



HAL
open science

Open-Set Likelihood Maximization for Few-Shot Learning

Malik Boudiaf, Etienne Bennequin, Myriam Tami, Antoine Toubhans, Pablo Piantanida, Céline Hudelot, Ismail Ben Ayed

► **To cite this version:**

Malik Boudiaf, Etienne Bennequin, Myriam Tami, Antoine Toubhans, Pablo Piantanida, et al.. Open-Set Likelihood Maximization for Few-Shot Learning. CVPR, 2023, Vancouver, Canada. 10.1109/cvpr52729.2023.02299 . hal-04327201

HAL Id: hal-04327201

<https://hal.science/hal-04327201v1>

Submitted on 6 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Open-Set Likelihood Maximization for Few-Shot Learning

Malik Boudiaf*
ÉTS Montreal

Etienne Bennequin*
Sicara - MICS[†]

Myriam Tami
MICS[†]

Antoine Toubhans
Sicara

Pablo Piantanida
ILLS - MILA - McGill - CNRS

Celine Hudelot
MICS[†]

Ismail Ben Ayed
ÉTS Montreal

Abstract

We tackle the *Few-Shot Open-Set Recognition (FSOSR)* problem, i.e. classifying instances among a set of classes for which we only have a few labeled samples, while simultaneously detecting instances that do not belong to any known class. We explore the popular transductive setting, which leverages the unlabelled query instances at inference. Motivated by the observation that existing transductive methods perform poorly in open-set scenarios, we propose a generalization of the maximum likelihood principle, in which latent scores down-weighting the influence of potential outliers are introduced alongside the usual parametric model. Our formulation embeds supervision constraints from the support set and additional penalties discouraging overconfident predictions on the query set. We proceed with a block-coordinate descent, with the latent scores and parametric model co-optimized alternately, thereby benefiting from each other. We call our resulting formulation *Open-Set Likelihood Optimization (OSLO)*. OSLO is interpretable and fully modular; it can be applied on top of any pre-trained model seamlessly. Through extensive experiments, we show that our method surpasses existing inductive and transductive methods on both aspects of open-set recognition, namely inlier classification and outlier detection. Code is available at <https://github.com/ebennequin/few-shot-open-set>.

1. Introduction

Few-shot classification consists in recognizing concepts for which we have only a handful of labeled examples. These form the *support set*, which, together with a batch of unlabeled instances (the *query set*), constitute a few-shot task.

Most few-shot methods classify the unlabeled query samples of a given task based on their similarity to the support instances in some feature space [36]. This implicitly assumes a *closed-set* setting for each task, i.e. query instances are supposed to be constrained to the set of classes explicitly defined by the support set. However, the real world is open and this closed-set assumption may not hold in practice, especially for limited support sets. Whether they are unexpected items circulating on an assembly line, a new dress not yet included in a marketplace’s catalog, or a previously undiscovered species of fungi, *open-set instances* occur everywhere. When they do, a closed-set classifier will falsely label them as the closest known class.

This drove the research community toward open-set recognition *i.e.* recognizing instances with the awareness that they may belong to unknown classes. In large-scale settings, the literature abounds of methods designed specifically to detect open-set instances while maintaining good accuracy on closed-set instances [1, 32, 51]. Very recently, the authors of [21] introduced a Few-Shot Open-Set Recognition (FSOSR) setting, in which query instances may not belong to any known class. The study in [21], together with other recent follow-up works [15, 16], exposed FSOSR to be a difficult task.

To help alleviate the scarcity of labeled data, transduction [38] was recently explored for few-shot classification [24], and has since become a prominent research direction, fueling a large body of works, e.g. [3, 4, 9, 14, 23, 26, 40, 43, 52], among many others. By leveraging the statistics of the query set, transductive methods yield performances that are substantially better than their inductive counterparts [4, 40] in the standard closed-set setting.

In this work, we seek to explore transduction for the FSOSR setting. We argue that theoretically, transduction has the potential to enable both classification and outlier detection (OD) modules to act symbiotically. Indeed, the classification module can reveal valuable structure of the

*Equal contribution. Corresponding authors: {malik.boudiaf.1@etsmtl.net, etienneb@sicara.com}

[†]MICS, CentraleSupélec, Université Paris-Saclay

inlier’s marginal distribution that the OD module seeks to estimate, such as the number of modes or conditional distributions, while the OD part indicates the “usability” of each unlabelled sample. However, transductive principles currently adopted for few-shot learning heavily rely on the closed-set assumption in the unlabelled data, leading them to match the classification confidence for open-set instances with that of closed-set instances. In the presence of outliers, this not only harms their predictive performance on closed-set instances, but also makes prediction-based outlier detection substantially harder than with simple inductive baselines.

Contributions. In this work, we aim at designing a principled framework that reconciles transduction with the open nature of the FSOSR problem. Our idea is simple but powerful: instead of finding heuristics to assess the *outlierness* of each unlabelled query sample, we treat this score as a latent variable of the problem. Based on this idea, we propose a generalization of the maximum likelihood principle, in which the introduced latent scores weigh potential outliers down, thereby preventing the parametric model from fitting those samples. Our generalization embeds additional supervision constraints from the support set and penalties discouraging overconfident predictions. We proceed with a block-coordinate descent optimization of our objective, with the closed-set soft assignments, *outlierness* scores, and parametric models co-optimized alternately, thereby benefiting from each other. We call our resulting formulation *Open-Set Likelihood Optimization* (OSLO). OSLO provides highly interpretable and closed-form solutions within each iteration for both the soft assignments, *outlierness* variables, and the parametric model. Additionally, OSLO is fully modular; it can be applied on top of any pre-trained model seamlessly.

