THE DEVIL IS IN THE WRONGLY-CLASSIFIED SAMPLES: TOWARDS UNIFIED OPEN-SET RECOGNITION

Anonymous authors
Paper under double-blind review

ABSTRACT

Open-set Recognition (OSR) aims to identify test samples whose classes are not seen during the training process. Recently, Unified Open-set Recognition (UOSR) has been proposed to reject not only unknown samples but also known but wrongly classified samples, which tends to be more practical in real-world applications. In this paper, we deeply analyze the UOSR task under different training and evaluation settings to shed light on this promising research direction. For this purpose, we first evaluate the UOSR performance of several OSR methods and show a significant finding that the uncertainty distribution of almost all these methods is actually closer to the expectation of UOSR than OSR. We show that the reason lies in the known but wrongly classified samples, as their uncertainty distribution is extremely close to unknown samples rather than known and correctly classified samples. Second, we analyze how the two training settings of OSR (*i.e.*, pre-training and outlier exposure) influence the UOSR. We find although they are both beneficial for distinguishing known and correctly classified samples from unknown samples, pre-training is also helpful for identifying known but wrongly classified samples while outlier exposure is not. In addition to different *training settings*, we also formulate a new *evaluation setting* for UOSR which is called few-shot UOSR, where only one or five samples per unknown class are available during evaluation to help identify unknown samples. We propose FS-KNNS for the few-shot UOSR to achieve state-of-the-art performance under all settings.

1 INTRODUCTION

Neural networks have achieved tremendous success in the closed-set classification (Deng et al., 2009), where the test samples share the same In-Distribution (InD) class set with training samples. Open-Set Recognition (OSR) (Scheirer et al., 2013) is proposed to tackle the challenge that some samples whose classes are not seen during training, which are Out-of-Distribution (OoD) data, may occur in the real world applications and should be rejected. However, some researchers have argued that the model should not only reject OoD samples but also InD samples that are Wrongly classified (InW), as the model gives the wrong answers for both of them. So Unified Open-set Recognition (UOSR) is proposed to only accept InD samples that are correctly classified (InC) and reject OoD and InW samples (Kim et al., 2021) simultaneously. The difference between the UOSR and OSR lies in the InW samples, where OSR is supposed to accept them while UOSR has the opposite purpose. Actually, UOSR is more useful in most real-world applications, but it receives little attention from the research community as it has been proposed very recently and lacks comprehensive systematic research. Therefore, we deeply analyze the UOSR problem in this work to fill this gap.

We first apply existing OSR methods for UOSR in Sec. 3, and then analyze UOSR under different *training settings* and *evaluation settings* in Sec. 4 and Sec. 5 respectively. In Sec. 3, several existing OSR methods are applied for UOSR, and we demonstrate a significant finding that *the uncertainty distribution of almost all existing methods designed for OSR is actually closer to the expectation of UOSR than OSR*. In other words, simply changing the uncertainty ground truth of InW samples from 0 (OSR) to 1 (UOSR) can bring a very large performance gain, as shown in Fig. 1 (a). We show that this phenomenon holds for different network architectures, datasets, and domains (image and video recognition). We find the devil is in the InW samples that have similar uncertainty distribution with OoD samples rather than InC samples.

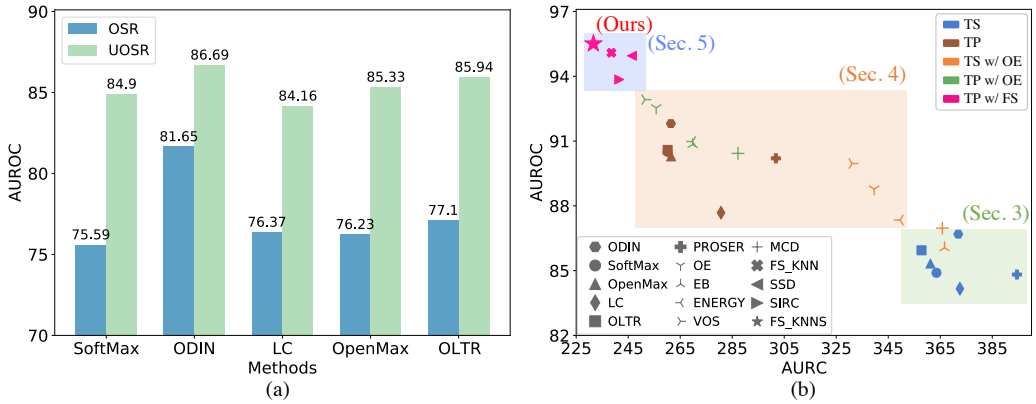

Figure 1: (a) shows that the UOSR performance is significantly better than OSR performance for the same method, which illustrates the uncertainty distribution of these OSR methods is actually closer to the expectation of UOSR than OSR. (b) shows the UOSR performance under different settings and the skeleton of this paper. Results are based on the ResNet50 backbone. CIFAR100 and TinyImageNet are InD and OoD datasets, respectively. (TS: Train from Scratch. TP: Train from Pre-training. OE: Outlier Exposure. FS: Few-shot.)

In Sec. 4, we introduce two *training settings* into UOSR, including pre-training (Hendrycks et al., 2019a) and outlier exposure (Hendrycks et al., 2019b; Yu & Aizawa, 2019; Thulasidasan et al., 2021), as they are both helpers of the OSR that introduce extra information beyond the training set. Pre-training is to use the weights that are trained on a large-scale dataset for better down-task performance, and outlier exposure is to introduce some background data without labels into training to help the model classify InD and OoD samples. We find both of them have better performance for InC/OoD discrimination, which explains why they are beneficial for OSR. However, pre-training is also helpful for InC/InW discrimination, while outlier exposure has a comparable or even worse performance to distinguish InC and InW samples. The performance of UOSR can be regarded as the comprehensive results of InC/OoD and InC/InW discrimination so that both techniques can boost the performance of UOSR. We build up a comprehensive UOSR benchmark that involves both pre-training and outlier exposure settings, as shown in Fig. 1 (b).

In addition to the two aforementioned *training settings*, we introduce a new *evaluation setting* into UOSR in Sec. 5. We formulate the few-shot UOSR, similar to SSD (Sehwag et al., 2022) that proposes few-shot OSR, where 1 or 5 samples per OoD class are introduced for reference to better identify OoD samples. We first develop a KNN-based baseline (Sun et al., 2022) FS-KNN for the few-shot UOSR. Although InC/OoD discrimination is improved due to the introduced OoD reference samples, the InC/InW discrimination is severely harmed compared to SoftMax baseline (Hendrycks & Gimpel, 2017). To alleviate this problem, we propose FS-KNNS that dynamically fuses the FS-KNN with SoftMax uncertainty scores to keep high InC/InW and InC/OoD performance simultaneously. Our FS-KNNS achieves state-of-the-art performance under all settings in the UOSR benchmark, as shown in Fig. 1 (b), even without outlier exposure during training. Note that InC/OoD performances are comparable between FS-KNNS and FS-KNN, but their distinct InC/InW performances makes FS-KNN better at OSR and FS-KNNS better at UOSR, which illustrates the difference between few-shot OSR and UOSR and the importance of InW samples during evaluation.

2 Towards Unified Open-set Recognition

In this section, we first formalize the UOSR problem and then discuss the relation between UOSR and other uncertainty-related tasks, as the core of UOSR is also how to estimate better uncertainty scores to reject InW and OoD samples.

Unified Open-set Recognition. Suppose the training dataset is $\mathcal{D}_{\text{train}} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N} \subset \mathcal{X} \times \mathcal{C}$, where $\mathcal{X}$ refers to the input space, *e.g.*, images or videos, and $\mathcal{C}$ refers to the InD sets. In closed-set recognition, all test samples come from the InD sets, *i.e.*, $\mathcal{D}_{\text{test}}^{\text{closed}} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{M} \subset \mathcal{X} \times \mathcal{C}$. In OSR and UOSR, the test samples may come from OoD sets $\mathcal{U}$ which are not overlap with InD sets $\mathcal{C}$, so

we have $\mathcal{D}_{\text{test}}^{\text{open}} = \mathcal{D}_{\text{test}}^{\text{closed}} \cup \mathcal{D}_{\text{test}}^{\text{unknown}}$, where $\mathcal{D}_{\text{test}}^{\text{unknown}} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{M'} \subset \mathcal{X} \times \mathcal{U}$. The InD test samples $\mathcal{D}_{\text{test}}^{\text{closed}}$ can be divided into two splits based on whether the sample is correctly classified or wrongly classified, *i.e.*, $\mathcal{D}_{\text{test}}^{\text{closed}} = \mathcal{D}_{\text{test}}^{\text{closed-c}} \cup \mathcal{D}_{\text{test}}^{\text{closed-w}}$, where $\mathcal{D}_{\text{test}}^{\text{closed-c}} = \{(\mathbf{x}_i, y_i) | \hat{y}_i = y_i\}_{i=1}^{M}$, $\mathcal{D}_{\text{test}}^{\text{closed-w}} = \{(\mathbf{x}_i, y_i) | \hat{y}_i \neq y_i\}_{i=1}^{M}$, and $\hat{y}_i$ refers to the model classification results of sample $\mathbf{x}_i$. The goal of UOSR is to reject InW and OoD samples and accept InC samples, so the ground truth uncertainty u of $\mathcal{D}_{\text{test}}^{\text{closed-w}}$ and $\mathcal{D}_{\text{test}}^{\text{unknown}}$ is 1 while for $\mathcal{D}_{\text{test}}^{\text{closed-c}}$ it is 0, as shown in Table 1. The key of UOSR is how to estimate the uncertainty $\hat{u}$ to be close to the ground truth uncertainty u.

