

Few-shot Open-set Recognition by Transformation Consistency

Minki Jeong Seokeon Choi Changick Kim

Korea Advanced Institute of Science and Technology, Daejeon, Republic of Korea

{rhm033, seokeon, changick}@kaist.ac.kr

Abstract

In this paper, we attack a few-shot open-set recognition (FSOSR) problem, which is a combination of few-shot learning (FSL) and open-set recognition (OSR). It aims to quickly adapt a model to a given small set of labeled samples while rejecting unseen class samples. Since OSR requires rich data and FSL considers closed-set classification, existing OSR and FSL methods show poor performances in solving FSOSR problems. The previous FSOSR method follows the pseudo-unseen class sample-based methods, which collect pseudo-unseen samples from the other dataset or synthesize samples to model unseen class representations. However, this approach is heavily dependent on the composition of the pseudo samples. In this paper, we propose a novel unknown class sample detector, named *SnaTCHer*, that does not require pseudo-unseen samples. Based on the transformation consistency, our method measures the difference between the transformed prototypes and a modified prototype set. The modified set is composed by replacing a query feature and its predicted class prototype. *SnaTCHer* rejects samples with large differences to the transformed prototypes. Our method alters the unseen class distribution estimation problem to a relative feature transformation problem, independent of pseudo-unseen class samples. We investigate our *SnaTCHer* with various prototype transformation methods and observe that our method consistently improves unseen class sample detection performance without closed-set classification reduction.

1. Introduction

Recently, deep neural networks show outstanding performance in various computer vision problems. There are many reasons for this achievement, but there is no doubt that a large volume of high-quality datasets has been a great aid. However, in real-world applications, a large amount of high-quality data is not always available for various reasons, such as a high labeling cost of experts or the need to collect rare data.

Few-shot learning (FSL) methods [28, 6, 27, 35] are pro-

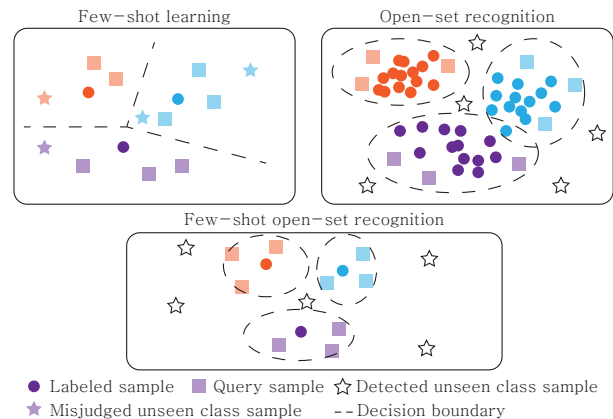


Figure 1. A visualization of the few-shot open-set recognition task. Few-shot learning methods fail to recognize unseen class samples, and open-set recognition methods require a large amount of datasets. Few-shot open-set recognition is a generalized few-shot learning task, where the model has to identify unseen class queries while classifying seen class queries correctly.

posed to reduce the data dependency. It assumes a severe condition where a few labeled data are available for training, such as one to five for each class. Various methods show remarkable improvements in FSL problems, however, these methods are limited to closed-set problems where training samples and testing samples share the same class pool. On the contrary, in real-world scenarios, there could be irregular inputs to network models that could damage the reliability of models. For instance, let us consider an automatic face tag system in images for social network services. For a new picture, the system should tag the face region of a friend or ask the user for manual tagging if the person in the image is not on the friend list. The system requires two crucial features for better user experience: correct tagging (*i.e.*, correct closed-set classification) and correct asking (*i.e.*, correct unseen class sample detection). However, the nature of FSL methods forces to tag the person as one in the friend list.

Open-set recognition (OSR) manages the unseen class detection problem in the large-scaled data situations. It aims to detect unseen class samples (*i.e.*, unseens) from




Figure 2. Influence of pseudo-unseen data configurations on FSOSR tasks. Error bars represent 95% confidence intervals. Closest samples pseudo-unseen data from the nearest classes to current prototypes, and Farthest collects from the farthest classes. The gathered pseudo-unseen samples are trained with a prediction entropy maximization loss widely used in pseudo-unseen based methods [5, 20, 16].

seen class samples (*i.e.*, seens), while maintaining the classification capability [2, 7, 21, 29, 23]. These OSR methods utilize the characteristics of seens from an informative dataset to organize an unseen class sample detector. However, the rich data information of seen classes is not guaranteed in FSL problems. These OSR methods suffer from performance degrading under the FSL condition for many reasons, such as overfitting.

In this paper, we attack the few-shot open-set recognition (FSOSR) problem. FSOSR aims to distinguish unseen class samples from seen class samples while maintaining the classification capability, using a few labeled supports. We illustrate the concept of FSOSR in Fig. 1. The previous FSOSR study [16] deals with this problem by training the feature extractor with pseudo-unseen class samples aggregated from additional non-overlapped classes. However, we observe that the pseudo-unseen based approach is heavily dependent on the quality of pseudo-unseen samples as illustrated in Fig. 2. Furthermore, these methods assume that unseen class instances are visually similar to the pseudo-unseen samples, which is not guaranteed in real-world situations. These problems are critical for FSL cases since the target task distribution is unknown during training the model with the base data.