Empirically, we show that OSLO significantly surpasses its inductive and transductive competitors alike for both outlier detection and closed-set prediction. Applied on a wide variety of architectures and training strategies and without any re-optimization of its parameters, OSLO’s improvement over a strong baseline remains large and consistent. This modularity allows our method to fully benefit from the latest advances in standard image recognition. Before diving into the core content, let us summarize our contributions:

1. To the best of our knowledge, we realize the first study and benchmarking of transductive methods for the Few-Shot Open-Set Recognition setting. We reproduce and benchmark five state-of-the-art transductive methods.
2. We introduce Open-Set Likelihood Optimization (OSLO), a principled extension of the Maximum Likelihood framework that explicitly models and handles the presence of outliers. OSLO is interpretable and modular *i.e.* can be applied on top of any pre-trained model seamlessly.
3. Through extensive experiments spanning five datasets and a dozen of pre-trained models, we show that OSLO consistently surpasses both inductive and existing transductive methods in detecting open-set instances while competing with the strongest transductive methods in classifying closed-set instances.

2. Related Works

Few-shot classification (FSC) methods. Many FSC works involve episodic training [41], in which a neural network acting as a feature extractor is trained on artificial tasks sampled from the training set. This replication of the inference scenario during training is intended to make the learned representation more robust to new classes. However, several recent works have shown that simple fine-tuning baselines are competitive in comparison to sophisticated episodic methods, e.g. [6, 12], motivating a new direction of few-shot learning research towards the development of model-agnostic methods that do not involve any specific training strategy [9].

Transductive FSC. Transductive FSC methods leverage statistics of the query set as unlabeled data to improve performance, through model fine-tuning [9], Laplacian regularization [52], clustering [20], mutual information maximization [4, 40], prototype rectification [23], or optimal transport [2, 14, 18], among other transduction strategies. The idea of maximizing the likelihood of both support and query samples under a model parameterized by class prototypes is proposed by [47] for few-shot segmentation. However, their method relies on the closed-set assumption. Differing from previous works, our framework leverages an additional latent variable, the *inlierness* score.

Open-set recognition (OSR). OSR aims to enable classifiers to detect instances from unknown classes [32]. Prior works address this problem in the large-scale setting by augmenting the SoftMax activation to account for the possibility of unseen classes [1], generating artificial outliers [11, 27], improving closed-set accuracy [39], or using placeholders to anticipate novel classes’ distributions with adaptive decision boundaries [51]. All these methods involve the training of deep neural networks on a specific class set. Therefore, they are not fully fit for the few-shot setting. In this work, we use simple yet effective adaptations of OpenMax [1] and PROSER [51] as strong baselines for FSOSR.

Few-shot open-set recognition. In the few-shot setting, methods must detect open-set instances while only a few closed-set instances are available. [21] use meta-learning on pseudo-open-set tasks to train a model to maximize the classification entropy of open-set instances. [16] use transformation consistency to measure the divergence between a query image and the set of class prototypes. [15] use an attention mechanism to generate a negative prototype for outliers. These methods require the optimization of a separate

model with a specific episodic training strategy.

Nonetheless, as we show in sec:experiments, they bring marginal improvement over simple adaptations of standard OSR methods to the few-shot setting. In comparison, our method doesn't require any specific training and can be plugged into any feature extractor without further optimization.

3. Few-Shot Open-Set Recognition

Model training. Let us denote the raw image space \mathcal{X} . As per the standard Few-Shot setting, we assume access to a *base* dataset $\mathcal{D}_{\text{base}} = \{(\mathbf{x}_i, y_i)\}_{i=1 \dots |\mathcal{D}_{\text{base}}|}$ with base classes \mathbb{C}_{base} , such that $\mathbf{x}_i \in \mathcal{X}$ and $y_i \in \mathbb{C}_{\text{base}}$. We use $\mathcal{D}_{\text{base}}$ to train a feature extractor ϕ_{θ} . Our method developed later in sec:open_set_likelihood, freezes ϕ_{θ} and performs inference directly on top of the extracted features for each task.

Testing. Given a set of *novel* classes $\mathbb{C}_{\text{novel}}$ disjoint from base classes *i.e.* $\mathbb{C}_{\text{novel}} \cap \mathbb{C}_{\text{base}} = \emptyset$, a K -way FSOSR task is formed by sampling a set of K *closed-set* classes $\mathbb{C}_{\text{CS}} \subset \mathbb{C}_{\text{novel}}$, a support set of labeled instances $\mathbb{S} = \{(\mathbf{x}_i, y_i) \in \mathcal{X} \times \mathbb{C}_{\text{CS}}\}_{i=1}^{|\mathbb{S}|}$ and a query set $\mathbb{Q} = \{\mathbf{x}_i \in \mathcal{X}\}_{i=|\mathbb{S}|+1}^{|\mathbb{S}|+|\mathbb{Q}|}$. In the standard few-shot setting, the unknown ground-truth query labels $\{y_i\}_{i=|\mathbb{S}|+1}^{|\mathbb{S}|+|\mathbb{Q}|}$ are assumed to be restricted to closed-set classes *i.e.* $\forall i, y_i \in \mathbb{C}_{\text{CS}}$. In FSOSR, however, query labels may also belong to an additional set $\mathbb{C}_{\text{OS}} \subset \mathbb{C}_{\text{novel}}$ of *open-set* classes *i.e.* $\forall i > |\mathbb{S}|, y_i \in \mathbb{C}_{\text{CS}} \cup \mathbb{C}_{\text{OS}}$ with $\mathbb{C}_{\text{CS}} \cap \mathbb{C}_{\text{OS}} = \emptyset$. For easy referencing, we refer to query samples from the closed-set classes \mathbb{C}_{CS} as *inliers* and to query samples from open-set classes \mathbb{C}_{OS} as *outliers*. For each query image \mathbf{x}_i , the goal of FSOSR is to simultaneously assign a closed-set prediction and an *outlierness* (or equivalently *inlierness*) score.

Transductive FSOSR. As a growing part of the Few-Shot literature, Transductive Few-Shot Learning assumes that unlabelled samples from the query set are observed at once, such that the structure of unlabelled data can be leveraged to help constrain ambiguous few-shot tasks. Transductive methods have achieved impressive improvements over inductive methods in standard closed-set FSC [4, 9, 14, 52]. We expect that transductive methods can help us improve overall open-set performance. While we find this to generally hold for closed-set predictive performance, we empirically show in sec:experiments that accuracy gains systematically come along significant outlier detection degradation, indicating that transductive methods are not equipped to handle open-set recognition. In the following, we take up the challenge of designing a transductive optimization framework that leverages the presence of outliers to improve its performance.




