l(c_j|x_i)}, \quad (2)$$

where $l(\mathbf{c}|x_i) = h_{ot}(z_i)$ denotes the logit of x_i belonging to c_j . The energy \mathbf{e} reflects the aggregate probability distribution of samples across clusters. Afterward, ET will encourage samples with higher energy scores to be assigned

to the same cluster, while the ones with lower energy scores, which means they have larger uncertainty in the cluster distribution, will tend to be split uniformly among the K clusters. And the matrix \mathbf{Q} based on \mathbf{e} can be denoted as:

$$\alpha = \frac{1}{K} \cdot \mathbf{1}_K, \quad \beta = \frac{\mathbf{e}}{\sum_{i=1}^N e_i}. \quad (3)$$

And the energy transport cost can be written as:

$$\mathbf{P}_{\text{en}} = \mathbf{P} \cdot \mathbf{e} \in \mathbb{R}^{K \times N}, \quad (4)$$

where \mathbf{e} is first broadcast into a $K \times N$ matrix and then multiplied element-wise by \mathbf{P} . Hence, the Wasserstein distance between α and β is defined as $\text{OT}(\alpha, \beta) = \min_{\mathbf{Q} \in \Pi(\alpha, \beta)} -\langle \mathbf{Q}, \mathbf{P}_{\text{en}} \rangle$, where $\langle \cdot, \cdot \rangle$ denotes the Frobenius dot-product. In this case, it is worth noting that \mathbf{P}_{en} cannot be naively regarded as the conventional cost matrix in the optimal transport problem, since the elements in \mathbf{P}_{en} represent the probability of a sample being assigned to a certain cluster. The larger the element value, the lower the cost of transferring the sample to the corresponding cluster. To avoid solving this linear programming problem which needs a large computational cost, we introduce the entropic regularization term $H(\mathbf{Q})$ into Wasserstein distance [5] and express the optimization problem as:

$$\text{OT}(\alpha, \beta) = \min_{\mathbf{Q} \in \Pi(\alpha, \beta)} -\langle \mathbf{Q}, \mathbf{P}_{\text{en}} \rangle + \varepsilon H(\mathbf{Q}), \quad (5)$$

where $\varepsilon > 0$ and $H(\mathbf{Q}) = \sum_{ji} Q_{ji} \log Q_{ji}$. In this way, the optimal \mathbf{Q} has been shown to be written as:

$$\mathbf{Q}^{\frac{1}{\varepsilon}} = \text{Diag}(\mathbf{u}) \mathbf{P}_{\text{en}} \text{Diag}(\mathbf{v}), \quad (6)$$

where exponentiation is element-wise, and \mathbf{u} and \mathbf{v} can be solved much faster by Sinkhorn's Algorithm [5]. The assignment matrix \mathbf{Q} maps N samples onto K clusters.

Since then, we have obtained the cluster indexes of all samples $\mathcal{C} = \{c_1, c_2, \dots, c_N\} \in \{1, 2, \dots, K\}$. We collect unlabeled ID samples and assign pseudo labels to them according to the proportion of each label in a cluster [28]. Specifically, we let samples that belong to k -th cluster at the t -th epoch form the set D_k as denoted in Eq. (7).

$$D_k^{(t)} = \{x_i | c_i^{(t)} = k\}. \quad (7)$$

At epoch t , we define the set of ID-class labels $\mathcal{Y}^{(t)} = \mathcal{Y}_L \cup \mathcal{Y}_{ot}^{(t)} \in \{0, 1, \dots, M-1\}$, where M is the number of ID classes, and \mathcal{Y}_L is formed by the ground-truth labels from D_L , while the ID pseudo labels assigned to unlabeled samples by ET constitute $\mathcal{Y}_{ot}^{(t)}$. Note that $\mathcal{Y}_{ot}^{(t)} = \emptyset$ when $t = 0$ and it will be updated during training, while \mathcal{Y}_L is unchanged. Then we calculate the proportion of samples belonging to class $y \in \mathcal{Y}^{(t)}$ in cluster k as follows:

$$\text{rate}_{k,y}^{(t)} = \frac{|D_{k,y}^{(t)} = \{x_i | c_i^{(t)} = k, y_i = y\}|}{|D_k^{(t)}|}. \quad (8)$$