To this end, we propose a novel unknown detector, named SnaTCHer, that does not require pseudo-unseen samples for training. Figure 3 shows the concept of SnaTCHer. Our method utilizes the transformation consistency [4], where similar samples remain close after the transformation. The transformer is trained to form a task-adaptive feature space using class representation vectors (*i.e.*, prototypes). For the unseen detection, a query firstly selects its closest transformed prototype. Then, the se-




Figure 3. A visualization of our method. Each colored box represents a prototype or a query feature. SnaTCHer replaces a predicted class prototype to the query sample in the set of prototypes (\mathcal{P}^k and \mathcal{P}^u), then it measures differences between the prototype set and the replaced set after the feature transformation $\mathbf{T}(\cdot)$. SnaTCHer rejects samples by the distance from the transformed prototype set $\mathbf{T}(\mathcal{P})$. Our method alters estimating feature distribution of unseen class samples for the detection to the relative feature transformation problem.

lected class prototype is replaced with the query feature vector. Since unseen samples tend to form a distinctive feature space from known samples, the transformed features after the replacement are more likely to be far away from the transformed prototypes than that of seen class samples by the transformation consistency. The difference after the transformation is used to identify unseen class samples. Note that SnaTCHer does not require additional pseudo-unseen samples to train the unseen sample detector. Our method shifts the training paradigm of the unseen sample detector from estimating unseen sample distributions to training a feature transformer that uses relationships between features. This approach is more straightforward than estimating the unseen sample distribution directly.

We evaluate our method with various transformation methods [36, 35, 30, 1, 3] including our normalization-based method. Our SnaTCHer significantly improves the unseen detection capability, without classification performance degradation of various few-shot learning models. Furthermore, we propose a cross-domain FSOSR benchmark that compares inter-domain robustness of FSOSR methods. Our method achieves the best performance in the cross-domain evaluation either.

To summarize, our contribution is three-fold:

- We propose a novel unseen sample detector for FSOSR, named SnaTCHer, based on the transformation consistency. Our SnaTCHer improves the unseen class sample detection capability without pseudo-unseen samples.

- We show the limitations of pseudo-unseen class sample-based methods on FSOSR tasks. These methods heavily depend on the pseudo-unseen data configuration.
- We conduct extensive experiments of our method on various benchmarks, and show our method achieves the best performance in the unseen sample detection without performance loss of the classification.

2. Related Work

2.1. Few-shot learning

Few-shot learning methods can be divided into two categories: adaptation methods and metric-based methods. There are two main goals for adaptation methods. The first is to find initial parameters which work well across novel tasks, and the other is to fastly adapt the parameters using a few supports. MAML [6] and its variants [27] divide the update process into the outer loop and the inner loop to find good initial parameters. The parameters adapt to current task in the inner loop, and the outer loop updates the parameters across tasks using the adapted parameters. Bronskill *et al.* [3] propose a new normalization technique for fast adaptation in meta-learning tasks. They propose to fuse different normalization methods to overcome the limitation of the conventional batch normalization [11] method in meta-learning tasks.

Metric-based methods aim to find a good distance function that measures the same class instances closer than different class instances. Vinyals *et al.* [32] show that neural networks can measure similarities between samples. ProtoNet [28] expands MatchingNet [32] to a problem with multiple samples per class. It introduces a prototype, which represents a class. The distance between prototypes and a query vector is used for the classification. Recent methods utilize task-specific prototypes. CTM [15] creates task-adaptive weights for the prototypes, and FEAT [35] applies a self-attention-based transformation on prototypes.

Since these FSL methods focus on good classification performance under the closed-set settings, the unseen detection capability is not guaranteed. We examine various FSL methods and show that our method records better detection capability than a naive combination of FSL and OSR methods.

2.2. Open-set recognition

Before deep neural networks rise, OSR methods rely on image feature descriptors to divide unseen class samples from seen class samples [13, 8]. OpenMax [2] employs deep neural networks to OSR by combining Extreme Value Theory with neural networks. Neal *et al.* [20] introduce the pseudo-unseen image generation method. A classifier with

an unseen category is trained with both the synthesized images and the training samples. Recently, generative model-based methods gain popularity. C2AE [21] and CGDL [29] train an autoencoder using all training samples. The reconstruction error for a given query from the autoencoder is used to detect unseen class samples. Perera *et al.* [23] propose to train a classifier with both an input image and its reconstruction. They concatenate the input image and the reconstructed image along the channel dimension, and use it to classify the image. Since unseen class samples are failed to reconstruct themselves properly, the concatenated input gives lower classification probability. These OSR methods requires to train an unseen sample detector from a training data [2, 21, 29, 23], or synthesized pseudo-unseen samples [20]. Both approaches need rich training data, which does not fit to the FSL situation.

2.3. Few-shot open-set recognition

Compared to FSL and OSR, FSOSR has been hardly explored. Recently, PEELER [16] modifies a ProtoNet-based few-shot learning model for FSOSR cases. On the top of training the distance-based classifier, it adds an open-set loss term for pseudo-unseen samples which are additionally sampled from the base data. Our method is in line with PEELER in terms of the metric-based FSOSR approach. However, our method does not require pseudo-unseen samples during training. The pseudo-unseen approach is heavily dependent to the quality of pseudo-unseens, which could be not enough to represent the true unseen class sample distribution.

FSOSR and Generalized zero-shot learning (GZSL) [17] are slightly different. GZSL is similar to FSOSR in that *target class* samples are unavailable during the training stage. However, the semantic relationships between *source classes* and *target classes* are available under the GZSL problems. On the other hand, in FSOSR, the unseen sample indicates an unobserved alien class sample. The class information is unavailable in FSOSR. Also, the goal of FSOSR is detecting out-of-distribution instances as unseen samples, while GZSL aims to classify *target class* samples to their corresponding classes.