Figure 1. **Intuition behind OSLO.** Standard transductive likelihood (left) tries to enforce high likelihood for all points, including outliers. OSLO (right) instead treats the *outlierness* of each sample as a latent variable to be solved alongside the parametric model. Besides yielding a principled *outlierness* score for open-set detection, it also allows the fitted parametric model to effectively disregard samples deemed outliers, and therefore provide a better approximation of underlying class-conditional distributions.

4. Open-Set Likelihood

In this section, we introduce OSLO, a novel extension of the standard likelihood designed for transductive FSOSR. Unlike existing transductive methods, OSLO explicitly models and handles the potential presence of outliers, which allows it to outperform inductive baselines on both aspects of the open-set scenario.

Observed variables. We start by establishing the observed variables of the problem. As per the traditional setting, we observe images from the support set $\{\mathbf{x}_i\}_{i=1}^{|\mathbb{S}|}$ and their associated labels $\{y_i\}_{i=1}^{|\mathbb{S}|}$. The transductive setting also allows us to observe images from the query set. For notation convenience, we concatenate all images in $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^{|\mathbb{S}|+|\mathbb{Q}|}$.

Latent variables. Our goal is to predict the class of each sample in the query set \mathbb{Q} , as well as their *inlierness*, *i.e.* the model's belief in a sample being an inlier or not. This naturally leads us to consider latent class assignments $\mathbf{z}_i \in \Delta^K$ describing the membership of sample i to each closed-set class, with $\Delta^K = \{\mathbf{z} \in [0, 1]^K : \mathbf{z}^T \mathbf{1} = 1\}$ the K -dimensional simplex. Additionally, we consider latent *inlierness* scores $\xi_i \in [0, 1]$ close to 1 if the model considers sample i as an inlier. For notation convenience, we consider latent assignments and *inlierness* scores for all samples, including those from the support, and group everything in $\mathbf{Z} = \{\mathbf{z}_i\}_{i=1}^{|\mathbb{S}|+|\mathbb{Q}|}$ and $\boldsymbol{\xi} = \{\xi_i\}_{i=1}^{|\mathbb{S}|+|\mathbb{Q}|}$. Note that support samples are inliers, and we know their class. Therefore $\forall i \leq |\mathbb{S}|$, the constraints $\mathbf{z}_i = \mathbf{y}_i$ and $\xi_i = 1$ will be imposed, where \mathbf{y}_i is the one-hot encoded version of y_i .

Parametric model. The final ingredient we need to formulate is a parametric joint model over observed features and assignments. Following standard practice, we model the joint distribution as a balanced mixture of standard Gaussian distri-

butions, parameterized by the centroids $\boldsymbol{\mu} = \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K\}$:

$$p(\mathbf{x}, k; \boldsymbol{\mu}) = p(k)p(\mathbf{x}|k) \propto \exp\left(-\frac{\|\phi_{\theta}(\mathbf{x}) - \boldsymbol{\mu}_k\|^2}{2}\right) \quad (1)$$

As mentioned in sec:fsosr_setting, the feature extractor's parameters θ are kept frozen, and only $\boldsymbol{\mu}$ will be optimized.

Objective. Using the i.i.d. assumption, we start by writing the standard likelihood objective:

$$p(\mathbf{X}, \mathbf{Z}; \boldsymbol{\mu}) = \prod_{i=1}^{|\mathcal{S}|+|\mathcal{Q}|} \prod_{k=1}^K p(\mathbf{x}_i, k; \boldsymbol{\mu})^{z_{ik}} \quad (2)$$

Without loss of generality, we consider the log-likelihood:

$$\log(p(\mathbf{X}, \mathbf{Z}; \boldsymbol{\mu})) = \sum_{i=1}^{|\mathcal{S}|+|\mathcal{Q}|} \sum_{k=1}^K z_{ik} \log(p(\mathbf{x}_i, k; \boldsymbol{\mu})) \quad (3)$$

Eq. (3) tries to enforce a high likelihood of all samples under our parametric model p . This becomes sub-optimal in the presence of outliers, which should ideally be disregarded. fig:likelihood illustrates this phenomenon on a toy 2D drawing. To downplay this issue, we introduce *Open-Set Likelihood Optimization* (OSLO), a generalization of the standard likelihood framework, which leverages latent *inlierness* scores to weigh samples:

$$\mathcal{L}_O(\mathbf{X}, \mathbf{Z}, \boldsymbol{\xi}; \boldsymbol{\mu}) = \sum_{i=1}^{|\mathcal{S}|+|\mathcal{Q}|} \xi_i \sum_{k=1}^K z_{ik} \log(p(\mathbf{x}_i, k; \boldsymbol{\mu})) \quad (4)$$

Eq (4) can be interpreted as follows: samples believed to be inliers *i.e.* $\xi_i \approx 1$ will be required to have high likelihood under our model p , whereas outliers won't. Note that $\boldsymbol{\xi}$ is treated as a variable of optimization, and is co-optimized alongside $\boldsymbol{\mu}$ and \mathbf{Z} . Finally, to prevent overconfident latent scores, we consider a *penalty* term on both \mathbf{Z} and $\boldsymbol{\xi}$:

$$\mathcal{L}_{\text{soft}} = \sum_{i=|\mathcal{S}|+1}^{|\mathcal{S}|+|\mathcal{Q}|} \lambda_z \mathcal{H}(z_i) + \lambda_{\xi} \mathcal{H}(\xi_i) \quad (5)$$

where $\boldsymbol{\xi}_i = [1 - \xi_i, \xi_i]$, and $\mathcal{H}(\mathbf{p}) = -\mathbf{p}^{\top} \log(\mathbf{p})$ denotes the entropy, which encourages smoother assignments.