The UOSR is proposed by (Kim et al., 2021) very recently, so it has not attracted many researchers to this problem yet. SIRC (Xia & Bouganis, 2022) augments SoftMax (Hendrycks & Gimpel, 2017) baseline for the better UOSR performance. Build upon these existing methods, we make a deep analysis of UOSR in this work. In Table 1, we compare different settings of the related uncertainty estimation tasks, and the detailed discussions are as follows.

Table 1: Comparison of uncertainty-related task settings. Cls: Classification. 0 and 1 refer to the corresponding ground truth uncertainty u, and u is not fixed in MC.

	InC	InW	OoD	Ordinal Rank	InD Cls
SP	0	1	✗	✓	✓
AD	0	0	1	✓	✗
OSR	0	0	1	✓	✓
UOSR	0	1	1	✓	✓
MC	-	-	✗	✗	✓

Selective Prediction (SP). Apart from the classical classification, SP also tries to estimate which sample is wrongly classified (Corbière et al., 2019; Moon et al., 2020; Granese et al., 2021), so the ground truth uncertainty of InW is 1. SP is constrained under the closed-set scenario and does not consider OoD samples during evaluation, as shown in Table 1, which is the key difference with UOSR. We involve some SP methods in the UOSR benchmark in Sec. 4.

Anomaly Detection (AD). AD is to detect anomaly patches within an image or anomaly events within a video (Deng & Li, 2022; Zaheer et al., 2022). Therefore, AD is defined within a test sample, and there is no InD samples classification task. Unlike other tasks that assign uncertainty on the image or video level, AD assigns uncertainty on the pixel or frame level.

Open-set Recognition (OSR) and Out-of-distribution Detection (OoDD). The task settings of OSR and OoDD are same, but their datasets might be different. Both of them aim to accept all InD samples no matter they are InC or InW samples, and reject OoD samples. OSR divides one dataset into two splits and uses one of them as InD data to train the model, while another split is regarded as OoD samples (Scheirer et al., 2013). In contrast, OoDD uses a whole dataset as InD data and regards another dataset as OoD data (Hendrycks & Gimpel, 2017). However, in the recent works about OSR in the video domain (Bao et al., 2021), it utilizes the OoDD setting rather than the OSR setting. In this work, we use OSR to represent the task setting and use one dataset as InD data and another dataset as OoD data. As mentioned before, the distinction between UOSR and OSR is InW samples, where OSR aims to accept them, and UOSR aims to reject them, so the ground truth uncertainty of InW samples is 0 in OSR and 1 in UOSR, as shown in Table 1. *Better InC/OoD performance is beneficial for both UOSR and OSR, but higher InC/InW discrimination is preferred by UOSR but not wanted by OSR.*

Model Calibration (MC). All tasks mentioned above solve the uncertainty ordinal ranking problem, *i.e.*, the ground truth uncertainty of each type of data is fixed, and the performance will be better if the estimated uncertainty is closer to the ground truth uncertainty (Geifman et al., 2019). In contrast, MC uses uncertainty to measure the probability of correctness (Guo et al., 2017). For a perfect calibrated model, if we consider all samples with a confidence of 0.8, then 80% of samples are correctly classified. In this case, 20% InW samples with a confidence of 0.8 are perfect, but in other uncertainty-related tasks, InW samples are supposed to have a confidence of either 1 or 0. MC also does not consider OoD samples. See Appendix A for more relation between UOSR and MC.

3 OSR Approaches for UOSR

In this section, we evaluate the existing OSR approaches for the UOSR problem, and show that the InW samples play a crucial role when evaluating the uncertainty quality. Specifically, simply changing the ground truth uncertainty u of InW samples from 0 to 1 can bring a large performance boost, as shown in Fig. 2. Then we provide the comprehensive experiment results and discussion of this phenomenon.

Table 2: Uncertainty distribution analysis in image domain with ResNet50. OoD dataset: TinyImageNet. AUORC (%) is reported.

Methods	InC/OoD	InC/InW	InW/OoD	OSR	UOSR
SoftMax	84.69	85.68	50.64	75.59	84.90
ODIN	88.35	80.76	64.36	81.65	86.69
LC	84.60	82.58	55.14	76.37	84.16
OpenMax	85.16	85.96	51.27	76.23	85.33
OLTR	85.99	85.74	52.22	77.10	85.94
PROSER	87.04	77.84	62.57	79.23	84.82

Table 3: Uncertainty distribution analysis in video domain with TSM backbone. OoD dataset is HMDB51. AUORC (%) is reported.

Methods	InC/OoD	InC/InW	InW/OoD	OSR	UOSR
OpenMax	88.56	82.53	64.47	82.27	85.95
Dropout	85.36	85.86	48.92	75.75	85.58
BNN SVI	89.93	88.85	55.44	80.10	89.44
SoftMax	88.73	88.01	54.19	79.72	88.42
RPL	89.22	88.06	55.74	79.67	88.70
DEAR	89.33	85.47	56.78	80.00	87.56

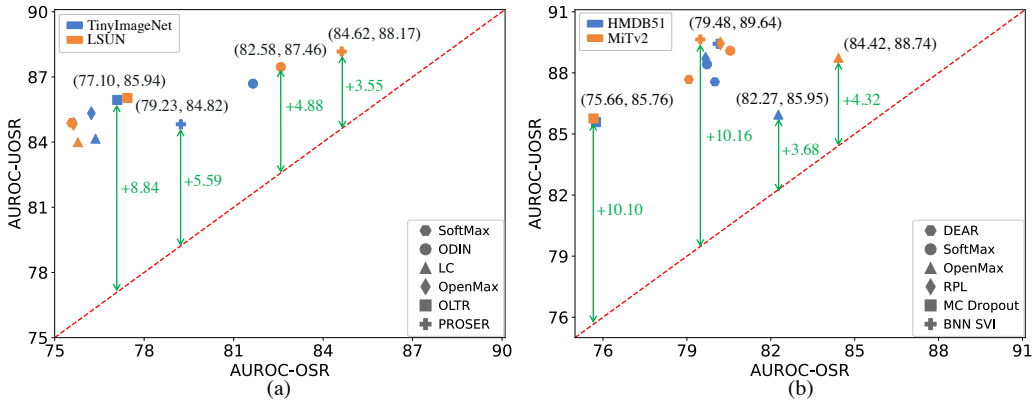

Figure 2: (a) and (b) show the relation between UOSR and OSR performance in the image and video domain under ResNet50 and TSM backbones. Different color indicates different OoD datasets. The red-dotted diagonal is where UOSR has the same AUROC as OSR. Green arrows show the performance gap between UOSR and OSR for the same method.

Applied Methods. We reproduce several classical OSR methods and evaluate their UOSR and OSR performance in our experiments, including SoftMax (Hendrycks & Gimpel, 2017), ODIN (Liang et al., 2018), LC (DeVries & Taylor, 2018), OpenMax (Bendale & Boult, 2016), OLTR (Liu et al., 2019) and PROSER (Zhou et al., 2021) in the image domain, as well as DEAR (Bao et al., 2021), RPL (Chen et al., 2020), MC Dropout (Gal & Ghahramani, 2016) and BNN SVI (Krishnan et al., 2018) in the video domain.

Datasets. In the image domain, the InD dataset is CIFAR100 (Krizhevsky et al., 2009) and the OoD datasets are TinyImageNet (Le & Yang, 2015) and LSUN (Yu et al., 2015). In the video domain, we follow (Bao et al., 2021) and adopt UCF101 (Soomro et al., 2012) as InD dataset, while the OoD datasets are HMDB51 (Kuehne et al., 2011) and MiTv2 (Monfort et al., 2021).

Experiments settings. We train the network using the InD dataset and evaluate the UOSR and OSR performance based on the ground truth in Table 1. The evaluation metric is AUORC (Hendrycks & Gimpel, 2017) which is a threshold-free value. The AUROC reflects the distinction quality of two uncertainty distributions. We adopt VGG13 (Simonyan & Zisserman, 2015) and ResNet50 (He et al., 2016) as the network backbone in the image domain and TSM (Lin et al., 2019) and I3D (Carreira & Zisserman, 2017) in the video domain. More experiment settings are in Appendix C.