When $rate_{k,c}^{(t)}$ is over a threshold τ , all unlabeled samples in cluster k will be added into labeled set D_L to form the updated $D_L^{(t)}$ denoted as Eq. (9), and $\mathcal{Y}^{(t)}$ is composed of the labels from $D_L^{(t)}$.

$$D_L^{(t)} = D_L \cup \{x_i | x_i \in D_k^{(t)}, rate_{k,y}^{(t)} > \tau\}. \quad (9)$$

Meanwhile, the remaining unlabeled samples form D_U at epoch t is denoted as $D_U^{(t)}$. With the updated training set, the classification loss $L_{cls}^{(t)}$ and equalization loss $L_{unif}^{(t)}$ can be written as:

$$\begin{aligned} L_{cls}^{(t)} &= - \sum_{x_i \in \mathcal{D}_L^{(t)}} \sum_{y_i \in \mathcal{Y}^{(t)}} y_i \cdot \log(p(\mathbf{y}|x_i)), \\ L_{unif}^{(t)} &= - \sum_{x_i \in \mathcal{D}_U^{(t)}} \frac{1}{M} \cdot \log(p(\mathbf{y}|x_i)), \end{aligned} \quad (10)$$

where $p(\mathbf{y}|x_i)$ means the prediction probability of x_i over all ID classes, and y_i should be converted to a one-hot vector denoting the corresponding label of x_i . Moreover, $\frac{1}{M}$ denotes the uniform posterior distribution over all of M ID classes and $L_{unif}^{(t)}$ will force samples in $D_U^{(t)}$ to uniformly distribute among M ID classes [11].

3.4. Inter-Cluster Extension Strategy

Due to the fact that supervised learning only produces semantic representations that meet the minimum necessary for classification [27], to widen the discrepancy between ID/OOD on energy metric and thus strengthen the ability of energy transport cost in Eq. (4) to guide ID/OOD samples to different clusters, we desire to obtain an improved feature representation by conducting an unsupervised training module called inter-cluster extension strategy.

With labels unavailable, given a batch of data $\{x_{(l,i)}\}_{i=1 \dots B_1}$ and $\{x_{(u,i)}\}_{i=1 \dots B_2}$ in D_L and D_U respectively, we combine them to form a batch of joint data $\{x_i\}_{i=1 \dots B} = \{x_{(l,i)}\} \cup \{x_{(u,i)}\}$, where $B = B_1 + B_2$. Then we obtain two augmentation versions of them denoted as $x_i^0 = A^0(x_i)$ and $x_i^1 = A^1(x_i)$ through two kinds of data augmentations A^0 and A^1 , respectively. Then the augmented data will be fed into an encoder to be extracted feature as:

$$\begin{aligned} z_i^0 &= h_m(e_\theta^0(x_i^0)), \\ z_i^1 &= h_m(e_\theta^1(x_i^1)), \end{aligned} \quad (11)$$

where e_θ^0 and e_θ^1 are two encoders with parameters θ , and h_m is an MLP. Additionally, a memory queue is created to store z_i^1 at different iterations and the n latest batches of z_i^1 compose a dynamic queue. The dynamic queue at iteration t' is described as:

$$Z^{1,(t')} = \{z_i^1\}^{(t')} \cup \{z_i^1\}^{(t'-1)} \dots \cup \{z_i^1\}^{(t'-n+1)}. \quad (12)$$

Considering the absence of labels, we obtain the desired representation by maximizing the cosine similarity of the same samples in $\{z_i^0\}^{(t')}$ and $\{z_i^1\}^{(t')}$, so we employ InfoNCE loss as the objective:

$$L_{rep} = - \sum_{i=1}^N \log \frac{\exp(\cos(z_i^0, z_i^1))}{\sum_{j=1}^{n \cdot B} \exp(\cos(z_i^0, z_j))}, \quad (13)$$

where $z_j \in Z^{1,(t')}$ and \cos denotes cosine similarity. The enhanced representation obtained by L_{rep} will be mapped to the lower-dimensional logit space, where we can obtain the class-specific energy score better reflecting the semantic discrepancy of samples from different classes and assign a more discriminative cluster distribution to them.