Optimization. We are now ready to formulate OSLO's optimization problem:

$$\begin{aligned} \max_{\boldsymbol{\mu}, \mathbf{Z}, \boldsymbol{\xi}} \quad & \mathcal{L}_O(\mathbf{Z}, \boldsymbol{\xi}, \boldsymbol{\mu}) + \mathcal{L}_{\text{soft}}(\mathbf{Z}, \boldsymbol{\xi}) \\ \text{s.t} \quad & \mathbf{z}_i \in \Delta^K, \quad \xi_i \in [0, 1] \quad \forall i \\ & \mathbf{z}_i = \mathbf{y}_i, \quad \xi_i = 1, \quad i \leq |\mathcal{S}| \end{aligned} \quad (6)$$

Problem (6) is strictly convex with respect to each variable when the other variables are fixed. Therefore, we proceed with a block-coordinate ascent, which alternates three iterative steps, each corresponding to a closed-form solution for one of the variables.

Proposition 1. OSLO's optimization problem (6) can be minimized by alternating the following updates, with σ denoting the sigmoid operation:

$$\xi_i^{(t+1)} = \begin{cases} 1 & \text{if } i \leq |\mathcal{S}| \\ \sigma\left(\frac{1}{\lambda_{\xi}} \sum_{k=1}^K z_{ik}^{(t)} \log p(\mathbf{x}_i, k; \boldsymbol{\mu}^{(t)})\right) & \text{else} \end{cases}$$

$$\mathbf{z}_i^{(t+1)} \propto \begin{cases} \mathbf{y}_i & \text{if } i \leq |\mathcal{S}| \\ \exp\left(\frac{\xi_i^{(t+1)}}{\lambda_z} \log p(\mathbf{x}_i, \cdot; \boldsymbol{\mu}^{(t)})\right) & \text{else} \end{cases}$$

$$\boldsymbol{\mu}_k^{(t+1)} = \frac{1}{\sum_{i=1}^{|\mathcal{S}|+|\mathcal{Q}|} \xi_i^{(t+1)} z_{ik}^{(t+1)}} \sum_{i=1}^{|\mathcal{S}|+|\mathcal{Q}|} \xi_i^{(t+1)} z_{ik}^{(t+1)} \phi_{\theta}(\mathbf{x}_i)$$

The proof of proposition 1 is performed by derivation of $\mathcal{L}_O(\mathbf{Z}, \boldsymbol{\xi}, \boldsymbol{\mu}) + \mathcal{L}_{\text{soft}}(\mathbf{Z}, \boldsymbol{\xi})$ and deferred to the supplementary material. The optimal solution for the *inlierness* score ξ_i appears very intuitive, and essentially conveys that samples with high likelihood under the current parametric model should be considered inliers. We emphasize that **beyond providing a principled outlierness score, as $1 - \xi_i$, the presence of ξ_i allows to refine and improve the closed-set parametric model.** In particular, ξ_i acts as a sample-wise temperature in the update of \mathbf{z}_i , encouraging outliers ($\xi_i \approx 0$) to have a uniform distribution over closed-set classes. Additionally, those samples contribute less to the update of closed-set prototypes $\boldsymbol{\mu}$, as each sample's contribution is weighted by ξ_i .

5. Experiments

5.1. Experimental setup

Baselines. One goal of this work is to fairly evaluate different strategies to address the FSOSR problem. In particular, we benchmark 4 families of methods: (i) popular Outlier Detection methods, e.g. Nearest-Neighbor [28], (ii) Inductive Few-Shot classifiers, e.g. SimpleShot [42] (iii) Inductive Open-Set methods formed by standard methods such as OpenMax [1] and Few-Shot methods such as Snatcher [16] (iv) Transductive classifiers, e.g. TIM [4], that implicitly rely on the closed-set assumption, and finally (v) Transductive Open-Set introduced in this work through OSLO. Following [16], closed-set few-shot classifiers are turned into open-set classifiers by considering the negative of the maximum probability as a measure of outlierness. Furthermore, we found that applying a center-normalize transformation $\psi_{\mathbf{v}} : \mathbf{x} \mapsto (\mathbf{x} - \mathbf{v}) / \|\mathbf{x} - \mathbf{v}\|_2$ on the features extracted by ϕ_{θ} benefited all methods. Therefore, we apply it to the features before applying any method, using an inductive *Base centering* [42] for inductive methods $\mathbf{v}_{\text{Base}} = \frac{1}{|\mathcal{D}_{\text{base}}|} \sum_{\mathbf{x} \in \mathcal{D}_{\text{base}}} \phi_{\theta}(\mathbf{x})$, and a transductive

Table 1. **Standard Benchmarking.** Evaluating different families of methods on the FSOSR problem on *mini-ImageNet* and *tiered-ImageNet* using a ResNet-12. For each column, a light-gray standard deviation is indicated, corresponding to the maximum deviation observed across methods for that metric. Best methods are shown in bold. Results marked with * are reported from their original paper.