3. Proposed Method

3.1. Problem setup

Throughout this paper, we interchangeably use the terms seen and known, and unseen and unknown. For given N classes with K labeled samples (*i.e.*, supports) per class, a task (or episode) is represented as an N -way K -shot problem. We denote the support set as $\mathcal{D}^S = \{\mathbf{x}_i^S, y_i^S\}_{i=1}^{NK}$, where $\mathbf{x}_i \in \mathcal{X}^S$ and $y_i \in \mathcal{Y}^S$ indicate an instance and its label, respectively. Different from conventional FSL, a query set in FSOSR includes unknown queries that are

not included in \mathcal{Y}^S . We denote the known query set as $\mathcal{D}^K = \{\mathbf{x}_i^K \in \mathcal{X}^K, y_i^K \in \mathcal{Y}^S\}_{i=1}^{NQ}$, where Q is the number of query samples for each class. The unknown query set is represented as $\mathcal{D}^U = \{\mathbf{x}_i^U \in \mathcal{X}^U, y_i^U \in \mathcal{Y}^U\}_{i=1}^{N^U}$, where $\mathcal{Y}^S \cap \mathcal{Y}^U = \phi$, and N^U is the number of unknowns. We employ the episodic learning approach [28] to train our model, where each mini-batch for training mimics the FSOSR task.

3.2. Revisit previous methods

At first, we briefly explain the outline of metric-based FSL methods [28, 35] and pseudo-unseen sample-based OSR methods [20, 7, 16].

A feature extractor $F(\cdot)$ creates support features and query features. With the support features, a class representation of each class (*i.e.*, prototype) is created by taking an average of class-wise features:

$$\mathbf{p}_c = \frac{1}{K} \sum_{\mathbf{x}_i^S \in \mathcal{X}_c^S} F(\mathbf{x}_i^S), \quad (1)$$

where \mathcal{X}_c^S is a set of support instances labeled to class c . The metric-based methods introduce a softmax classification using the distance between a query and the prototypes:

$$p(y = k | \mathbf{x}, \mathcal{P}) = \frac{\exp(-\text{dist}(F(\mathbf{x}), \mathbf{p}_c))}{\sum_{i \in \mathcal{Y}^S} \exp(-\text{dist}(F(\mathbf{x}), \mathbf{p}_i))}, \quad (2)$$

where $p(\cdot | \cdot, \cdot)$ is a classification probability, $\text{dist}(\cdot, \cdot)$ is a distance function, \mathbf{x} indicates the query, and $\mathcal{P} = \{\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_{N-1}\}$ indicates a set of prototypes. Usually, Euclidean distance or cosine distance is employed to measure distances. The class predictions are utilized to update the feature extractor by classification losses such as the cross-entropy loss [28].

Pseudo-unseen class-based OSR methods use pseudo-unseen instances to train an unknown sample detector. The pseudo-unseen samples are either selected from existing datasets [5, 16] or synthesized from a trained generator [7, 20]. However, training the unseen sample detector with pseudo-unseen instances assumes that the unseens are limited to the pseudo-unseen sample distributions [7, 20]. This assumption is not guaranteed in various real-world environments. Moreover, the pseudo-unseen data configuration greatly affects the model’s performance as illustrated in Fig. 2. Recently, reconstruction-based OSR methods [21, 23, 29] are proposed to exclude the pseudo-unseen dependency. They train an encoder-decoder structure, where the encoder encodes a given image to a feature vector, and the decoder reproduces the input image from the feature vector. Since the reconstruction model is fitted to the known classes and unknowns are distinct from knowns in the high-dimensional feature space, the reconstruction module fails to reproduce unseen samples. However, training the generative model requires many training data, which is unavailable in FSL. Under FSL settings,

Algorithm 1: SnaTCHer details

Input: A set of prototypes \mathcal{P} , a query feature \mathbf{q} , a distance function $\text{dist}(\cdot, \cdot)$, a transformation function $\mathbf{T}(\cdot)$, a threshold η

- 1 c : The predicted query class;
- 2 $\mathcal{P}^q = \mathcal{P} - \{\mathbf{p}_c\} + \{\mathbf{q}\}$;
- 3 $d_{\text{SnaTCHer}} = \sum_{c \in \mathcal{Y}^S} \text{dist}(\mathbf{p}_c^{q'}, \mathbf{p}'_c)$,
where $\mathbf{p}_c^{q'} \in \mathbf{T}(\mathcal{P}^q)$ and $\mathbf{p}'_c \in \mathbf{T}(\mathcal{P})$;
- 4 **if** $d_{\text{SnaTCHer}} > \eta$ **then**
- 5 | The query is unknown;
- 6 **else**
- 7 | The query is known;
- 8 **end**

there are no class overlaps between base classes and evaluation classes. Therefore, the meta-trained autoencoder with the base classes fails to reconstruct the evaluation class instances, which degrades the unseen sample detection performance.