Our overall objective can be expressed as Eq. (14) with the wights γ and λ .

$$L = L_{cls}^{(t)} + \gamma L_{unif}^{(t)} + \lambda L_{rep}. \quad (14)$$

3.5. T-Energy Score

Considering the parallel structure of h_{ot} and h_{cls} , which both map the same features to the logit space through different fully-connected layers, there is a certain degree of correlation between the two logits. So we use the energy score from the logit $l(\mathbf{y}|x_i) = h_{cls}(z_i)$ outputted by h_{cls} as OOD score to detect OOD samples after training. The energy can be readily calculated through the $\log\text{sumexp}$ operator as $energy(x_i) = \log \sum_{m=1}^M e^{l(y_m|x_i)}$, where $l(y_m|x_i)$ is the logit of x_i belonging to the ID class y_m , and ID samples have higher energy score. We observe that OOD samples are more concentrated at minimum energy compared to methods training with only ID samples (e.g., [18]) due to the involvement of training OOD data in SCOOD. In this case, a large temperature value T can smooth the energy distribution and thereby widen the discrepancy between ID/OOD on the energy metric. Therefore, we implement the temperature scaling proposed in ODIN [17] to compute energy score using temperature-scaled logits instead of mapping them to softmax score like ODIN, which is consistent with the source of the uncertainty metric in the proposed ET. The T-energy score is formulated as:

$$T\text{-energy}(x_i) = T \cdot \log \sum_{m=1}^M e^{l(y_m|x_i)/T}. \quad (15)$$

4. Experiments

4.1. Benchmarks

We evaluate our method on the realistic SCOOD benchmarks [28] proposed recently. SCOOD benchmarks contain two benchmarks: CIFAR-10 benchmark and CIFAR-100 benchmark which regard CIFAR-10 and CIFAR-100 [13] as the labeled set D_L , respectively. Tiny-ImageNet [14] mixed

Table 1. **Comparison between the previous SOTA methods and ours on the two SCOOD benchmarks.** We report the averaged values on 6 OOD datasets and detailed results are shown in Appendix. \uparrow/\downarrow indicates higher/lower value is better. The best results are in **bold**.

Benchmarks	Method	FPR95 \downarrow	AUROC \uparrow	AUPR-In/Out \uparrow	CCR@FPR \uparrow			
					10^{-4}	10^{-3}	10^{-2}	10^{-1}
CIFAR-10 Benchmark	ODIN [17]	52.00	82.00	73.13 / 85.12	0.36	1.29	6.92	39.37
	EBO [18]	50.03	83.83	77.15 / 85.11	0.49	1.93	9.12	46.48
	OE [11]	50.53	88.93	87.55 / 87.83	13.41	20.25	33.91	68.20
	MCD [31]	73.02	83.89	83.39 / 80.53	5.41	12.3	28.02	62.02
	UDG [28]	36.22	93.78	93.61 / 92.61	13.87	34.48	59.97	82.14
	Ours	8.53	96.47	97.10 / 95.65	40.31	63.95	77.35	86.27
CIFAR-100 Benchmark	ODIN [17]	81.89	77.98	78.54 / 72.56	1.84	5.65	17.77	46.73
	EBO [18]	81.66	79.31	80.54 / 72.82	2.43	7.26	21.41	49.39
	OE [11]	80.06	78.46	80.22 / 71.83	2.74	8.37	22.18	46.75
	MCD [31]	85.14	74.82	75.93 / 69.14	1.06	4.60	16.73	41.83
	UDG [28]	75.45	79.63	80.69 / 74.10	3.85	8.66	20.57	44.47
	Ours	41.05	82.44	84.37 / 76.47	11.70	19.47	34.05	50.97

with ID and OOD data is utilized as extra unlabeled training set D_U on SCOOD benchmarks. In test time, CIFAR-100 is used as one of the OOD datasets for CIFAR-10 benchmark while CIFAR-10 is included in the OOD set of CIFAR-100 benchmark. Additionally, five other datasets including Texture [4], SVHN [19], Tiny-ImageNet [14], LSUN [30] and Places365 [32] are collected to form the complete test set T of the two benchmarks, which contain both ID and OOD samples. SCOOD benchmarks re-split these samples into T^I and T^O according to their true semantics.