		<i>mini-ImageNet</i>							
Strategy	Method	1-shot				5-shot			
		Acc <small>±0.72</small>	AUROC <small>±0.79</small>	AUPR <small>±0.69</small>	Prec@0.9 <small>±0.47</small>	Acc <small>±0.44</small>	AUROC <small>±0.73</small>	AUPR <small>±0.61</small>	Prec@0.9 <small>±0.56</small>
OOD detection	<i>k</i> -NN [28]	-	70.86	70.43	58.23	-	76.22	76.36	61.48
	IForest [22]	-	55.59	55.24	52.18	-	62.80	61.62	54.77
	OCVSM [33]	-	69.67	69.71	57.35	-	68.49	65.60	59.24
	PCA [35]	-	67.23	66.50	56.67	-	75.24	75.53	60.73
	COPOD [19]	-	50.60	51.85	50.92	-	51.63	52.65	51.31
	HBOS	-	58.26	57.41	53.06	-	61.11	60.18	54.30
Inductive classifiers	SimpleShot [42]	65.90	64.99	63.78	55.77	81.72	70.61	70.06	57.91
	Baseline ++ [6]	65.81	65.15	63.85	55.87	81.86	66.37	65.58	56.33
	FEAT [48]	67.23	52.45	54.44	50.00	82.00	53.25	56.48	50.00
Inductive Open-Set	PEELER* [21]	65.86	60.57	-	-	80.61	67.35	-	-
	TANE-G* [15]	68.11	72.41	-	-	83.12	79.85	-	-
	SnatcherF [16]	67.23	70.10	69.74	58.02	82.00	76.57	76.97	61.64
	OpenMax [1]	65.90	71.34	70.86	58.67	82.23	77.42	77.63	62.35
	PROSER [51]	65.00	68.93	68.84	57.03	80.08	74.98	75.58	60.11
Transductive classifiers	LaplacianShot [52]	70.59	53.13	54.59	52.06	82.94	57.17	57.90	52.56
	BDCSPN [23]	69.35	57.95	58.58	52.71	82.66	61.27	62.17	53.26
	TIM-GD [4]	67.53	62.46	61.05	54.83	82.49	67.19	66.15	56.70
	PT-MAP [14]	66.32	59.05	58.67	53.74	78.12	62.78	62.48	54.67
	LR-ICI [43]	68.24	49.96	51.61	50.45	81.77	51.82	53.49	50.80
Transductive Open-Set	OSLO (ours)	71.73	74.92	74.61	60.95	83.40	82.59	82.34	66.98
		<i>tiered-ImageNet</i>							
		<small>±0.74</small>	<small>±0.76</small>	<small>±0.71</small>	<small>±0.52</small>	<small>±0.52</small>	<small>±0.68</small>	<small>±0.75</small>	<small>±0.57</small>
OOD detection	<i>k</i> -NN [28]	-	73.97	73.15	60.74	-	80.22	80.06	65.47
	IForest [22]	-	54.57	54.24	51.85	-	62.31	60.82	54.72
	OCVSM [33]	-	71.22	71.17	58.81	-	71.20	68.23	61.09
	PCA [35]	-	68.30	67.02	57.66	-	76.26	76.41	61.81
	COPOD [19]	-	50.87	51.95	51.07	-	52.62	53.48	51.44
	HBOS	-	57.54	56.67	52.98	-	60.91	59.95	54.15
Inductive classifiers	SimpleShot [42]	70.27	69.78	67.89	58.54	84.94	77.38	76.28	63.21
	Baseline ++ [6]	70.21	69.73	67.80	58.50	85.10	73.77	72.39	61.05
	FEAT [48]	69.94	52.49	56.74	50.00	83.96	53.30	59.81	50.00
Inductive Open-Set	PEELER* [21]	69.51	65.20	-	-	84.10	73.27	-	-
	TANE-G* [15]	70.58	73.53	-	-	85.38	81.54	-	-
	SnatcherF [16]	69.94	74.02	73.33	60.79	83.96	81.90	81.67	66.89
	OpenMax [1]	70.27	72.40	71.91	59.91	85.79	77.91	78.42	63.07
	PROSER [51]	68.48	70.07	69.87	57.99	83.34	75.84	76.56	61.12
Transductive classifiers	LaplacianShot [52]	75.66	57.82	58.41	53.67	86.23	63.75	63.65	55.36
	BDCSPN [23]	74.07	62.13	61.84	54.53	85.65	67.41	67.57	56.30
	TIM-GD [4]	72.56	68.08	65.97	57.84	85.70	74.67	73.06	61.59
	PT-MAP [14]	71.13	64.48	62.94	56.25	82.81	71.08	69.89	59.11
	LR-ICI [43]	73.80	49.32	51.41	50.35	85.21	51.65	53.85	50.79
Transductive Open-Set	OSLO (ours)	76.64	79.06	79.07	64.36	86.35	86.92	87.28	71.98




Figure 2. **OSLO improves open-set performances on a wide variety of tasks.** Relative 1-shot performance of the best methods of each family w.r.t the *Strong baseline* using a ResNet-12, across a set of 5 scenarios, including 3 with domain-shift. Each vertex represents one scenario, e.g. *tiered*→Fungi (x) means the feature extractor was pre-trained on *tiered*-ImageNet, test tasks are sampled from Fungi, and the *Strong Baseline* performance is x . For each method, the average relative improvement across the 5 scenarios is reported in parenthesis in the legend. The same charts are provided in the supplementary materials for the 5-shot setting and using a WideResNet backbone.

Task centering [14] $v_{Task} = \frac{1}{|\mathcal{S}_{UQ}|} \sum_{x \in \mathcal{S}_{UQ}} \phi_{\theta}(x)$ for all transductive methods. Since features are normalized, we empirically found it beneficial to re-normalize centroids $\mu_k \leftarrow \mu_k / \|\mu_k\|_2$ after each update from Prop. 1, which we show in the Appendix remains a valid minimizer of Eq. (6) when adding the constraint $\|\mu_k\|_2 = 1$.

Hyperparameters. For all methods, we define a grid over salient hyper-parameters and tune over the validation split of *mini*-ImageNet. To avoid cumbersome per-dataset tuning, and evaluate the generalizability of methods, we then keep hyper-parameters fixed across all other experiments.

Architectures and checkpoints. To provide the fairest comparison, all non-episodic methods are tuned and tested using off-the-shelf pre-trained checkpoints. All results except fig:barplots are produced using the pre-trained ResNet-12 and Wide-ResNet 28-10 checkpoints provided by the authors from [48]. As for episodically-finetuned models required by Snatcher [16] and FEAT [48], checkpoints are obtained from the authors’ respective repositories. Finally, to challenge the model-agnosticity of our method, we resort to an additional set of 10 ImageNet pre-trained models covering three distinct architectures: ResNet-50 [13] for CNNs, ViT-B/16 [10] for vision transformers, and Mixer-B/16 [37] for MLP-Mixer. These models are taken from the excellent TIMM library [45].