3.3. Transformation consistency based unseen class sample detector

Our method is inspired by the concept of reconstruction-based OSR methods, which measures differences after data processing. Since the feature extractor is trained with a distance-based classifier in the metric-based FSL methods, a feature from the same class is closer than that of the different classes. Therefore, the distance between a query feature and its class prototype is closer than the distance between an unknown class feature and the prototype. On top of that, let us assume that we have a trained feature transformation function that modifies prototypes to be more distinguishable. Then, the transformer gives similar outputs to adjacent features and distinguishable outputs to distinctive inputs. Therefore, we can use the prototype-dependent transformation to score unseen class samples (*i.e.*, distinctive samples). This idea is similar to the transformation consistency regularization [18], which is widely used in various computer vision tasks. Based on the concept mentioned above, we propose a new unknown detection model, named SnaTCHer. The details are explained in Algorithm 1. With a distance-based classifier and a prototype transformer, SnaTCHer predicts the class of the input query. Then, SnaTCHer replaces the predicted class prototype with the query feature. The transformer modifies the prototype set and the replaced set, then measure the difference between them. Finally, for a certain threshold, we determine whether the query is an unseen query or not. Note that training the transformation function is free from pseudo-unseens. It aims to construct current task-specific feature space with known samples.

3.4. Task-adaptive transformation function

In the previous sub-section, we explained the details of SnaTCHer. However, an important question remains: what is the appropriate transformation function $T(\cdot)$ that changes \mathcal{P} to $\mathcal{P}' = T(\mathcal{P})$? Since we measure the difference between transformed features, the transformer should understand feature distributions. In other words, it has to gather the same class samples while scattering different class samples. Furthermore, the transformation function should be symmetric, which means that the function is independent of the prototype order (*i.e.*, a set-to-set function). Several FSL methods use transformation functions that modify prototypes to be more distinguishable for better performance. Creating more distinct prototypes meets the first condition. Moreover, they are robust to the prototype orders. This satisfies the second condition. In this subsection, we explain various prototype transformation function choices for SnaTCHer, including our normalization-based method¹.

DeepSets [36] proposes to combine symmetric operations, such as identity or max-pooling operations, to achieve the permutation robustness. We utilize a modified version of DeepSets introduced in [35] for FSL.

$$\begin{aligned} \mathbf{p}'_c &= \mathbf{p}_c + g([\mathbf{p}_c; \max_{\mathbf{p}'_c \in \mathcal{P}'_c} h(\mathbf{p}'_c)]), \\ \mathcal{P}' &= \{\mathbf{p}'_0, \mathbf{p}'_1, \dots, \mathbf{p}'_{N-1}\}, \end{aligned} \quad (3)$$

where $g(\cdot)$ and $h(\cdot)$ are non-linear MLP, $[\cdot; \cdot]$ indicates a channel-wise concatenation operation, and $\mathcal{P}'_c = \mathcal{P} - \{\mathbf{p}_c\}$ is a complementary set of $\{\mathbf{p}_c\}$.

Transformer [31, 35] projects the input feature into three different spaces (*i.e.*, query, key, and value) with learnable transformation matrices. With the projected features, the overall process is formulated as follows:

$$\begin{aligned} R(\mathbf{P}) &= \frac{1}{N} (W_V \mathbf{P}) \left(s \left((W_Q \mathbf{P})^T (W_K \mathbf{P}) / \sqrt{m} \right) \right)^T, \\ \mathbf{P}' &= \sigma(\mathbf{P} + R(\mathbf{P})), \mathcal{P}' = \{\mathbf{P}'_0, \mathbf{P}'_1, \dots, \mathbf{P}'_{N-1}\}, \end{aligned} \quad (4)$$

where \mathbf{P} indicates stacked prototypes, \mathbf{P}'_c is a modified prototype of class c , m means the feature size, W_Q , W_K , $W_V \in \mathbb{R}^{m \times m}$ are learnable transform matrices, $\sigma(\cdot)$ indicates layer normalization [1], and $s(\cdot)$ is the softmax function.

In addition to the symmetric feature transformation modules, we propose to use normalization functions directly to modify prototypes. The generalization capabilities of the normalization methods help form distinctive feature space for a novel episode. We explain normalization methods with an input feature tensor $\mathbf{t} \in \mathbb{R}^{D \times H \times W}$, where D , H , W are the channel size, height, and width of the input, respectively. The normalization functions transform the input

feature with a calculated mean $\boldsymbol{\mu}$ and a variance $\boldsymbol{\sigma}^2$:

$$\bar{\mathbf{t}} = \mathbf{a} \frac{\mathbf{t} - \boldsymbol{\mu}}{\sqrt{\boldsymbol{\sigma}^2 + \epsilon}} + \mathbf{b}, \quad (5)$$

where $\bar{\mathbf{t}}$ is the normalized feature, \mathbf{a} and \mathbf{b} are learnable transformation parameters, and ϵ is a constant for numerical stability. We describe various normalization methods to calculate the mean and variance.

Layer normalization [1] (LN) computes the normalization statistics over all elements in the input:

$$\boldsymbol{\mu}^{(L)} = \frac{1}{DHW} \sum_{t_i \in \mathbf{t}} t_i, \quad \boldsymbol{\sigma}^{2(L)} = \frac{1}{DHW} \sum_{t_i \in \mathbf{t}} (t_i - \boldsymbol{\mu}^{(L)})^2. \quad (6)$$

We apply layer normalization to individual prototypes to make the transformation symmetric. The normalized prototypes form $\mathcal{P}' = \{\bar{\mathbf{p}}_0^{(L)}, \bar{\mathbf{p}}_1^{(L)}, \dots, \bar{\mathbf{p}}_{N-1}^{(L)}\}$, where $\bar{\mathbf{p}}_c^{(L)}$ is the normalized prototype of class c .