Results. We provide the UOSR and OSR performance of different methods in Fig. 2. It shows that all data points are above the red dotted diagonal, illustrating that the UOSR performance is significantly better than the OSR performance of the same method, and this relationship holds across different datasets, domains (image and video), and network architectures (See Appendix B for results of more network backbones). The performance gap between UOSR and OSR of the same method can be very large, such as 8.84% for OLTR when TinyImageNet is the OoD dataset, and 10.16% for BNN SVI when the OoD dataset is MiTv2. Therefore, existing OSR methods have uncertainty distributions that are actually closer to the expectation of UOSR than OSR.

Analysis. To better understand our findings, we provide a detailed analysis of the uncertainty distribution relationships between InC, InW, and OoD samples in Tables 2 and 3. Note that higher

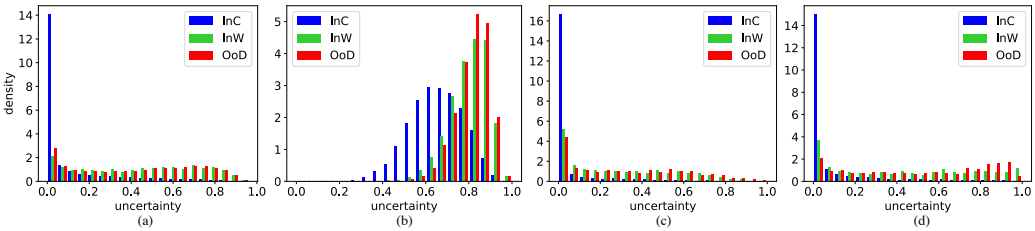

Figure 3: (a) and (b) are the SoftMax and ODIN methods in the image domain, while (c) and (d) are the SoftMax and DEAR methods in the video domain. OoD datasets are TinyImageNet for the image domain and HMDB51 for the video domain.

AUROC means a better distinction between two uncertainty distributions, and AUROC=50% means two distributions overlap with each other. From Tables 2 and 3 we can clearly see that AUROC of InC/OoD and InC/InW are significantly higher than InW/OoD, and AUROC of InW/OoD is very close to 50%. Therefore, the uncertainty distribution of InC samples is distinguishable from OoD and InW samples, and there is a lot of overlap between the uncertainty distributions of InW and OoD samples. We provide several uncertainty distribution visualizations in Fig. 3, in which we can see that InW and OoD samples share a similar uncertainty distribution, while InC samples have smaller uncertainty in general. Therefore, the devil is in the InW samples for OSR, as InW samples have a similar uncertainty distribution to OoD samples rather than InC samples, which is contradictory to the ground truth of OSR. In contrast, the uncertainty distribution of InW samples is consistent with the expectation of UOSR, so existing OSR methods are actually more suitable for UOSR than OSR.

Why? We provide some empirical reasons for the finding that InW samples share a similar uncertainty distribution with OoD samples rather than InC samples. The model assigns low uncertainty for InC samples as they are well fitted by the model, and the cross-entropy loss encourages the SoftMax probability of the corresponding class to be close to 1. However, InW and OoD samples are 'hard' samples which are not well fitted by the model, so the cross-entropy loss cannot constrain the predictions to be confidence and their predictions tend to be more uncertain (Vaze et al., 2022; Ranjan et al., 2017). The fact that AUROC of InW/OoD is close to 50% in Table 2 and 3 and the qualitative results in Fig. 3 can prove our claims.

A better closed-set model is better for UOSR. Recently, (Vaze et al., 2022) found that better closed-set performance means better OSR performance. We provide a deeper explanation of this finding and show that this conclusion also holds for UOSR. Table 4 shows that we improve the closed-set accuracy through data augmentation and longer training (Vaze et al., 2022). The open-set method is the SoftMax baseline with ResNet50 backbone, and the OoD dataset is TinyImageNet. For OSR, we can see that AUROC of InC/OoD is significantly better than OSR, which indicates the uncertainty distribution of InW samples are contradictory with the expectation of OSR. So less InW samples and better InC/OoD performance are two reasons for better OSR performance when closed-set accuracy is higher. For UOSR, both the InC/InW and InC/OoD performance are improving with the growth of closed-set accuracy, which brings better UOSR performance.

Table 4: Relation between closed-set accuracy *Acc.* (%) and open-set performance. *Aug:* Augmentation; *Ep:* Epoch. AUROC (%) is reported.

Aug.	Ep.	Acc.	InC/OoD	InC/InW	OSR	UOSR
✗	100	59.41	81.65	82.47	68.64	81.89
✓	100	68.83	84.18	85.13	73.71	84.40
✓	300	73.28	84.69	85.68	75.59	84.90

4 Pre-training and Outlier Exposure

After directly applying existing OSR methods for UOSR in Sec. 3, we explore two additional training settings in this section, including pre-training (Hendrycks et al., 2019a) and outlier exposure (Hendrycks et al., 2019b), which are effective methods to improve the OSR performance because of introduced extra information beyond the training set. Pre-training is to use large-scale pre-trained weights for initialization for better down-task performance. Outlier exposure is to introduce some unlabeled outlier data (OD) during training and regard these outlier data as the proxy as OoD data to improve the open-set performance. We find both of them also have a positive effect on UOSR, but for different reasons.

Table 5: UOSR benchmark in the image domain under the ResNet50 model. InD dataset is CIFAR100 while the OoD dataset is TinyImageNet. $\dagger, \ddagger, \diamondsuit$ refer to OSR-based, SP-based, UOSR-based methods. OD: Outlier Data. AUORC (%), AURC ($\times 10^3$) and Acc. (%) are reported.

		w/o pre-training				w/ pre-training			
		UOSR			OSR	UOSR			OSR
Methods	**OD**	**Acc.↑**	**AURC↓**	**AUROC↑**	**AUROC↑**	**Acc.↑**	**AURC↓**	**AUROC↑**	**AUROC↑**
SoftMax†	✗	73.28	363.55	84.90	75.59	86.44	260.14	90.50	85.93
ODIN†	✗	72.08	371.92	86.69	81.65	86.48	261.42	91.81	91.47
LC†	✗	72.08	372.55	84.16	76.37	86.48	280.69	87.68	85.16
OpenMax†	✗	73.64	361.22	85.33	76.23	86.43	261.53	90.29	85.86
OLTR†	✗	73.69	357.79	85.94	77.10	86.22	260.17	90.59	86.05
PROSER†	✗	68.08	394.48	84.82	79.23	81.32	301.78	90.20	89.29
BCE‡	✗	73.29	369.91	84.29	74.59	86.39	257.74	90.79	86.50
TCP‡	✗	71.80	369.17	85.17	75.46	86.83	261.97	89.95	85.67
DOCTOR‡	✗	72.08	378.45	84.48	75.06	86.48	262.05	90.24	85.75
SIRC(MSP, $\|z\|_1$)$^\diamondsuit$	✗	73.13	358.93	85.77	76.67	86.44	260.08	90.53	85.98
SIRC(MSP, Res.)$^\diamondsuit$	✗	73.13	348.40	87.26	80.29	86.44	247.75	92.99	90.39
SIRC($-\mathcal{H}$, $\|z\|_1$)$^\diamondsuit$	✗	73.13	355.62	86.57	78.16	86.44	257.28	91.23	87.27
SIRC($-\mathcal{H}$, Res.)$^\diamondsuit$	✗	73.13	346.21	87.71	81.43	86.44	244.55	93.68	91.62
OE	✓	73.54	339.59	88.78	84.35	85.43	255.77	92.54	90.65
EB	✓	71.82	366.65	86.06	81.95	85.06	269.79	90.84	90.38
Energy	✓	73.10	349.56	87.34	79.34	84.66	269.74	90.98	87.09
VOS	✓	73.18	331.12	89.96	85.78	85.93	251.44	92.92	91.34
MCD	✓	70.88	365.82	86.97	79.88	81.96	287.21	90.43	86.59

Pre-training settings. In the image domain, we use BiT pre-training weights (Kolesnikov et al., 2020) for ResNet50 and ImageNet (Deng et al., 2009) pre-training weights for VGG13. In the video domain, Kinetics400 (Carreira & Zisserman, 2017) pre-trained weights are used for initialization.

Outlier exposure settings. In the image domain, we use 300K Random Images dataset (Hendrycks et al., 2019b) as outlier data. In the video domain, outlier data comes from Kinetics400, and we filter out the overlapping classes with InD and OoD datasets. We implement several outlier exposure based methods, including OE (Hendrycks et al., 2019b), EB (Thulasidasan et al., 2021), Energy (Liu et al., 2020), VOS (Du et al., 2022) and MCD (Yu & Aizawa, 2019).