Datasets and tasks. We experiment with a total of 5 vision datasets. As standard FSC benchmarks, we use the *mini*-ImageNet [41] dataset with 100 classes and the larger *tiered*-ImageNet [29] dataset with 608 classes. We also experiment on more challenging cross-domain tasks formed by using 3 finer-grained datasets: the Caltech-UCSD Birds 200 [44] (CUB) dataset, with 200 classes, the FGVC-Aircraft dataset [25] with 100 classes, and the Fungi classification challenge [34] with 1394 classes. Following standard FSOSR protocol, support sets contain $|\mathcal{C}_{CS}| = 5$ closed-set classes with 1 or 5

instances, or *shots*, per class, and query sets are formed by sampling 15 instances per class, from a total of ten classes: the five closed-set classes and an additional set of $|\mathcal{C}_{OS}| = 5$ open-set classes. We follow this setting for a fair comparison with previous works [16, 21] which sample open-set query instances from only 5 classes. We also report results in supplementary materials for a more general setting in which open-set query instances are sampled indifferently from all remaining classes in the test set.

5.2. Results

Simplest inductive methods are competitive. The first surprising result comes from analyzing the performances of standard OOD detectors on the FSOSR problem. Fig. 1 shows that k -NN and PCA outperform, by far, arguably more advanced methods that are OCVSM and Isolation Forest. This result contrasts with standard high-dimensional benchmarks [50] where k -NN falls typically short of the latter, indicating that the very difficult challenge posed by FSOSR may lead advanced methods to overfit. In fact, Fig. 2 shows that across 5 scenarios, the combination SimpleShot [42]+ k -NN [28] formed by the simplest FS-inductive classifier and the simplest inductive OOD detector is a strong baseline that outperforms all specialized open-set methods. We refer to this combination as *Strong baseline* in Figures 2 and 4. Additional results for the Wide-ResNet architecture are provided in the supplementary material.

Transductive methods still improve accuracy but degrade outlier detection. As shown in Table 1, most transductive classifiers still offer a significant boost in closed-set accuracy, even in the presence of outliers in the query set. Note that this contrasts with findings from the semi-supervised literature, where standard methods drop below the baseline in the presence of even a small fraction of outliers [8, 17, 31, 49]. We hypothesize that the deliberate under-parametrization of




Figure 3. **OSLO improves performance even with few queries.** We study the closed-set (accuracy) and open-set (AUROC) performance of transductive methods depending on the size of the query set on *tiered*-ImageNet in the 1-shot and 5-shot settings. The total size $|Q|$ of the query set is obtained by multiplying the number of queries per class N_Q by the number of classes in the task (*i.e.* 5) and adding as many outlier queries *e.g.* $N_Q = 1$ corresponds to 1 query per class and 5 open-set queries *i.e.* $|Q| = 10$. We add the inductive method k -NN + SimpleShot to compare with a method that is by nature independent of the number of queries. The results for *mini*-ImageNet are provided in the supplementary materials.

few-shot methods –typically only training a linear classifier–, required to avoid overfitting the support set, partly explains such robustness. However, transductive methods still largely underperform in outlier detection, with AUROCs as low as 52 % (50% being a random detector) for LaplacianShot. Note that the *outlierness* score for these methods is based on the negative of the maximum probability, therefore this result can be interpreted as transductive methods having artificially matched the prediction confidence for outliers with the confidence for inliers.

OSLO achieves the best trade-off. Benchmark results in Fig. 1 show that OSLO surpasses the best transductive methods in terms of closed-set accuracy, while consistently outperforming existing out-of-distribution and open-set detection competitors on outlier detection ability. Interestingly, while the gap between closed-set accuracy of transductive methods and inductive ones typically contracts with more shots, the outlier detection performance of OSLO remains largely superior to its inductive competitors even in the 5-shot scenario, where a consistent 3-6% gap in AUROC and AUPR with the second-best method can be observed. We accumulate further evidence of OSLO’s superiority by introducing 3 additional cross-domain scenarios in Fig. 2, corresponding to a base model pre-trained on *tiered*-ImageNet, but tested on CUB, Aircraft, and Fungi datasets. In such challenging scenarios, where both feature and class distributions shift, OSLO remains competitive in closed-set accuracy and largely outperforms other methods in outlier detection.

OSLO benefits from more query samples. A critical question for transductive methods is the dependency of their performance on the size of the query set. Intuitively, a larger query set will provide more unlabeled data and thus lead to better results. We exhibit this relation in fig:variate-query-

Table 2. **OSLO’s ablation study** along two factors described in `sec:ablation_study`. Results are reproduced on the 1 – shot scenario on *mini* – ImageNet, with a ResNet – 12.

(i) <i>Inlierness</i> latent	Acc	AUROC
Ignore (3)	69.42	64.97
Leverage (4)	71.73	74.92
(ii) Optimization steps	Acc	AUROC
At initialization	66.63	71.76
After optimization	71.73	74.92

tiered by spanning the number of queries per class from 1 to 30. We observe that the closed-set accuracy of most transductive methods is stable across this span in the 5-shot scenario. In the 1-shot scenario, OSLO gains from additional queries but stays above the baseline even with a small number of queries. Interestingly enough, OSLO is the only transductive method to improve its outlier detection ability when the number of queries increases.

OSLO steps toward model-agnosticity. We evaluate OSLO’s *model-agnosticity* by its ability to maintain consistent improvement over the *Strong Baseline*, regardless of the model used, and without hyperparameter adjustment. In that regard, we depart from the standard ResNet-12 and cover 3 largely distinct architectures, each encoding different inductive biases. To further strengthen our empirical demonstration of OSLO’s model-agnosticity, for each architecture, we consider several training strategies spanning different paradigms – unsupervised, supervised, semi and semi-weakly supervised – and using different types of data –image, text–. Results in fig:barplots show the relative improvement of OSLO w.r.t the strong baseline in the 1-shot




Figure 4. **OSLO’s improvement is consistent across many architectures and training strategies.** To evaluate model-agnosticity, we compare OSLO to the Strong baseline on challenging 1-shot Fungi tasks. We experiment across 3 largely distinct architectures: ResNet-50 (CNN) [13], ViT-B/16 (Vision Transformer) [10], and Mixer-B/16 (MLP-Mixer) [37]. For each architecture, we include different types of pre-training, including Supervised (Sup.), Semi-Supervised, Semi-Weakly Supervised (SW Sup.) [46], DINO [5], SAM [7], MIIL [30]. Improvements over the baseline are consistently significant and generally higher than those observed with the ResNet-12 in fig:spider_{charts}.

scenario on the * \rightarrow Fungi benchmark. Without any tuning, OSLO remains able to leverage the strong expressive power of large-scale models, and even consistently widens the gap with the strong baseline, achieving a remarkable performance of 79% accuracy and 81% AUROC with the ViT-B/16. This set of results testifies to how easy obtaining highly competitive results on difficult specialized tasks can be by combining OSLO with the latest models.