4.2. Evaluation Metrics.

Following UDG [28], we use five metrics to evaluate the performance of our method. **FPR@TPR95%** presents the ratio of falsely identified OOD when most (*i.e.*95%) ID samples are correctly recognized. **AUROC** is the area under the receiver operating characteristic curve, evaluating the OOD detection performance. **AUPR-In/Out** calculates the area under the precision-recall curve. AUPR-In/OUT denotes ID/OOD samples as positive. **CCR@FPR n** shows classification accuracy when the ratio of falsely identified OOD equals n , which can evaluate ID classification and OOD detection capability simultaneously.

4.3. Results on SCOOD Benchmarks

We compare the results of our proposed approach with the previous state-of-the-art (SOTA) OOD detection methods in Tab. 1. ODIN [17] and EBO [18] do not require extra unlabeled training data, and OE [11], MCD [31], and UDG [28] use Tiny-ImageNet as unlabeled training data D_U . All experiments use ResNet-18 [9]. We only report the average metric values on 6 OOD test datasets for each benchmark limited by space. Results show that our proposed uncertainty-aware optimal transport scheme

Table 2. **Ablation study on CIFAR-10 benchmark.** We take UDG [28] as baseline method. Our completed method combining ET and L_{rep} is written as ‘Ours’ for short. All experiments are tested in the T-energy score mentioned in Sec. 3.5 for fair comparisons. For brevity, we refer to each experiment by its index.

Strategy	FPR95 \downarrow	AUROC \uparrow	AUPR-In/Out \uparrow	ACC \uparrow
1: UDG [28]	21.57	92.44	91.70 / 92.12	92.28
2: UDG + L_{rep}	17.75	94.77	94.80 / 93.13	93.07
3: ET	11.39	96.05	96.18 / 94.57	93.40
4: Ours	8.53	96.47	97.10 / 95.65	93.71




Figure 4. **The number of accurately assigned labels during training in EXP#1 – #4** shows that the ET has a significant improvement over IDF strategy in assigning exact labels to unlabeled ID samples, and L_{rep} can further increase this advantage. EXP#1 is the UDG baseline, EXP#2 introduces the L_{rep} into UDG, EXP#3 is the proposed ET alone, and EXP#4 is the complete version of our method.

consistently obtains the best results across all metrics, especially FPR@95 is significantly improved on both benchmarks. In particular, although MCD [31] introduces extra unlabeled training data and achieves satisfying results

across traditional OOD benchmarks, its maximization of the entropy discrepancy between ID/OOD samples makes the model pay excessive attention to the low-level covariate shifts, which eventually leads to limited generalization in other OOD sources and failure on the SCOOD task.

4.4. Ablation Study and Qualitative Analysis

Effectiveness of uncertainty-aware optimal transport scheme. Here we analyze the effect of each major component, including the energy-based transport (ET) mechanism and inter-cluster extension strategy L_{rep} in Tab. 2. EXP#1 is the UDG [28] method which we use as the baseline, and then we introduce the L_{rep} into UDG to carry out EXP#2. To evaluate the effectiveness of the proposed ET, we replace the IDF strategy in UDG with ET and implement EXP#3, and EXP#4 which combines ET with the L_{rep} is the complete version of our method.

From the comparison between EXP#2 and EXP#1, it can be seen when L_{rep} is introduced in UDG and combined with IDF strategy [28] based on K-means clustering, the result is improved compared with the baseline. It is intuitive that an enhanced feature representation produced by L_{rep} increases the ID/OOD discrepancy and facilitates clustering-based label assignment in the feature space. However, limited by the non-robust feature distance to covariate shifts and lacking the practical guidance in clustering task, simply applying L_{rep} in SCOOD fails to overcome the impact of covariate shift, so the IDF strategy still has shortcomings in extracting ID semantics from the unlabeled set, as proved in Fig. 4. When we replace the IDF strategy with the ET based on the energy uncertainty and implement the Exp#3, the result shows that the OOD detection ability is enhanced by a large margin, as FPR@95 gets a significant 10.18% improvement than baseline. After combining ET and L_{rep} , the results of EXP#4 show that the performance of the model is at its best when obtaining a more discriminative uncertainty derived from a enhanced representation to promote the discrimination of inter-cluster distribution. In Fig. 4 we also report the comparison among EXP#1 – #4 in the number of accurately assigned labels, we can see from where both the ET and L_{rep} have brought apparent ascension in the aspect of increasing the accuracy of the label assignment. Combined with the results in Tab. 2 and Fig. 4, it is again demonstrated that more accurate semantic label assignment during training will promote the performance of the model, since the model benefits from exploring the semantic discrepancy and knowledge hidden in the unlabeled set more sufficiently.