UOSR benchmark settings. We include OSR-based, SP-based, and UOSR-based methods to build a comprehensive UOSR benchmark. We implement several Selective Prediction (SP) based methods, including BCE, TCP (Corbière et al., 2019), CRL (Moon et al., 2020), and DOCTOR (Granese et al., 2021). Although these SP methods are originally designed to differentiate InC and InW samples, we adopt them in the UOSR setting to build a more comprehensive benchmark. The only UOSR-based method is SIRC (Xia & Bouganis, 2022), which combines two uncertainty scores for better UOSR performance. The evaluation metrics are AURC and AUROC (Kim et al., 2021). More details can be found in Appendix C.

Results. The UOSR benchmarks for image and video domains are Table 5 and 6 respectively. See Appendix D for the results of different model architectures. First, we can see the AUROC of UOSR is higher than OSR for almost all the methods under both settings, *i.e.*, pre-training and outlier exposure, which further strengthens the conclusion that the uncertainty distribution of OSR methods is closer to the ground truth of UOSR. For instance, AUROC is 87.34 for UOSR and 79.34 for OSR under the Energy method w/o pre-training in Table 5. Second, pre-training and outlier exposure can effectively boost the UOSR and OSR performance , *e.g.*, the pre-training boost the AUROC of UOSR/OSR from 84.90/75.59 to 90.50/85.93 for the SoftMax, and outlier exposure method VOS has 89.96/85.78 AUROC of UOSR/OSR compared to 84.90/75.59 of SoftMax in Table 5.

Analysis. To better understand why pre-training and outlier exposure are helpful for UOSR and OSR, we provide the InC/InW and InC/OoD discrimination performance of each method in Fig. 4. We can see that almost all outlier exposure methods are in the Q2, which means that outlier exposure methods have lower InC/InW AUROC and higher InC/OoD AUROC than the SoftMax baseline. For OSR, both lower InC/InW and higher InC/OoD AUROC are beneficial, but only higher InC/OoD AUROC is wanted for UOSR, while lower InC/InW AUROC is not. This can explain why some of

Table 6: UOSR benchmark in the video domain under the TSM model. InD dataset is UCF101 while the OoD dataset is HMDB51. †, ‡, ◊ refer to OSR-based, SP-based, UOSR-based methods. OD: Outlier Data. AUORC (%), AURC ($\times 10^3$) and Acc. (%) are reported.

		w/o pre-training				w/ pre-training			
		UOSR			OSR	UOSR			OSR
Methods	**OD**	**Acc.↑**	**AURC↓**	**AUROC↑**	**AUROC↑**	**Acc.↑**	**AURC↓**	**AUROC↑**	**AUROC↑**
OpenMax†	✗	73.92	185.81	85.95	82.27	95.32	75.75	91.22	90.89
MC Dropout†	✗	73.63	184.66	85.58	75.75	95.06	79.53	90.30	88.23
BNN SVI†	✗	71.51	181.45	89.44	80.10	94.71	69.89	93.58	91.81
SoftMax†	✗	73.92	173.14	88.42	79.72	95.03	68.08	93.94	91.75
RPL†	✗	71.46	186.18	88.70	79.67	95.59	72.88	92.44	90.53
DEAR†	✗	71.33	215.80	87.56	80.00	94.41	102.01	91.50	91.49
BCE‡	✗	69.90	223.57	83.27	78.96	93.66	110.42	83.83	81.64
CRL‡	✗	71.80	183.76	88.75	78.57	95.22	67.61	93.36	91.38
DOCTOR‡	✗	72.01	182.04	88.73	79.76	95.06	65.61	93.89	91.80
SIRC(MSP, $\|z\|_1$)◊	✗	73.59	173.32	88.71	80.33	95.00	65.97	93.74	91.42
SIRC(MSP, Res.)◊	✗	73.59	174.76	88.17	78.74	95.00	65.43	93.83	91.73
SIRC($-\mathcal{H}$, $\|z\|_1$)◊	✗	73.59	172.30	89.27	81.61	95.00	66.11	94.06	91.95
SIRC($-\mathcal{H}$, Res.)◊	✗	73.59	175.79	88.50	79.62	95.00	69.06	94.07	92.18
OE	✓	74.42	162.36	90.19	86.29	94.71	67.93	94.33	93.40
EB	✓	70.90	212.01	85.32	86.47	94.66	67.83	94.40	93.06
Energy	✓	74.00	187.82	86.10	84.51	95.27	65.68	94.44	93.62
VOS	✓	74.68	172.71	87.98	87.09	94.79	64.99	94.97	93.72
MCD	✓	72.49	168.83	91.26	85.57	93.47	71.19	95.34	93.68

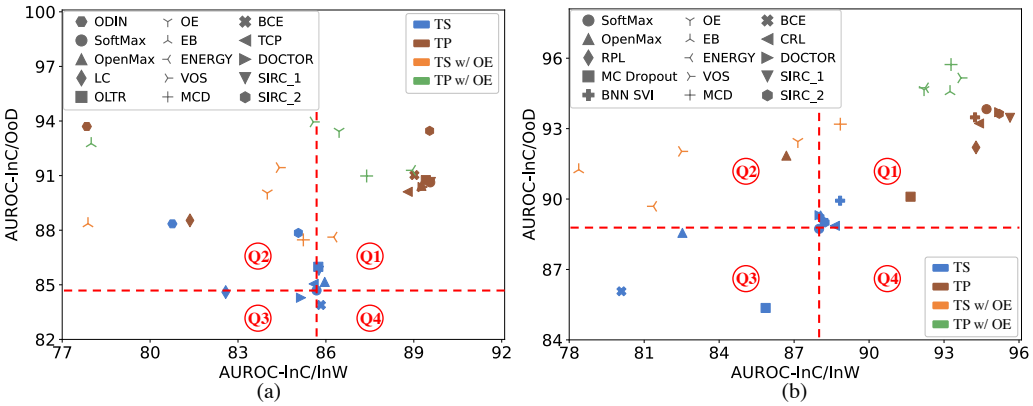

Figure 4: (a) and (b) plot the InC/InW and InC/OoD discrimination in the image and video domain. We set the SoftMax method training from scratch as the original point and divide the coordinate system into 4 quadrants (Q1 to Q4).

the UOSR performances with outlier exposure are comparable or even worse than the baseline, such as EB, Energy, and VOS in the video domain Table 5, but all of the OSR performances are increased. In contrast, pre-training is helpful for both InC/InW and InC/OoD AUROC, as methods with pre-training are in the Q1. Therefore, outlier exposure and pre-training bring the UOSR performance gain for different reasons. InC/OoD performance is improved by both of the techniques, but pre-training is also helpful for InC/InW, but outlier exposure is not. This can be explained by the fact that the model may see additional InD samples during pre-training so that the closed-set accuracy and InC/InW discrimination are improved. In contrast, outlier exposure only provides OoD data but no more InD data, so only InC/OoD performance is improved.

5 Few-shot Unified Open-set Recognition

In addition to the analysis of two useful training settings in Sec. 4, we introduce a new evaluation setting into UOSR in this section. Inspired by the recent work SSD (Sehwag et al., 2022) which proposes the few-shot OSR, we introduce the few-shot UOSR, where 1 or 5 samples per OoD class can be introduced for reference to better distinguish OoD samples. The introduced reference OoD

Table 7: Results of few-shot UOSR in the image domain. Model is ResNet50 with pre-training. InD and OoD datasets are CIFAR100 and TinyImageNet. AUORC (%) and AURC ($\times 10^3$) are reported.

	5-shot					1-shot				
	AURC↓	AUROC↑				AURC↓	AUROC↑			
Methods	**UOSR**	**UOSR**	**OSR**	**InC/OoD**	**InC/InW**	**UOSR**	**UOSR**	**OSR**	**InC/OoD**	**InC/InW**
SoftMax	260.14	90.51	85.93	90.64	89.58	260.14	90.51	85.93	90.64	89.58
KNN	245.48	93.45	92.34	94.86	83.08	245.48	93.45	92.34	94.86	83.08
FS-KNN	238.54	95.09	95.91	97.20	79.58	239.71	94.67	94.91	96.52	81.04
SSD	246.42	94.95	**97.89**	**98.32**	70.14	245.99	94.16	**96.51**	**97.16**	72.09
FS-KNN+S	239.49	94.14	91.64	95.27	85.85	240.30	93.97	91.36	94.99	86.45
FS-KNN*S	255.82	91.26	87.10	91.47	**89.67**	255.69	91.30	87.16	91.51	**89.69**
SIRC	241.58	93.85	91.42	94.44	89.58	239.64	94.20	92.03	94.87	89.26
FS-KNNS	**231.61**	**95.51**	94.16	96.54	87.98	**234.84**	**94.91**	93.10	95.84	88.08

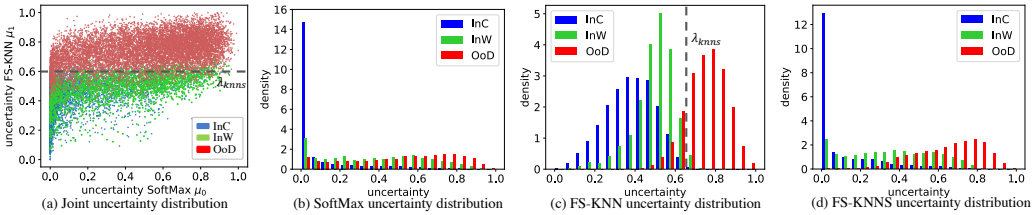

(a) Joint uncertainty distribution (b) SoftMax uncertainty distribution (c) FS-KNN uncertainty distribution (d) FS-KNNS uncertainty distribution

Figure 5: Uncertainty scores of each test sample (a) and uncertainty distribution of SoftMax (b), FS-KNN (c), and FS-KNNS (d).

datasets are marked as $\mathcal{D}_{test}^{ref} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N'} \subset \mathcal{X} \times \mathcal{U}$. We show that few-shot UOSR has distinct challenges from few-shot OSR because of InW samples, and we propose our FS-KNNS method to achieves the state-of-the-art UOSR performance even without outlier exposure during training.