Instance normalization [30] (IN) normalizes features across each channel:

$$\begin{aligned} \boldsymbol{\mu}_d^{(I)} &= \frac{1}{HW} \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} \mathbf{t}_{dhw}, \\ \boldsymbol{\sigma}_d^{2(I)} &= \frac{1}{HW} \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} (\mathbf{t}_{dhw} - \boldsymbol{\mu}_d^{(I)})^2, \end{aligned} \quad (7)$$

where $\mathbf{t}_{dhw} \in \mathbf{t}$ is an element of the feature. IN is frequently used in style-robust feature extraction methods such as image style transfer [10] and person re-ID [12], and proved that it acts like a style normalization [22, 10]. Inspired by previous observations, we apply IN to stacked prototypes to normalize task-specific features.

TaskNorm [3] uses a transductive and a non-transductive normalization simultaneously to compensate each other. It calculates $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$ from moments of both normalization methods:

$$\begin{aligned} \boldsymbol{\mu}^{(T)} &= \alpha \boldsymbol{\mu}_{BN} + (1 - \alpha) \boldsymbol{\mu}_+, \\ \boldsymbol{\sigma}^{2(T)} &= \alpha (\boldsymbol{\sigma}_{BN}^2 + (\boldsymbol{\mu}_{BN} - \boldsymbol{\mu}^{(T)})^2) \\ &\quad + (1 - \alpha) (\boldsymbol{\sigma}_+^2 + (\boldsymbol{\mu}_+ - \boldsymbol{\mu}^{(T)})^2), \end{aligned} \quad (8)$$

where $\boldsymbol{\mu}_{BN}$ and $\boldsymbol{\sigma}_{BN}^2$ are calculated from batch normalization [11], $\boldsymbol{\mu}_+$ and $\boldsymbol{\sigma}_+^2$ are from a non-transductive normalization such as LN or IN, and $\alpha \in [0, 1]$ is a meta-learned weight parameter. We use LN to calculate $\boldsymbol{\mu}_+$ and $\boldsymbol{\sigma}_+^2$ since prototypes are feature vectors.

Layer-Task Normalization (LTN). Inspired by TaskNorm and style-normalized feature extraction methods [12], we design a feature transformation function for SnaTCHer. It combines IN and LN with a task-dependent weight, formulated as follows:

$$\begin{aligned} \mathbf{p}'_c &= \gamma (\alpha(\mathbf{P}) \cdot \bar{\mathbf{p}}_c^{(L)} + (1 - \alpha(\mathbf{P})) \cdot \bar{\mathbf{p}}_c^{(I)}) + \boldsymbol{\beta}, \\ \mathcal{P}' &= \{\mathbf{p}'_0, \mathbf{p}'_1, \dots, \mathbf{p}'_{N-1}\}, \end{aligned} \quad (9)$$

¹Here we share the transformation function symbol $T(\cdot)$ and the set of transformed prototypes symbol \mathcal{P}' across all methods for convenience.

where $\bar{p}_c^{(L)}$ and $\bar{p}_c^{(I)}$ indicate normalized prototypes of class c with LN and IN, respectively. Here we do not apply the affine transformations to the normalized prototypes. $\alpha(\cdot) \in [0, 1]$ is a weight generator, $\gamma \in \mathbb{R}^D$ and $\beta \in \mathbb{R}^D$ are learnable transformation parameters. IN normalizes task-related information similar to the style normalization. However, removing task-specific knowledge may lose class-specific information. Therefore, we introduce the weighted summation with LN to compensate for the task-normalization. The balance weight is generated from a meta-learned weight generator. It is a symmetric function that consumes prototypes, where the transformation is independent of the stack order of prototypes. A concept to combine multiple normalization methods to compensate each other is similar to BIN [19] and TaskNorm [3]. However, our inspiration to fuse task-normalized and instance-normalized features for few-shot settings is different from previous works. Also, the task-adaptive meta-weight generation approach is different.

3.5. Training loss

With a transformation function $T(\cdot)$, the prediction probability in Eq. 2 is redefined as follows:

$$p(y = k | \mathbf{x}, \mathcal{P}') = \frac{\exp(-\text{dist}(F(\mathbf{x}), \mathbf{p}'_c))}{\sum_{i \in \mathcal{Y}^S} \exp(-\text{dist}(F(\mathbf{x}), \mathbf{p}'_i))}, \quad (10)$$

where $\mathbf{p}'_c \in \mathcal{P}'$ indicates the transformed prototype. We utilized Euclidean distance with a temperature value of 64 for the distance function, and employed the cross-entropy loss function with Eq. 10 to train networks. Moreover, the transformation should cluster each class after the transformation to satisfy our assumption for the unseen class instance rejection. Therefore, we add a regularization term \mathcal{R} that gathers class-wise instances after transformation.

$$\mathcal{R} = CE(p(\mathbf{x}', \mathcal{C}), y)_{(\mathbf{x}, y) \in \mathcal{D}^K \cup \mathcal{D}^S} \quad (11)$$

where $CE(\cdot, \cdot)$ indicates the cross-entropy function, $\mathbf{x}' \in T(\mathcal{X}^S \cup \mathcal{X}^K)$ is a transformed feature, and \mathcal{C} is a set of class centers after the transformation. With the regularization term, the objective loss function is defined as follows:

$$\mathcal{L}_{total} = CE(p(\mathbf{x}, \mathcal{P}'), y)_{(\mathbf{x}, y) \in \mathcal{D}^K} + \lambda \mathcal{R}, \quad (12)$$

where λ is a weight hyperparameter.