In summary, the performance of the model boosts mainly due to the contribution of ET, which provides effective guidance for the cluster distribution of samples based on semantics via energy metric. Then L_{rep} further enlarges the advantages over baseline by providing a more discriminate

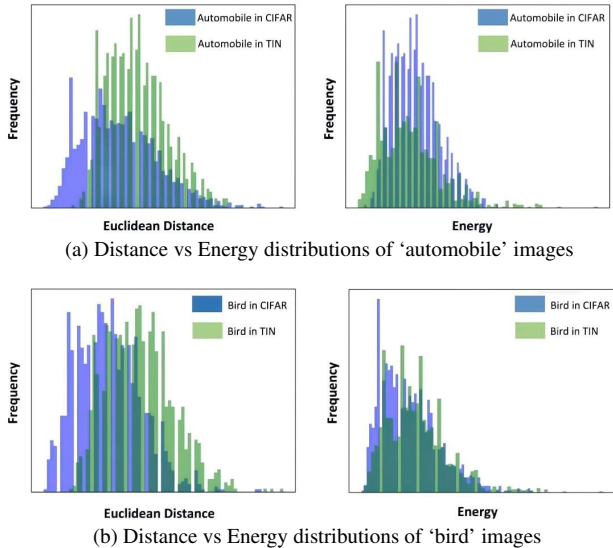


Figure 5. **The effectiveness of energy metric.** In ‘automobile’ and ‘bird’ classes, the energy metric still maintains better consistency with the interference of covariate shifts. CIFAR and TIN denote CIFAR-10 and Tiny-ImageNet datasets, respectively.

Table 3. **The effectiveness of energy in ET.** V-OT and En-OT denotes the methods introducing *variance* and *entropy* as uncertainty like ET, respectively. The performance of our proposed ET outperforms the two methods on all metrics.

	FPR95 ↓	AUROC ↑	AUPR-IN/OUT ↑	ACC ↑
V-OT	13.46	96.11	96.70 / 95.43	93.15
En-OT	10.94	95.86	96.75 / 94.67	93.10
ET	8.53	96.47	97.10 / 95.65	93.71

energy metric between ID/OOD.

The consistency and effectiveness of energy metric in ET. In this work, we propose the energy-based transport (ET) mechanism to explore the semantic discrepancy and knowledge hidden in the unlabeled set more sufficiently. Here, we first demonstrate the robustness of the energy metric to covariate shifts on more classes. We employ a model trained only with CIFAR-10 to output the minimum Euclidean distance and energy distributions for samples belonging to the same class but from different datasets (e.g., CIFAR-10 and Tiny-ImageNet). Except for the distributions of the ‘dog’ images shown in the histograms of Fig. 2, we perform the same experiment with the ‘automobile’ and ‘bird’ images and report the distribution difference in Fig. 5. In ‘automobile’ and ‘bird’ categories, the energy metric still better overcomes the interference of the covariate shifts and maintains the consistency of the semantic compared with feature distance. Afterward, we consider other commonly-used uncertainty metrics to replace the energy metric,

Table 4. **The impact of extra training data.** EXP#1 uses a model trained only on CIFAR-10 (CIFAR) with CE loss, while EXP#2 – #4 introduce Tiny-ImageNet (TIN) as extra unlabeled data and distinguish unlabeled ID/OOD samples in training. Extra training data consistently promotes the OOD detection performance but also brings a drop to classification accuracy (ACC). Our method (including both ET and L_{rep} , written as ‘Ours’ for short) minimizes the degradation on ACC.

D	Strategy	FPR95 ↓	AUROC ↑	AUPR-IN/OUT ↑	ACC ↑
CIFAR	1: CE loss	34.96	85.72	84.58 / 83.24	94.94
CIFAR	2: UDG [28]	21.57	92.44	91.70 / 92.12	92.28
+	3: ET	10.21	96.25	96.92 / 94.76	93.48
TIN	4: Ours	8.53	96.47	97.10 / 95.65	93.71

including the *variance* and *entropy* of samples over cluster distribution, and replace the energy metric in ET with them to guide label assignment. The result in Tab. 3 indicates that our proposed ET outperforms both methods based on other uncertainty, confirming the effective guidance of energy in the assignment process.