Baselines and dataset settings. SSD is the only existing few-shot OSR method that utilizes the Mahalanobis distance with InD training samples $\mathcal{D}_{train}$ and OoD reference samples $\mathcal{D}_{test}^{ref}$. KNN (Sun et al., 2022) utilizes the feature distance with samples in $\mathcal{D}_{train}$ as uncertainty scores and shows it is better than Mahalanobis distance. So we slightly modify KNN to adapt to the few-shot UOSR and find the modified FS-KNN can already beat SSD. $\mathcal{D}_{test}^{ref}$ comes from the validation set of TinyImageNet and training set of UCF101 in the image and video domain. We repeat the evaluation process until all OoD reference samples are used and calculate the mean as the final few-shot result.

FS-KNN. Given a test sample $\mathbf{x}^*$ and a feature extractor f, the feature of $\mathbf{x}^*$ is $\mathbf{z}^* = f(\mathbf{x}^*)$. The cosine similarity set between $\mathbf{z}^*$ and the feature set of $\mathcal{D}_{train}$ is $\mathcal{S}_{train} = \{d_i\}_{i=1}^{N}, d_i = (\mathbf{z}^* \cdot \mathbf{z}_i)/(\|\mathbf{z}^*\| \cdot \|\mathbf{z}_i\|), \mathbf{z}_i = f(\mathbf{x}_i), \mathbf{x}_i \in \mathcal{D}_{train}$. $\mathcal{S}_{test}^{ref}$ is similar with $\mathcal{S}_{train}$ except $\mathbf{x}_i$ comes from $\mathcal{D}_{test}^{ref}$. We believe the uncertainty should be higher if the test sample is more similar to reference OoD samples and not similar with InD samples. Therefore, the uncertainty score is

$$\hat{u}_{fs-knn} = 1 - topK(S_{train}) + topK(S_{test}^{ref}), \tag{1}$$

where $topK$ means the Kth largest value. The performances of SoftMax, KNN, FS-KNN, and SSD are in Table 7 and 8. Pre-training is used for better performance. FS-KNN has better InC/OoD performance than KNN as reference OoD samples are introduced. Although the overall UOSR performance of FS-KNN is better than the SoftMax baseline, *the InC/InW performance is significantly sacrificed, which is also an important aspect of UOSR.* For example, InC/InW performance drops from 89.58 to 79.58 in the 5-shot results of Table 7. So we naturally ask a question: *Can we improve the InC/OoD performance based on the introduced OoD reference samples while keeping similar InC/InW performance with SoftMax baseline?* This is the key difference between few-shot OSR and few-shot UOSR, as low InC/InW performance is wanted in OSR but not preferred in UOSR.

FS-KNNS. Inspired by SIRC (Xia & Bouganis, 2022), we aim to find a way to fuse SoftMax and FS-KNN scores so that the mixed score can keep the high InC/OoD performance of FS-KNN and meanwhile has the comparable InC/InW performance with SoftMax. The uncertainty distributions of SoftMax and FS-KNNS are depicted in Fig. 5. We find that OoD samples have larger uncertainty than InW and InC samples in the FS-KNN method, but the uncertainty of InC samples overlaps a lot with InW samples, which explains the reason that InC/OoD performance is high but InC/InW

Table 8: Results of few-shot UOSR in the video domain. Model is TSM with pre-training. InD and OoD datasets are UCF101 and HMDB51. AUORC (%) and AURC ($\times 10^3$) are reported.

	5-shot					1-shot				
	AURC↓	AUROC↑				AURC↓	AUROC↑			
Methods	**UOSR**	**UOSR**	**OSR**	**InC/OoD**	**InC/InW**	**UOSR**	**UOSR**	**OSR**	**InC/OoD**	**InC/InW**
SoftMax	66.55	93.66	91.44	93.46	**94.99**	66.55	93.66	91.44	93.46	**94.99**
KNN	73.73	93.38	92.99	94.11	88.44	73.73	93.38	92.99	94.11	88.44
FS-KNN	68.44	95.04	**96.14**	**96.71**	83.74	73.19	93.60	**94.32**	95.09	83.53
SSD	70.36	93.91	95.41	95.97	79.97	75.61	92.48	93.34	94.07	81.75
FS-KNN+S	65.64	95.44	94.69	96.13	90.78	68.98	94.54	93.66	95.14	90.48
FS-KNN*S	65.29	94.09	92.13	93.96	**94.99**	65.42	94.04	92.08	93.91	94.97
SIRC	62.41	95.11	93.79	95.13	94.94	63.11	94.88	93.49	94.88	94.84
FS-KNNS	**60.00**	**96.19**	95.37	96.40	94.82	**62.65**	**95.35**	94.26	**95.48**	94.49

performance is low. In contrast, InW and OoD samples share similar uncertainty in the SoftMax method, which brings higher InC/InW performance. Therefore, we want to keep the uncertainty of InW samples in the SoftMax method, as well as the uncertainty of OoD samples in the FS-KNN method. In this way, the mixed scores obtain the high InC/OoD performance from FS-KNN while keeping the comparable InC/InW performance of SoftMax. We call this method FS-KNN with SoftMax (FS-KNNS), and the uncertainty is

$$\hat{u}_{fs-knns} = u_0 + \frac{1}{1 + e^{-\alpha(u_1 - \lambda_{knns})}} u_1, \tag{2}$$

where u_0 and u_1 refer to the uncertainty score of SoftMax and FS-KNN, respectively. λ_{knns} is a threshold to determine when the weight of u_1 becomes large, and α is a coefficient to control the change rate of the weight. $\hat{u}_{fs-knns}$ will be largely influenced by u_1 when $u_1 > \lambda_{knns}$. A proper λ_{knns} should be located between the InW and OoD samples, as shown in Fig. 5 (a) and (c). In this way, the uncertainty of InW samples is mainly controlled by SoftMax, and OoD samples are strengthened by FS-KNN. Ablation study about hyperparameters K, α, and λ_{knns} is in Appendix E.

Discussion. The uncertainty distribution of FS-KNNS is shown in Fig. 5 (d), which shows the uncertainty of InC and InW samples are similar to SoftMax (b), but the uncertainty of OoD samples is larger. From Table 7 and 8, we can see our FS-KNNS has significantly better InC/InW performance than FS-KNN (87.98 and 79.58 under ResNet50), and meanwhile keeps the high InC/OoD performance, so the overall UOSR performance is better than both SoftMax and FS-KNN. We also try three score fusion methods, including FS-KNN+S, FS-KNN*S, and SIRC, but these methods are general score fusion methods, while our method is specifically designed for UOSR, so our FS-KNNS surpass them. Our FS-KNNS is totally based on the existing model trained by the classical cross-entropy loss, so there is no extra effort and no outlier data during training, and our FS-KNNS still has better performance than all outlier exposure methods and achieves state-of-the-art UOSR performance, as shown in Fig. 1 (b) and 6. Note that the best choice for OSR (FS-KNN or SSD) may not be the best choice for UOSR (FS-KNNS), as their expectation of InC/InW is contradictory, which shows the necessity of our proposed few-shot UOSR.

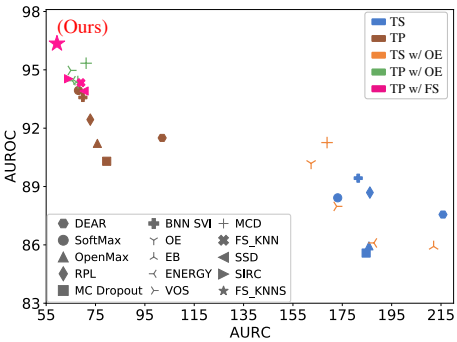

Figure 6: UOSR performance under all settings of TSM backbone in the video domain. OoD dataset is HMDB51.