4. Experiments

4.1. Datasets and evaluation methods

We use miniImageNet [32] and tieredImageNet [25], which are widely used for few-shot learning evaluations. Both datasets are the subsets of the ImageNet [26] dataset. There are 100 classes of 84×84 natural RGB images in the

miniImageNet dataset. For each class, 600 images exist. We splitted the 100 classes following the commonly used dataset split for few-shot learning [24]. We set 64 classes for base classes, 16 classes for validation, and 20 classes for evaluation. The tieredImageNet dataset is composed of 608 classes of 84×84 natural RGB images. These images are divided into 351, 97, and 160 classes for training, validation, and testing, respectively.

Following the previous FSOSR study [16], we use the closed-set classification accuracy (Acc) and the AUROC of unknown class sample detection (AUROC) for evaluation. The accuracy measures the correct classification ratio using the seen class samples, and the AUROC measures unseen class instance detection capability using both seen and unseen class samples. We set five classes as known classes and the other non-overlapped five classes as unknown classes to compose a single 5-way episode during the experiments. We collected 15 instances for each class as queries, which leads to 75 known queries and 75 unknown queries for a 5-way episode.

4.2. Implementation details

We utilized the ResNet-12 [9] based architecture for the feature extraction, following previous FSL methods [14, 35]. The feature extractor creates a 640-dimensional feature vector for an input image through the last average pooling layer. We pre-trained the extractor with a simple classifier to classify all classes in the base dataset. The classifier is a single fully-connected layer that creates logits of the classes. The feature extractor architecture is shared across all transformation methods. The weight generator in LTN consumes the stacked prototypes and results in a scalar weight. It is composed of four one-dimensional convolutional layers with an one-dimensional average pooling layer. The pooling layer is placed after the second convolution to make the generator robust to the prototype order. The convolutional layers in the generator use the LeakyReLU [34] function as non-linearity except the final layer, which uses the sigmoid function to limit the weight range from zero to one.

We used stochastic gradient descent to train networks over 20,000 episodes. The initial learning rate is set to 0.002 for transformers and 0.0002 for the feature extractor with learning rate decaying. The loss weight λ is set to 0.1. During training, we selected the best model with the validation data to compare with the other methods on the test data. The comparison results are calculated over 600 evaluation episodes. Please see the supplementary material for further details.

4.3. Transformer comparison

We evaluate the unseen class sample detection performance of the transformation methods on 5-way episodes of the tieredImageNet dataset, and present the results in

Table 1. For each transformation method, we compare SnaTCHer with two other unseen class sample detection baselines. The first method utilizes the classification probability to detect unseen samples, which is widely used in OSR methods [2, 5, 20, 16]. The other method uses the distance between a query and its predicted class prototype. This method is based on our assumption on the feature distributions of knowns and unknowns. We denote the former as Probability, and the latter as Distance in Table 1. The result shows that the feature distance-based detection methods show better performance than Probability. Since the softmax function utilizes relative differences of logit vector elements, it sacrifices the absolute feature distance information which is important in the FSOSR problem. Compared with Distance, SnaTCHer shows better performance since it considers the relationships between a query and all prototypes, thanks to the transformation. Based on the comparison result, we select FEAT, TaskNorm, and LTN to compare with the other methods. We name them SnaTCHer-F, SnaTCHer-T, and SnaTCHer-L, respectively.

4.4. Comparison with the other methods

We compare our approach with the state-of-the-art FSL methods (ProtoNet [28], FEAT [35]), the OSR methods (NN [13], OpenMax [2]), and the FSOSR method [16]. Since the FSL and OSR methods are not designed for FSOSR cases, we slightly changed them for the comparison. For all FSL methods, we determined the unknown detection score with the negative of the predicted class probability. For OpenMax, we fitted the Weibull models over the training episodes and used them to detect unknowns in the evaluation stage. For NN, we used a ProtoNet-based classifier to calculate logits. Please see the supplementary material for the details.

Table 2 shows the comparison result. OSR methods have poor closed-set accuracy, and FSL methods show low unseen detection capabilities than the FSOSR methods. This result shows the limitations of FSL methods and OSR methods on FSOSR. FSL methods are fitted to the closed-set classification tasks, which leads to poor unknown instance detection capability. On the contrary, OSR methods require large-scale datasets, which degrades its performance under the few-shot condition. Compared with the previous FSOSR method, our method accomplishes both higher classification accuracy and better unknown detection AUROC simultaneously. Our SnaTCHer accomplishes unseen detection by training the task-adaptive feature transformer without pseudo-unseen samples. This paradigm shift enables improving both the classification accuracy and unseen sample detection capability.