Effectiveness of extra data and classification accuracy decay. The comprehensive survey [29] indicates that extra data generally promotes the OOD detection performance of models, but the classification performance of classification-based OOD detection methods will inevitably decrease due to the interferences of OOD samples. To explore the impact of the extra training data on the performance of OOD detection and ID classification, we first evaluate the model only trained on CIFAR-10 with standard cross-entropy loss (CE loss) and then introduce Tiny-ImageNet as extra unlabeled training data and implement EXP#2 – #4. The results are reported in Tab. 4.

By comparing EXP#2 – #4 and EXP#1, we can see when the model is exposed to the extra unlabeled set mixed with ID and OOD and distinguishes the unlabeled ID/OOD samples during training as UDG and the ET do, the OOD detection performance will be comprehensively improved because the model keeps focusing on the semantic shifts between ID/OOD thus learning their discrepancy better. However, since unlabeled ID samples cannot be collected as fully as ideal, some of them are classified into the wrong categories, which hurts the classification performance of the model. Our method brings a more accurate label assignment derived from the practical guidance of energy score on the cluster distribution, and mitigates accuracy decay induced by extra OOD samples compared to UDG.

Influence of OOD scores. We notice that when replacing max softmax probability (MSP) [10] with T-energy as OOD score, UDG drops by about 15% on FPR@95. To exclude the possibility that our approach only gains from T-energy score, we test the performance of UDG and our

Table 5. **Influence of OOD scores.** MSP denotes max softmax probability. Our method outperforms the UDG on all OOD scores and achieves the best on the T-energy.

OOD Score	Method	FPR95 ↓	AUROC ↑	AUPR-IN/OUT ↑
MSP	UDG [28]	36.22	93.78	92.61 / 92.94
	Ours	13.55	95.56	96.89 / 94.96
Energy	UDG [28]	34.90	90.65	91.56 / 91.11
	Ours	12.86	96.05	97.04 / 95.01
T-Energy	UDG [28]	21.57	92.44	91.70 / 92.12
	Ours	8.53	96.47	97.10 / 95.65

approach on three kinds of OOD scores, including MSP, energy and T-energy. As shown in Tab. 5, our method outperforms the baseline no matter which OOD score is chosen. In particular, when evaluating UDG on energy and T-energy, the AUROC and AUPR both suffer a drop. However, the proposed ET explores the ID/OOD semantic discrepancy based on the energy, and in turn promotes the ability of energy score to discriminate ID/OOD by continuously optimizing label assignment, so our method obtains improved performance when being tested with energy score. Additionally, temperature scaling further widens the gap between ID and OOD, helping the model achieves the best performance on all metrics.

5. Conclusion

In this work, we focus on addressing the fundamental challenge of OOD detection tasks and further explore how to fully understand the semantic discrepancy between the ID/OOD samples on SCOOD benchmarks. We reveal that the key to success in the realistic SCOOD task is to allocate as many ID samples in the unlabeled set correctly as possible. To this end, we propose a novel uncertainty-aware optimal transport scheme, which introduces class-specific energy scores as guidance for effective label assignment. Experimental results show that our method achieves better performance than previous state-of-the-art methods on SCOOD benchmarks.

Limitations. In addition to temperature scaling, other techniques such as feature clipping applied in ReAct [24] also enhance the performance of energy score, so how to obtain an OOD score that best fits the SCOOD task can be further explored. Moreover, a setting highly related to SCOOD has been proposed in [12] and formulated as a constrained optimization problem. We will also theoretically analyze these practical OOD settings in our future work.

Acknowledgments. This work is supported by National Key R&D Program of China under Grant 2020AAA0105701 and Ant Group through Ant Research Intern Program.