Ablation study. We perform an ablation study on the important ingredients of OSLO. As a core contribution of our work, we show in Tab. 2 that the presence *and* optimization of the latent *inlieriness* scores is crucial. In particular, the closed-form latent score ξ yields strong outlier recognition performance, even at *initialization* (i.e. after the very first update from Prop. 1). Interestingly, refining the parametric model without accounting for ξ in \mathbf{Z} and μ ’s updates (i.e. standard likelihood) allows the model to fit those outliers, leading to significantly worse outlier detection, from 71.76% to 64.97%. On the other hand, accounting for ξ , as proposed in OSLO, improves the outlier detection by more than 3% over the initial state, and closed-set accuracy by more than 5%. In the end, in a fully apples-to-apples comparison, OSLO outperforms its standard likelihood counterpart by more than 2% in accuracy and 10% in outlier detection. We strengthen this ablation study in the Appendix.

6. Discussion

Limitations. Unlike inductive methods, transductive methods are inevitably affected by the amount of unlabelled data

provided, which in real-world scenarios cannot necessarily be controlled. OSLO is no exception, and fewer query samples tend to decrease its performance. In the extreme case with only 1 sample per class, OSLO’s performance comes close to our inductive baseline. In those scenarios where unlabelled data is particularly scarce, the benefits brought by transduction remain therefore limited. As a second limitation, poorer representations appear to diminish OSLO’s competitive advantage in closed-set accuracy. In particular, OSLO’s closed-set accuracy stands more than 6% above the baseline on the *tiered* \rightarrow *tiered* scenario but reduces to 2% in the most challenging-domain scenario *tiered* \rightarrow *Aircraft*. Fig. 4 further corroborates this hypothesis, with OSLO’s accuracy outperforming the baseline’s by 9% with the best transformer, but by only 2.7% on the least performing model.

Conclusion. We presented OSLO, the first transductive method for FSOSR. OSLO extends the vanilla maximum likelihood objective in two important ways. First, it accounts for the constraints imposed by the provided supervision. More importantly, it explicitly models the potential presence of outliers in its very latent model, allowing it to co-learn the optimal closed-set model and outlier assignments. Beyond FSOSR, we believe OSLO presents a general, conceptually simple, and completely modular formulation to leverage unlabelled data in the potential presence of outliers. That, of course, naturally extends to other classification settings, such as large-scale open-set detection, but to other tasks as well, such as segmentation in which *background* pixels could be viewed as *outliers* with respect to closed-set classes. We hope OSLO inspires further work in that direction.

References

- [1] Abhijit Bendale and Terrance E Boult. Towards open set deep networks. In *Computer Vision and Pattern Recognition Conference (CVPR)*, 2016.
- [2] Etienne Bennequin, Victor Bouvier, Myriam Tami, Antoine Toubhans, and Céline Hudelot. Bridging few-shot learning and adaptation: New challenges of support-query shift. *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*, 2021.
- [3] Malik Boudiaf, Hoel Kervadec, Ziko Imtiaz Masud, Pablo Piantanida, Ismail Ben Ayed, and Jose Dolz. Few-shot segmentation without meta-learning: A good transductive inference is all you need? In *Computer Vision and Pattern Recognition Conference (CVPR)*, 2021.
- [4] Malik Boudiaf, Ziko Imtiaz Masud, Jérôme Rony, José Dolz, Pablo Piantanida, and Ismail Ben Ayed. Transductive information maximization for few-shot learning. In *Neural Information Processing Systems (NeurIPS)*, 2020.
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021.
- [6] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *International Conference on Learning Representations (ICLR)*, 2019.
- [7] Xiangning Chen, Cho-Jui Hsieh, and Boqing Gong. When vision transformers outperform resnets without pre-training or strong data augmentations. *International Conference on Learning Representations (ICLR)*, 2022.
- [8] Yanbei Chen, Xiatian Zhu, Wei Li, and Shaogang Gong. Semi-supervised learning under class distribution mismatch. In *Conference on Artificial Intelligence (AAAI)*, 2020.
- [9] Guneet Singh Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. A baseline for few-shot image classification. In *International Conference on Learning Representations (ICLR)*, 2020.
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations (ICLR)*, 2021.
- [11] ZongYuan Ge, Sergey Demyanov, Zetao Chen, and Rahil Garnavi. Generative openmax for multi-class open set classification. *arXiv*, 2017.
- [12] Micah Goldblum, Steven Reich, Liam Fowl, Renkun Ni, Valeriia Cherepanova, and Tom Goldstein. Unraveling meta-learning: Understanding feature representations for few-shot tasks. In *International Conference on Machine Learning (ICML)*, 2020.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition Conference (CVPR)*, 2016.
- [14] Yuqing Hu, Vincent Gripon, and Stéphane Pateux. Leveraging the feature distribution in transfer-based few-shot learning. In *International Conference on Artificial Neural Networks*, 2021.
- [15] Shiyuan Huang, Jiawei Ma, Guangxing Han, and Shih-Fu Chang. Task-adaptive negative envision for few-shot open-set recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7171–7180, 2022.
- [16] Minki Jeong, Seokeon Choi, and Changick Kim. Few-shot open-set recognition by transformation consistency. In *Computer Vision and Pattern Recognition Conference (CVPR)*, 2021.
- [17] Krishnateja Killamsetty, Xujiang Zhao, Feng Chen, and Rishabh Iyer. Retrieve: Coreset selection for efficient and robust semi-supervised learning. *Neural Information Processing Systems (NeurIPS)*, 34, 2021.
- [18] Michalis Lazarou, Tania Stathaki, and Yannis Avrithis. Iterative label cleaning for transductive and semi-supervised few-shot learning. In *ICCV*, 2021.
- [19] Zheng Li, Yue Zhao, Nicola Botta, Cezar Ionescu, and Xiyang Hu. Copod: copula-based outlier detection. In *International Conference on Data Mining (ICDM)*, 2020.
- [20] Moshe Lichtenstein, Prasanna Sattigeri, Rogerio Feris, Raja Giryes, and Leonid Karlinsky. Tafssl: Task-adaptive feature sub-space learning for few-shot classification. In *European Conference on Computer Vision (ECCV)*, 2020.
- [21] Bo Liu, Hao Kang, Haoxiang Li, Gang Hua, and Nuno Vasconcelos. Few-shot open-set recognition using meta-learning. In *Computer Vision and Pattern Recognition Conference (CVPR)*, 2020.
- [22] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *International Conference on Data Mining (ICMD)*, 2008.
- [23] Jinlu Liu, Liang Song, and Yongqiang Qin. Prototype rectification for few-shot learning. In *European Conference on Computer Vision (ECCV)*, 2020.
- [24] Yanbin Liu, Juho Lee, Minseop Park, Saehoon Kim, Eunho Yang, Sung Ju Hwang, and Yi Yang. Learning to propagate labels: Transductive propagation network for few-shot learning. *International Conference on Learning Representations (ICLR)*, 2019.
- [25] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv*, 2013.
- [26] Ségolène Martin, Malik Boudiaf, Emilie Chouzenoux, Jean-Christophe Pesquet, and Ismail Ben Ayed. Towards practical few-shot query sets: Transductive minimum description length inference. 2022.
- [27] Lawrence Neal, Matthew Olson, Xiaoli Fern, Weng-Keen Wong, and Fuxin Li. Open set learning with counterfactual images. In *European Conference on Computer Vision (ECCV)*, 2018.
- [28] Sridhar Ramaswamy, Rajeev Rastogi, and Kyuseok Shim. Efficient algorithms for mining outliers from large data sets. In *International Conference on Management of Data*, 2000.
- [29] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and