6 Conclusion

Rejecting both InW and OoD samples during evaluation has very practical value in most real-world applications, but UOSR is a newly proposed problem and lacks comprehensive analysis and research, so we look deeper into the UOSR problem in this work. We first demonstrate the uncertainty distribution of almost all existing OSR methods is actually closer to the expectation of UOSR than OSR, and then analyze two training settings including pre-training and outlier exposure. Finally, we introduce a new evaluation setting into UOSR, which is few-shot UOSR, and we propose a FS-KNNS method that achieves state-of-the-art performance under all settings.

Reproducibility Statement

We provide the uncertainty distribution result files and corresponding codes required to reproduce Tables 2, 3, 5, 6, 7 and 8 in the supplementary material folder *UOSR/UOSR_eval*. The users do not need to run the training or evaluation process for a deep learning model. Simply follow the *README_UOSR_eval.md* in the *UOSR/UOSR_eval* folder and all performance results can be printed in the terminal in the table form.

We also provide the training codes for all methods under all settings (pre-training and outlier exposure) based on the ResNet50 backbone in the supplementary material folder *UOSR/UOSR_train*. The corresponding guideline is in *README_UOSR_train.md*. The datasets and experiments details can be found in Appendix C.

How to conduct few-shot evaluation to obtain the corresponding uncertainty distribution result files is also included in the supplementary material folder *UOSR/UOSR_few_shot*. The corresponding guideline is in *README_UOSR_few_shot.md*. The ablation study of hyper parameters in our FS-KNNS is in Appendix E.

REFERENCES

Wentao Bao, Qi Yu, and Yu Kong. Evidential deep learning for open set action recognition. In *ICCV*, 2021.

Abhijit Bendale and Terrance E Boult. Towards open set deep networks. In *CVPR*, 2016.

Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.

Guangyao Chen, Limeng Qiao, Yemin Shi, Peixi Peng, Jia Li, Tiejun Huang, Shiliang Pu, and Yonghong Tian. Learning open set network with discriminative reciprocal points. In *ECCV*, 2020.

Charles Corbière, Nicolas Thome, Avner Bar-Hen, Matthieu Cord, and Patrick Pérez. Addressing failure prediction by learning model confidence. In *NeurIPS*, 2019.

Hanqiu Deng and Xingyu Li. Anomaly detection via reverse distillation from one-class embedding. In *CVPR*, 2022.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.

Terrance DeVries and Graham W Taylor. Learning confidence for out-of-distribution detection in neural networks. *arXiv preprint arXiv:1802.04865*, 2018.

Xuefeng Du, Zhaoning Wang, Mu Cai, and Yixuan Li. Vos: Learning what you don't know by virtual outlier synthesis. In *ICLR*, 2022.

Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 2016.

Yonatan Geifman, Guy Uziel, and Ran El-Yaniv. Bias-reduced uncertainty estimation for deep neural classifiers. In *ICLR*, 2019.

Federica Granese, Marco Romanelli, Daniele Gorla, Catuscia Palamidessi, and Pablo Piantanida. Doctor: A simple method for detecting misclassification errors. In *NeurIPS*, 2021.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *ICML*, 2017.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

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

Dan Hendrycks, Kimin Lee, and Mantas Mazeika. Using pre-training can improve model robustness and uncertainty. In *ICML*, 2019a.

Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. In *ICLR*, 2019b.

Jihyo Kim, Jiin Koo, and Sangheum Hwang. A unified benchmark for the unknown detection capability of deep neural networks. *arXiv preprint arXiv:2112.00337*, 2021.

Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *ECCV*, 2020.

Ranganath Krishnan, Mahesh Subedar, and Omesh Tickoo. Bar: Bayesian activity recognition using variational inference. In *NeurIPS*, 2018.

Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *ICCV*, 2011.

Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 2015.

Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. In *ICLR*, 2018.

Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, 2019.

Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. In *NeurIPS*, 2020.

Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X. Yu. Large-scale long-tailed recognition in an open world. In *CVPR*, 2019.

Mathew Monfort, Bowen Pan, Kandan Ramakrishnan, Alex Andonian, Barry A McNamara, Alex Lascelles, Quanfu Fan, Dan Gutfreund, Rogerio Feris, and Aude Oliva. Multi-moments in time: Learning and interpreting models for multi-action video understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.

Jooyoung Moon, Jihyo Kim, Younghak Shin, and Sangheum Hwang. Confidence-aware learning for deep neural networks. In *ICML*, 2020.

Rajeev Ranjan, Carlos D Castillo, and Rama Chellappa. L2-constrained softmax loss for discriminative face verification. *arXiv preprint arXiv:1703.09507*, 2017.

Walter J. Scheirer, Anderson de Rezende Rocha, Archana Sapkota, and Terrance E. Boult. Toward open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2013.

Vikash Sehwag, Mung Chiang, and Prateek Mittal. Ssd: A unified framework for self-supervised outlier detection. In *ICLR*, 2022.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.

Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.

Yiyou Sun, Yifei Ming, Xiaojin Zhu, and Yixuan Li. Out-of-distribution detection with deep nearest neighbors. In *ICML*, 2022.

Sunil Thulasidasan, Sushil Thapa, Sayera Dhaubhadel, Gopinath Chennupati, Tanmoy Bhattacharya, and Jeff Bilmes. An effective baseline for robustness to distributional shift. In *ICMLA*, 2021.

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

Guoxuan Xia and Christos-Savvas Bouganis. Augmenting softmax information for selective classification with out-of-distribution data. *arXiv preprint arXiv:2207.07506*, 2022.

Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, 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.

Qing Yu and Kiyoharu Aizawa. Unsupervised out-of-distribution detection by maximum classifier discrepancy. In *ICCV*, 2019.

Dahua Lin Yue Zhao, Yuanjun Xiong. Mmaction. https://github.com/open-mmlab/mmaction, 2019.

M Zaigham Zaheer, Arif Mahmood, M Haris Khan, Mattia Segu, Fisher Yu, and Seung-Ik Lee. Generative cooperative learning for unsupervised video anomaly detection. In *CVPR*, 2022.

Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Learning placeholders for open-set recognition. In *CVPR*, 2021.

A RELATION BETWEEN MODEL CALIBRATION AND UOSR

Different from Selective Prediction (SP), Anomaly Detection (AD), OSR, and UOSR, Model Calibration (MC) is not an uncertainty ordinal ranking problem. In other words, all other settings expect the uncertainty of a test sample to be either 0 or 1, while MC is not. A perfect calibrated model should meet (Guo et al., 2017)

$$P(\hat{y} = y | f(x) = p) = p \quad \forall p \in [0, 1] \tag{3}$$

where p is the confidence or the negative version of uncertainty $p = 1 - u$. For example, given 100 test samples whose confidence scores are all 0.8, then the model is perfectly calibrated if 80% samples are correctly classified. In this case, 20% samples with confidence 0.8 is consistent with the requirement of a perfect calibrated model, but the confidence is supposed to be either 0 or 1 in other settings. The performance of MC is evaluated by ECE, and the readers may refer to (Guo et al., 2017) for the formal definition. Smaller ECE means better calibrated model. We provide 5 examples to illustrate the relation between the performance of UOSR and MC in Fig. 7. Note that MC does not consider OoD samples during evaluation, so we also do not involve OoD samples in Fig. 7. Therefore, UOSR in Fig. 7 is also equal to SP. From (f) we can see that the performance of UOSR and model calibration is not perfectly positively correlated, *i.e.*, the best case for model calibration (c) is not the best case for UOSR and vice versa (e).

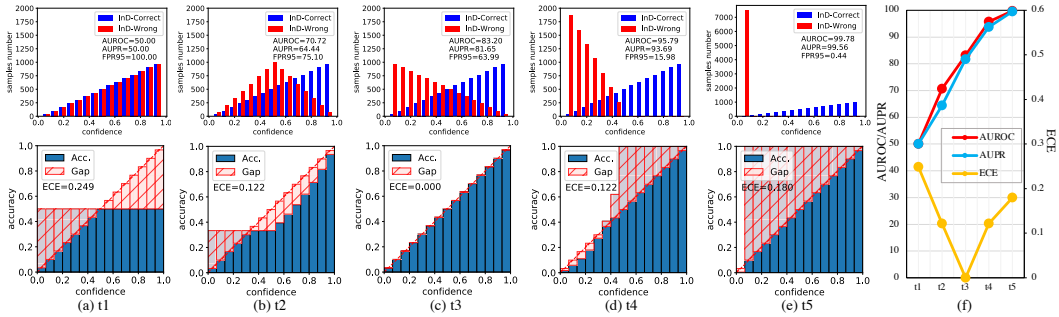

Figure 7: We provide 5 samples in (a)-(e), where we keep the confidence distribution of InC and change the confidence distribution of InW samples. The evaluation metrics of UOSR are AUROC and AUPR, and ECE is for the MC.