References

- [1] Martin Arjovsky, Soumith Chintala, and Leon Bottou. Wasserstein gan. In *arXiv preprint arXiv:1701.07875*, 2017. 3
- [2] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *Int. Conf. Learn. Represent.*, 2020. 3
- [3] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Adv. Neural Inform. Process. Syst.*, 2020. 3
- [4] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2014. 6
- [5] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Adv. Neural Inform. Process. Syst.*, 2013. 3, 4
- [6] Xuefeng Du, Zhaoning Wang, Mu Cai, and Yixuan Li. Vos: Learning what you don’t know by virtual outlier synthesis. In *International Conference on Learning Representations*. 3
- [7] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? The kitti vision benchmark suite. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2012. 1
- [8] Dandan Guo, Long Tian, He Zhao, Mingyuan Zhou, and Hongyuan Zha. Adaptive distribution calibration for few-shot learning with hierarchical optimal transport. In *Adv. Neural Inform. Process. Syst.*, 2022. 3
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016. 6
- [10] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *Int. Conf. Learn. Represent.*, 2017. 1, 8
- [11] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. In *Int. Conf. Learn. Represent.*, 2019. 1, 3, 5, 6
- [12] Julian Katz-Samuels, Julia B Nakhleh, Robert Nowak, and Yixuan Li. Training ood detectors in their natural habitats. In *International Conference on Machine Learning*, pages 10848–10865. PMLR, 2022. 8
- [13] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *CiteSeer*, 2009. 5
- [14] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. <http://cs231n.stanford.edu/tiny-imagenet-200.zip>, 2015. 5, 6
- [15] Kimin Lee, Honglak Lee, Kibok Lee, and Jinwoo Shin. Training confidence-calibrated classifiers for detecting out-of-distribution samples. In *6th International Conference on Learning Representations, ICLR 2018*, 2018. 3
- [16] Zhenpeng Li, Zhen Zhao, Yuhong Guo, Haifeng Shen, and Jieping Ye. Mutual learning network for multi-source domain adaptation. In *arXiv preprint arXiv:2003.12944*, 2020. 3
- [17] Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. In *Int. Conf. Learn. Represent.*, 2017. 1, 5, 6
- [18] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. In *Adv. Neural Inform. Process. Syst.*, 2020. 1, 4, 5, 6
- [19] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *Proceedings of NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011. 6
- [20] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 427–436, 2015. 1
- [21] Walter J Scheirer, Anderson de Rezende Rocha, Archana Sapkota, and Terrance E Boulton. Toward open set recognition. volume 35, pages 1757–1772. IEEE, 2012. 1
- [22] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International conference on information processing in medical imaging*, pages 146–157, 2017. 1
- [23] Kumar Sricharan and Ashok Srivastava. Building robust classifiers through generation of confident out of distribution examples. *arXiv preprint arXiv:1812.00239*, 2018. 3
- [24] Yiyu Sun, Chuan Guo, and Yixuan Li. React: Out-of-distribution detection with rectified activations. In *Adv. Neural Inform. Process. Syst.*, 2021. 8
- [25] Rosanna Turrisi, Rémi Flamary, Alain Rakotomamonjy, and Massimiliano Pontil. Multi-source domain adaptation via weighted joint distributions optimal transport. In *Uncertainty in Artificial Intelligence*, pages 1970–1980. PMLR, 2022. 3
- [26] Sachin Vernekar, Ashish Gaurav, Vahdat Abdelzad, Taylor Denouden, Rick Salay, and Krzysztof Czarnecki. Out-of-distribution detection in classifiers via generation. *arXiv preprint arXiv:1910.04241*, 2019. 3
- [27] Jim Winkens, Rudy Bunel, Abhijit Guha Roy, Robert Stanforth, Vivek Natarajan, Joseph R. Ledsam, Patricia MacWilliams, Pushmeet Kohli, Alan Karthikesalingam, Simon Kohl, Taylan Cemgil, S. M. Ali Eslami, and Olaf Ronneberger. Contrastive training for improved out-of-distribution detection. In *arXiv preprint arXiv:2007.05566*, 2020. 5
- [28] Jingkang Yang, Haoqi Wang, Litong Feng, Xiaopeng Yan, Huabin Zheng, Wayne Zhang, and Ziwei Liu. Semantically coherent out-of-distribution detection. In *Int. Conf. Comput. Vis.*, 2021. 1, 2, 3, 4, 5, 6, 7, 8
- [29] Jingkang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized out-of-distribution detection: A survey. In *arXiv preprint arXiv:2110.11334*, 2021. 3, 8
- [30] 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. 6

- [31] Qing Yu and Kiyoharu Aizawa. Unsupervised out-of-distribution detection by maximum classifier discrepancy. In *Int. Conf. Comput. Vis.*, 2019. [1](#), [3](#), [6](#)
- [32] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2017. [6](#)
- [33] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [1](#)