- Richard S Zemel. Meta-learning for semi-supervised few-shot classification. *International Conference on Learning Representations (ICLR)*, 2018.
- [30] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for the masses. *Neural Information Processing Systems (NeurIPS)*, 2021.
- [31] Kuniaki Saito, Donghyun Kim, and Kate Saenko. Open-match: Open-set semi-supervised learning with open-set consistency regularization. *Neural Information Processing Systems (NeurIPS)*, 2021.
- [32] Walter J Scheirer, Anderson de Rezende Rocha, Archana Sapkota, and Terrance E Boulton. Toward open set recognition. *PAMI*, 2012.
- [33] Bernhard Schölkopf, John C Platt, John Shawe-Taylor, Alex J Smola, and Robert C Williamson. Estimating the support of a high-dimensional distribution. *Neural computation*, 2001.
- [34] Brigit Schroeder and Yin Cui. Fgvcx fungi classification challenge 2018. 2018.
- [35] Mei-Ling Shyu, Shu-Ching Chen, Kanoksri Sarinnapakorn, and LiWu Chang. A novel anomaly detection scheme based on principal component classifier. Technical report, 2003.
- [36] Jake Snell, Kevin Swersky, and Richard S Zemel. Prototypical networks for few-shot learning. *Neural Information Processing Systems (NeurIPS)*, 2017.
- [37] Ilya O Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, et al. Mlp-mixer: An all-mlp architecture for vision. *Neural Information Processing Systems (NeurIPS)*, 2021.
- [38] Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 2013.
- [39] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Open-set recognition: A good closed-set classifier is all you need. In *International Conference on Learning Representations (ICLR)*, 2022.
- [40] Olivier Veilleux, Malik Boudiaf, Pablo Piantanida, and Ismail Ben Ayed. Realistic evaluation of transductive few-shot learning. *Neural Information Processing Systems (NeurIPS)*, 2021.
- [41] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *Neural Information Processing Systems (NeurIPS)*, 2016.
- [42] Yan Wang, Wei-Lun Chao, Kilian Q Weinberger, and Laurens van der Maaten. Simpleshot: Revisiting nearest-neighbor classification for few-shot learning. *arXiv*, 2019.
- [43] Yikai Wang, Chengming Xu, Chen Liu, Li Zhang, and Yanwei Fu. Instance credibility inference for few-shot learning. In *Computer Vision and Pattern Recognition Conference (CVPR)*, 2020.
- [44] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-ucsd birds 200. 2010.
- [45] Ross Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019.
- [46] I Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. Billion-scale semi-supervised learning for image classification. *arXiv*, 2019.
- [47] Boyu Yang, Chang Liu, Bohao Li, Jianbin Jiao, and Qixiang Ye. Prototype mixture models for few-shot semantic segmentation. In *European Conference on Computer Vision*, pages 763–778. Springer, 2020.
- [48] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions. In *Computer Vision and Pattern Recognition Conference (CVPR)*, 2020.
- [49] Qing Yu, Daiki Ikami, Go Irie, and Kiyoharu Aizawa. Multi-task curriculum framework for open-set semi-supervised learning. In *European Conference on Computer Vision (ECCV)*, 2020.
- [50] Yue Zhao, Zain Nasrullah, and Zheng Li. Pyod: A python toolbox for scalable outlier detection. *JMLR*, 2019.
- [51] Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Learning placeholders for open-set recognition. In *Computer Vision and Pattern Recognition Conference (CVPR)*, 2021.
- [52] Imtiaz Ziko, Jose Dolz, Eric Granger, and Ismail Ben Ayed. Laplacian regularized few-shot learning. In *International Conference on Machine Learning (ICML)*, 2020.