B EXISTING OSR METHODS ARE MORE SUITABLE FOR UOSR

In addition to Fig. 2, we provide the results of VGG13 backbone in the image domain and I3D backbone in the video domain. Fig. 8 further proves the finding that existing OSR methods are actually more suitable for UOSR holds for differnt model architectures and domains.

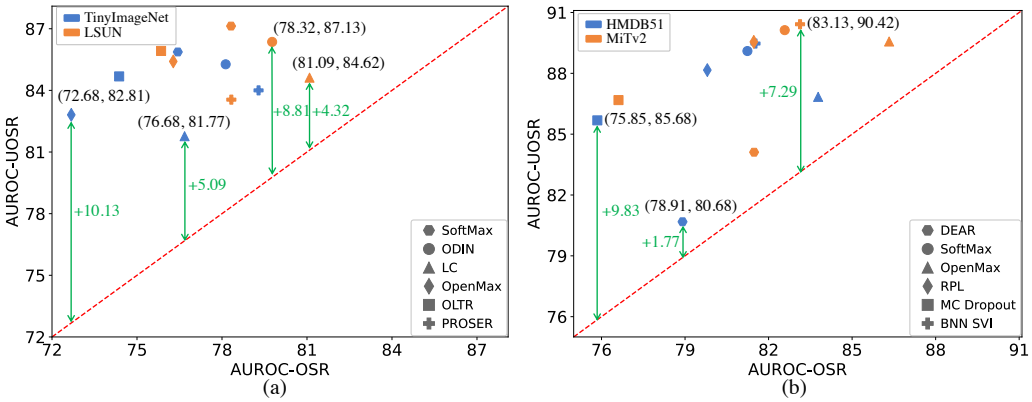

Figure 8: (a) and (b) are conducted using the VGG13 and I3D backbone in the image and video domain respectively. InD datasets are CIFAR100 and UCF101 for (a) and (b), and OoD datasets are shown with different colors.

C IMPLEMENTATION DETAILS

C.1 IMAGE DOMAIN

w/o pre-training. When we train the model from scratch, we find that setting the base learning rate as 0.1 and step-wisely decayed by 10 every 24000 steps with totally 120000 steps can achieve good enough closed-set performance. We use a linear warmup strategy to warmup the training in the first 500 steps. We use SGD with momentum, batch size 128 for all models.

w/ ImageNet-21k pre-training. When we fine tune the model with ImageNet pretrained model from (Kolesnikov et al., 2020), we set the base learning rate as 0.003 and step-wisely decayed every 3000 steps with totally 10000 steps. We use a linear warmup strategy to warmup the training in the first 500 steps. We use SGD with momentum, batch size 512 for all models. For outlier exposure methods, the batch size of outlier data is set to 128.

InD and OoD datasets. We follow the datasets setting in (Liang et al., 2018). The training InD dataset is CIFAR-100, which contains 100 classes with 50000 training images and 10000 test images. The OoD datasets for open-set evaluation are TinyImageNet and LSUN. The TinyImageNet dataset contains 10000 test images from 200 different classes, we use the TinyImageNet (resize) in (Liang et al., 2018). The Large-scale Scene UNderstanding (LSUN) dataset consists of 10000 images of 10 different scenes categories like classroom, conference room, dining room, etc. We use LSUN (resize) from (Liang et al., 2018). The size of images in both InD and OoD datasets is 32 $\times$ 32. We use the same strategy in image preprocessing stage as (Kolesnikov et al., 2020), if the resolution of the training image is lower than 96 $\times$ 96, we use 128 $\times$ 128 image cropping technique and random horizontal mirroring followed by 160 $\times$ 160 image resize. Test images and OoD images are directly resized to 128 $\times$ 128.

Outlier datasets. We use 300K Random Images dataset from (Hendrycks et al., 2019b) as outlier dataset for those outlier exposure methods. The 300K Random Images dataset is a debiased dataset with real images scraped from online. According to (Hendrycks et al., 2019b), all images that belong to CIFAR classes and images with divisive metadata have been removed.

C.2 VIDEO DOMAIN

w/o pre-training. When we train the model from scratch, we find that setting the base learning rate as 0.05 and step-wisely decayed every 160 epochs with totally 400 epochs can achieve good enough closed-set performance. The batch size is 256 for all methods.

w/ Kinetics400 pre-training. We follow (Bao et al., 2021) to set the base learning rate as 0.001 and step-wisely decayed every 20 epochs with totally 50 epochs. For those methods without outlier exposure, we fix the parameters of all Batch Normalization layers except the first one, and set the learning rate of the fully connected layer to be 10 times of the base learning rate. For those methods with outlier exposure, all parameters are updated with the same learning rate.

InD and OoD datasets. We follow the datasets setting in (Bao et al., 2021). The training InD dataset is UCF101, which contains 101 classes with 9537 training samples and 3783 test samples. The OoD datasets for open-set evaluation are HMDB51 and MiT-v2. We use the test sets of them which contain 1530 samples and 30500 samples respectively. For UCF101 and HMDB51, we follow the MMAction (Yue Zhao, 2019) to use the split 1 for training and evaluation, which is the same with (Bao et al., 2021). Note that in (Bao et al., 2021), they find some classes in HMDB51 overlap with those in UCF101 but they do not clean them. We remove the overlapping classes in UCF101 and HMDB51 so that OoD data does not contain any samples of InD classes. The classes we remove in HMDB51 and the corresponding same classes in UCF101 are in Table 9.

Table 9: Overlapping classes in HMDB51 and UCF101.

HMDB51	35, Shoot bow	29, Push up	15, Golf	26, Pull up
UCF101	2, Archery	71, PushUps	32, GolfSwing	69, PullUps
HMDB51	30, Ride bike	34, Shoot ball	43, Swing baseball	31, Ride horse
UCF101	10, Biking	7, Basketball	6, BaseballPitch	41, HorseRiding

Outlier datasets. We use Kinetics400 as our outlier datasets for those outlier exposure methods. To ensure that the classes of outlier data is not overlapping with InD data and OoD data, we remove corresponding classes in Kinetics400. The overlapping classes between Kinetics400 and UCF101/HMDB51 are too many to be listed here (129 overlapping classes). The available training sample ID list of Kinetics400 and all codes will be public. We pick up 271 classes from Kinetics400 and 25 samples in each class as outlier data.

D UOSR BENCHMARK AND FEW-SHOT RESULTS OF VGG13

We provide the UOSR benchmark and few-shot results of VGG13 in Table 10 and 11 respectively. Our FS-KNNS still achieves the best performance under all settings.

Table 10: Unified open-set recognition benchmark in the image domain. All methods are conducted under the VGG13 model. InD dataset is CIFAR100 while OoD dataset is TinyImageNet. $\dagger, \ddagger, \diamondsuit$ refer to OSR-based, SP-based, UOSR-based methods. OD: use Outlier Data in training.

		w/o pre-training				**w/ pre-training**			
		UOSR			**OSR**	**UOSR**			**OSR**
Methods	**OD**	**Acc.↑**	**AURC↓**	**AUROC↑**	**AUROC↑**	**Acc.↑**	**AURC↓**	**AUROC↑**	**AUROC↑**
SoftMax†	✗	75.07	341.06	85.87	76.44	74.69	311.70	91.01	85.13
ODIN†	✗	75.07	346.70	85.27	78.13	74.69	314.80	91.24	88.90
LC†	✗	75.07	365.96	81.77	76.68	74.69	339.96	88.27	88.95
OpenMax†	✗	74.52	368.58	82.81	72.68	75.05	312.70	90.47	84.25
OLTR†	✗	73.80	365.61	84.68	74.37	74.52	306.73	92.12	87.08
PROSER†	✗	70.95	376.63	84.00	79.29	71.11	367.16	86.45	85.97
BCE‡	✗	74.74	339.25	86.54	76.99	74.45	325.29	88.85	81.14
TCP‡	✗	75.09	340.41	85.94	76.18	74.83	313.80	90.57	84.35
DOCTOR‡	✗	75.07	340.69	85.95	76.72	74.69	310.29	91.37	85.97
SIRC(MSP, $\|z\|_1$)$^\diamondsuit$	✗	75.07	343.44	85.33	75.63	74.69	313.22	90.67	84.59
SIRC(MSP, Res.)$^\diamondsuit$	✗	75.07	336.87	86.77	78.23	74.69	309.67	91.46	86.14
SIRC($-\mathcal{H}$, $\|z\|_1$)$^\diamondsuit$	✗	75.07	343.53	85.25	76.37	74.69	309.19	91.56	86.82
SIRC($-\mathcal{H}$, Res.)$^\diamondsuit$	✗	75.07	335.71	86.93	79.13	74.69	305.34	92.39	88.34
OE	✓	71.71	312.44	93.71	89.76	73.38	306.72	93.34	92.77
EB	✓	74.19	340.66	87.67	85.73	72.76	340.81	89.40	89.79
Energy	✓	73.30	306.08	93.40	90.09	72.76	307.20	93.89	92.19
VOS	✓	71.73	312.75	93.68	90.76	73.08	306.37	93.65	92.97
MCD	✓	70.45	316.57	94.21	91.54	72.08	300.87	95.40	94.26

Table 11: Results of few-shot UOSR in the image domain. Model is VGG13 with pre-training. InD and OoD datasets are CIFAR100 and TinyImageNet respectively.