Among the three transformation methods, our SnaTCHer-L achieves high classification and detection capability simultaneously. The adaptive fusion of the

Model	1-shot	5-shot
Identity		
Probability	60.73 ± 0.80	64.96 ± 0.83
Distance	72.40 ± 0.76	79.22 ± 0.63
SnaTCHer	72.40 ± 0.76	79.22 ± 0.63
DeepSets [36]		
Probability	62.66 ± 0.75	72.46 ± 0.70
Distance	65.75 ± 0.72	75.24 ± 0.61
SnaTCHer	68.21 ± 0.78	75.63 ± 0.60
FEAT [35]		
Probability	63.60 ± 0.75	74.27 ± 0.68
Distance	69.36 ± 0.78	77.63 ± 0.65
SnaTCHer	74.28 ± 0.80	82.02 ± 0.64
LN [1]		
Probability	64.09 ± 0.76	69.29 ± 0.76
Distance	68.92 ± 0.82	68.92 ± 0.82
SnaTCHer	74.07 ± 0.82	79.74 ± 0.72
IN [30]		
Probability	60.94 ± 0.78	66.01 ± 0.73
Distance	61.89 ± 0.74	67.93 ± 0.71
SnaTCHer	71.22 ± 0.76	79.39 ± 0.61
TaskNorm [3]		
Probability	65.31 ± 0.72	67.16 ± 0.77
Distance	69.72 ± 0.76	74.91 ± 0.68
SnaTCHer	74.84 ± 0.79	82.03 ± 0.66
LTN		
Probability	64.94 ± 0.74	69.66 ± 0.74
Distance	69.80 ± 0.77	76.22 ± 0.70
SnaTCHer	74.95 ± 0.83	80.81 ± 0.68

Table 1. The AUROC (%) comparison of unknown detection methods on various metric-based few-shot learning methods. Identity indicates the identity transformation.

task-normalization and the layer-normalization effectively moves the prototypes to be more task-specific, which brings the performance improvements.

4.5. Further analysis

We illustrate histograms of the normalized distances between a query and its corresponding prototype in Fig. 4(a). As we expected, the known samples are closer to its prototypes than the unknowns. This observation supports our assumption that unknowns form distinguishable feature space. Figure 4(b) illustrates the histograms of the normalized unseen scores. As evident from the figure, SnaTCHer has a better separation than using the distances directly, because the difference is amplified via the relationship-based transformation. This result validates the larger AUROC values of SnaTCHer.

During the experiments, we fixed the number of unknown classes to five. However, the number of unknown samples varies in real-world applications. Therefore, we investigate the performance differences over various unseen data settings. As illustrated in Fig. 5, SnaTCHer is robust to the unknown data configuration for both cases.

Model	miniImageNet 5-way				tieredImageNet 5-way			
	1-shot		5-shot		1-shot		5-shot	
	Acc	AUROC	Acc	AUROC	Acc	AUROC	Acc	AUROC
ProtoNet [28]	64.01 ± 0.88	51.81 ± 0.93	80.09	60.39	68.26	60.73	83.40	64.96
FEAT [35]	67.02 ± 0.85	57.01 ± 0.84	82.02	63.18	70.52	63.54	84.74	70.74
NN [13]	63.82 ± 0.85	56.96 ± 0.75	80.12	63.43	67.73	62.70	83.43	69.77
OpenMax [2]	63.69 ± 0.84	62.64 ± 0.80	80.56	62.27	68.28	60.13	83.48	65.51
PEELER* [16]	58.31 ± 0.58	61.66 ± 0.62	75.08	69.85	—	—	—	—
PEELER [16]	65.86 ± 0.85	60.57 ± 0.83	80.61	67.35	69.51	65.20	84.10	73.27
SnaTCHer-F	67.02 ± 0.85	68.27 ± 0.96	82.02	77.42	70.52	74.28	84.74	82.02
SnaTCHer-T	66.60 ± 0.80	70.17 ± 0.88	81.77	76.66	70.45	74.84	84.42	82.03
SnaTCHer-L	67.60 ± 0.83	69.40 ± 0.92	82.36	76.15	70.85	74.95	85.23	80.81

Table 2. Average closed-set classification accuracies (%) and average unknown detection AUROCs (%) over 600 episodes. Please see the supplementary material for the confidence intervals. PEELER* is quoted from the paper, which has a ResNet-18 backbone.

Model	tieredImageNet-CUB				CUB-tieredImageNet				CUB-CUB			
	5-way 1-shot		5-way 5-shot		5-way 1-shot		5-way 5-shot		5-way 1-shot		5-way 5-shot	
	Acc	AUROC	Acc	AUROC	Acc	AUROC	Acc	AUROC	Acc	AUROC	Acc	AUROC
PEELER [16]	69.51	67.59	84.10	76.10	58.81	57.58	77.66	64.38	59.42	58.63	78.42	66.04
SnaTCHer-F	70.52	83.22	84.74	90.12	57.81	63.87	77.33	69.64	57.98	64.55	77.05	71.05
SnaTCHer-T	70.45	84.95	84.42	91.83	57.84	64.52	77.77	70.63	57.82	65.10	77.66	72.04
SnaTCHer-L	70.85	83.67	85.23	90.23	60.21	63.55	79.27	69.42	59.69	64.75	78.68	70.67

Table 3. Cross-domain FSOSR comparison results. Please see the supplementary material for the confidence intervals.




Figure 4. Normalized score histograms for knowns and unknowns over 600 5-way 5-shot tasks on tieredImageNet. We use SnaTCHer-L for the comparison. The scores are normalized by the maximum value of each episode.




Figure 5. Ablation study results on 5-way tieredImageNet episodes. We use SnaTCHer-L for the ablation study.