	5-shot					**1-shot**				
	AURC↓	**AUROC↑**				**AURC↓**	**AUROC↑**			
Methods	**UOSR**	**UOSR**	**OSR**	**InC/OoD**	**InC/InW**	**UOSR**	**UOSR**	**OSR**	**InC/OoD**	**InC/InW**
SoftMax	311.68	91.00	85.13	92.10	86.68	311.68	91.00	85.13	92.10	86.68
KNN	306.04	93.04	91.61	95.75	82.32	306.04	93.04	91.61	95.75	82.32
FS-KNN	295.80	95.44	98.87	**99.39**	79.83	296.75	95.27	98.48	99.16	79.87
SSD	321.31	92.57	**99.26**	99.36	65.72	315.46	93.11	**99.10**	**99.27**	68.77
FS-KNN+S	296.42	94.22	92.15	96.50	85.20	296.67	94.20	92.15	96.48	85.16
FS-KNN*S	303.01	92.85	89.42	94.38	**86.83**	302.84	92.88	89.45	94.41	**86.83**
SIRC	287.91	95.71	96.33	98.00	86.66	288.12	95.68	96.30	98.00	86.53
FS-KNNS	**285.85**	**95.95**	96.52	98.30	86.66	**287.02**	**95.75**	95.96	98.04	86.67

E ABLATION STUDY OF FEW-SHOT UOSR

Analyze of K. Similar with (Sun et al., 2022), we find the value of K influences the performance a lot. The ablation study of K is shown in Fig. 9. It shows that the best K for InC/OoD discrimination is between 3 and 5, and drops quickly after 7. In contrast, the InC/InW performance keeps increasing until 20 and then drops. The overall UOSR performance achieves the best when $K = 5$.

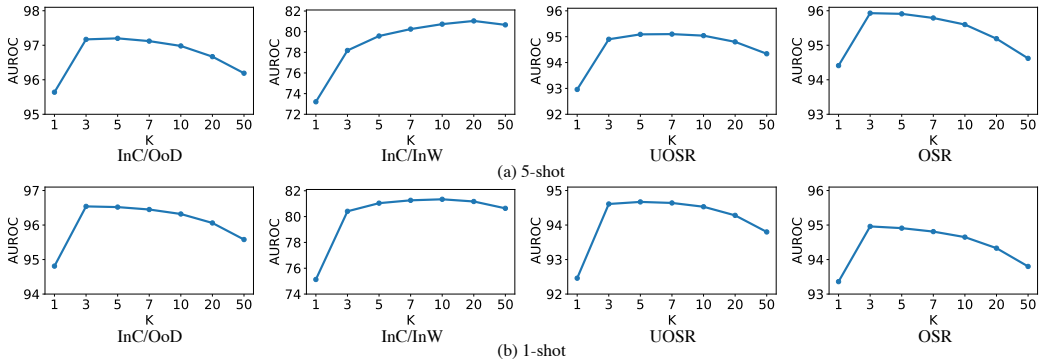

Figure 9: Ablation study of K used in FS-KNN. The backbone is ResNet50.

Analyze of α and λ_{knns}. Two hyper parameters of FS-KNNS including α and λ_{knns} are significant for the performance. λ_{knns} is a threshold to determine when the weight of u_1 becomes large, *i.e.*, $\hat{u}_{fs-knns}$ is important when $u_1 > \lambda_{knns}$. α is to control the change rate of the weight. The ideal λ_{knns} should be located between the uncertainty of InW and OoD samples as shown in Fig. 5 (a) and (c), so that the uncertainty of InW samples is still determined by SoftMax and the uncertainty of OoD samples is enlarged because of FS-KNN. However, which sample is InW or OoD is unknown during test, so we cannot determine the λ_{knns} based on the uncertainty distribution of test samples. (Xia & Bouganis, 2022) proposed to determine the hyper parameters of SIRC based on the training set, but we find that the uncertainty distribution $\hat{u}_{fs-knns}$ of training samples significantly different from the test InC samples, as shown in Fig. 10. Therefore, we seek help from the OoD reference

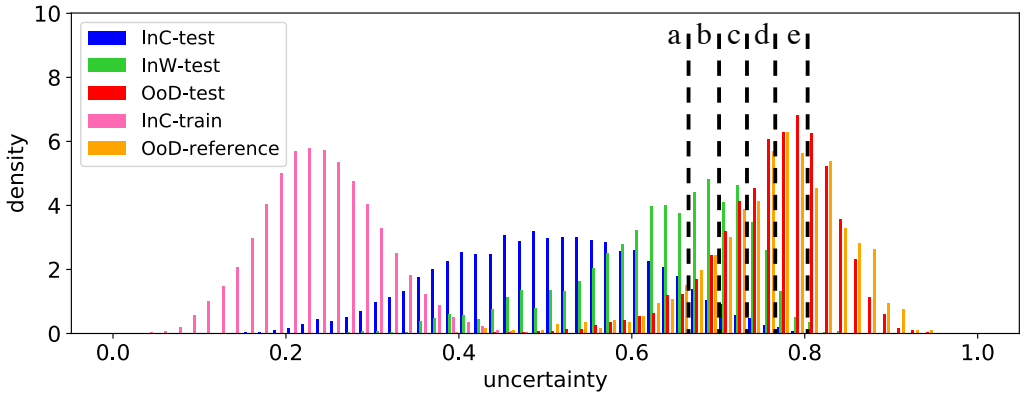

Figure 10: Uncertainty distribution of $\hat{u}_{fs-knns}$. InC-train samples have distinct uncertainty distribution with InC-test samples, but OoD reference samples share similar uncertainty distribution with OoD test samples. a to e correspond to the λ_{knns} when $\beta = 1.5, 1, 0.5, 0, -0.5$ in Eq. 4.

samples, as we find their uncertainty distribution is extremely similar with OoD test samples. We calculate the mean $\overline{u}$ and standard deviation σ of the $\hat{u}_{fs-knns}$ of OoD reference samples, and we aim to find a proper λ_{knns} through:

$$\lambda_{knns} = \overline{u} - \beta \cdot \sigma \tag{4}$$

We draw several λ_{knns} with different β in Fig. 10 (a to e). A smaller λ_{knns} or a larger β means more samples will be influenced by the FS-KNN uncertainty, so that the $\hat{u}_{fs-knns}$ of more OoD test samples will be strengthened by the FS-KNN uncertainty which brings better InC/OoD performance, but meanwhile more InW samples will also be influenced by the FS-KNN uncertainty which brings

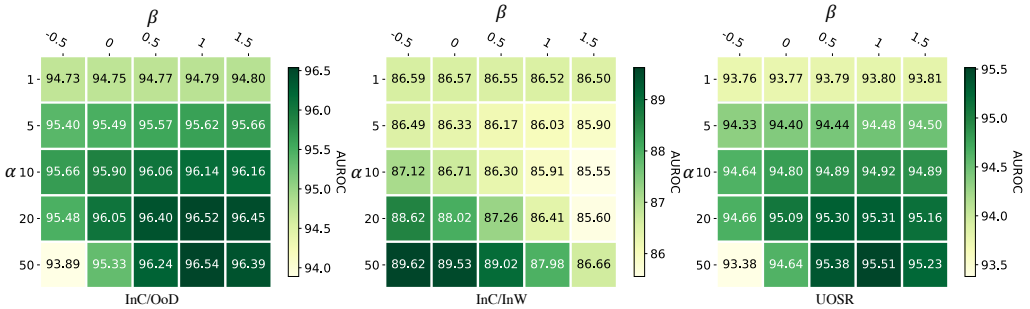

Figure 11: Ablation study of β and α in Eq. 2 and 4.

worse InC/InW performance, as shown in Fig. 11. In other words, λ_{knns} controls the trade-off between InC/OoD and InC/InW performance. Overall, $\beta = 1$ achieves the best UOSR performance which is a balanced result of InC/OoD and InC/InW trade-off. Larger α means the weight of FS-KNN uncertainty grows quickly when FS-KNN uncertainty is closed to λ_{knns}. We find that a smaller α makes the performance insensitive of λ_{knns}, and vice versa. So finally we pick α and β as 50 and 1 as the optimal hyper parameters.