4.6. Cross-domain FSOSR

Until now, we compared various methods on the same dataset with different splits. We extend the comparison to more general situations where the domain of the base data, seen data and unseen data of the episode are different. We call this problem a cross-domain FSOSR problem. For the unseen data domain, we use the CUB [33] dataset, which is composed of 200 classes of RGB bird images. The last 100 classes are used as a cross-domain dataset. We pre-

pare three evaluation scenarios for cross-domain FSOSR. The first case samples knowns from tieredImageNet, and collect unknown instances from the CUB dataset. We denote it as a tieredImageNet-CUB case. Similarly, we define a CUB-tieredImageNet case and a CUB-CUB case to assess the generalization capabilities of the FSOSR methods. For the latter two cases, we use the models trained with 5-way 5-shot tieredImageNet episodes for evaluations. Table 3 shows the result. When the known and unknown domains are different, it is easier to detect unknowns. Therefore AUROC values are large, even over 90% in the tieredImageNet-CUB 5-shot case. Overall, our methods show higher AUROC values with comparable or better classification performance than PEELER for all cases. This result indicates that our transformation consistency approach is universally applicable to various FSL problems.

5. Conclusion

In this paper, we have proposed a novel unseen class sample detection method, named SnaTCHer, to solve few-shot open-set recognition problems. SnaTCHer is model-free, and it is straightforward to apply existing metric-based few-shot learning methods. Moreover, our method does not require additional pseudo-unknown samples. Our extensive analysis with various transformation functions validates the effect of SnaTCHer. Furthermore, comparisons with the pseudo-unseen dependent method validates our approach to prevent pseudo-unseen samples is effective.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [2] Abhijit Bendale and Terrance E Boult. Towards open set deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1563–1572, 2016.
- [3] John Bronskill, Jonathan Gordon, James Requeima, Sebastian Nowozin, and Richard E Turner. Tasknorm: Rethinking batch normalization for meta-learning. *arXiv preprint arXiv:2003.03284*, 2020.
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [5] Akshay Raj Dhamija, Manuel Günther, and Terrance Boult. Reducing network agnostophobia. In *Advances in Neural Information Processing Systems*, pages 9157–9168, 2018.
- [6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- [7] Zongyuan Ge, Sergey Demyanov, Zetao Chen, and Rahil Garnavi. Generative openmax for multi-class open set classification. In *British Machine Vision Conference 2017*. British Machine Vision Association and Society for Pattern Recognition, 2017.
- [8] Chuanxing Geng, Sheng-jun Huang, and Songcan Chen. Recent advances in open set recognition: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [10] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017.
- [11] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [12] Xin Jin, Cuiling Lan, Wenjun Zeng, Zhibo Chen, and Li Zhang. Style normalization and restitution for generalizable person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3143–3152, 2020.
- [13] Pedro R Mendes Júnior, Roberto M De Souza, Rafael de O Werneck, Bernardo V Stein, Daniel V Pazinato, Waldir R de Almeida, Otávio AB Penatti, Ricardo da S Torres, and Anderson Rocha. Nearest neighbors distance ratio open-set classifier. *Machine Learning*, 106(3):359–386, 2017.
- [14] Kwonjoon Lee, Subhransu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10657–10665, 2019.
- [15] Hongyang Li, David Eigen, Samuel Dodge, Matthew Zeiler, and Xiaogang Wang. Finding task-relevant features for few-shot learning by category traversal. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–10, 2019.
- [16] Bo Liu, Hao Kang, Haoxiang Li, Gang Hua, and Nuno Vasconcelos. Few-shot open-set recognition using meta-learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8798–8807, 2020.
- [17] Shichen Liu, Mingsheng Long, Jianmin Wang, and Michael I Jordan. Generalized zero-shot learning with deep calibration network. In *Advances in Neural Information Processing Systems*, pages 2005–2015, 2018.
- [18] Aamir Mustafa and Rafal K Mantiuk. Transformation consistency regularization—a semi-supervised paradigm for image-to-image translation. *arXiv preprint arXiv:2007.07867*, 2020.
- [19] Hyeonseob Nam and Hyo-Eun Kim. Batch-instance normalization for adaptively style-invariant neural networks. In *Advances in Neural Information Processing Systems*, pages 2558–2567, 2018.
- [20] Lawrence Neal, Matthew Olson, Xiaoli Fern, Weng-Keen Wong, and Fuxin Li. Open set learning with counterfactual images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 613–628, 2018.
- [21] Poojan Oza and Vishal M Patel. C2ae: Class conditioned auto-encoder for open-set recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2307–2316, 2019.
- [22] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 464–479, 2018.
- [23] Pramuditha Perera, Vlad I Morariu, Rajiv Jain, Varun Manjunatha, Curtis Wigington, Vicente Ordonez, and Vishal M Patel. Generative-discriminative feature representations for open-set recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11814–11823, 2020.
- [24] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017.
- [25] Mengye Ren, Sachin Ravi, Eleni Triantafillou, Jake Snell, Kevin Swersky, Josh B. Tenenbaum, Hugo Larochelle, and Richard S. Zemel. Meta-learning for semi-supervised few-shot classification. In *International Conference on Learning Representations*, 2018.
- [26] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [27] Andrei A. Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization.

- In *International Conference on Learning Representations*, 2019.
- [28] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087, 2017.
 - [29] Xin Sun, Zhenning Yang, Chi Zhang, Keck-Voon Ling, and Guohao Peng. Conditional gaussian distribution learning for open set recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13480–13489, 2020.
 - [30] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
 - [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
 - [32] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638, 2016.
 - [33] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
 - [34] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015.
 - [35] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8808–8817, 2020.
 - [36] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. In *Advances in neural information processing systems*, pages 3391–3401, 2017.