

Learning Open Set Network with Discriminative Reciprocal Points

Guangyao Chen¹, Limeng Qiao¹, Yemin Shi¹, Peixi Peng^{1*}, Jia Li^{2,3}, Tiejun Huang^{1,3}, Shiliang Pu⁴, and Yonghong Tian^{1,3 *}

¹ Department of Computer Science and Technology, Peking University

² State Key Laboratory of Virtual Reality Technology and Systems, SCSE, Beihang University

³ Peng Cheng Laboratory, Shenzhen, China

⁴ Hikvision Research Institute, Hangzhou, China

{gy.chen, qiaolm, pxpeng, yhtian}@pku.edu.cn

Abstract. Open set recognition is an emerging research area that aims to simultaneously classify samples from predefined classes and identify the rest as 'unknown'. In this process, one of the key challenges is to reduce the risk of generalizing the inherent characteristics of numerous unknown samples learned from a small amount of known data. In this paper, we propose a new concept, *Reciprocal Point*, which is the potential representation of the extra-class space corresponding to each known category. The sample can be classified to known or unknown by the otherness with reciprocal points. To tackle the open set problem, we offer a novel open space risk regularization term. Based on the bounded space constructed by reciprocal points, the risk of unknown is reduced through multi-category interaction. The novel learning framework called **Reciprocal Point Learning (RPL)**, which can indirectly introduce the unknown information into the learner with only known classes, so as to learn more compact and discriminative representations. Moreover, we further construct a new large-scale challenging aircraft dataset for open set recognition: **Aircraft 300 (Air-300)**. Extensive experiments on multiple benchmark datasets indicate that our framework is significantly superior to other existing approaches and achieves state-of-the-art performance on standard open set benchmarks.

1 Introduction

Recent error rate for classification task by Deep Convolutional Neural Network has surpassed the human-level performance [9]. Such a significant improvement promotes its attempts in real-world applications. However, a robust recognition system should not only be able to identify all test instances of the seen or known classes, but also be able to handle unknown samples or novel events that have not been seen. Recent works on open set recognition [23] have formalized the process for performing recognition in the setting that labels the input as one of the known

*Corresponding author

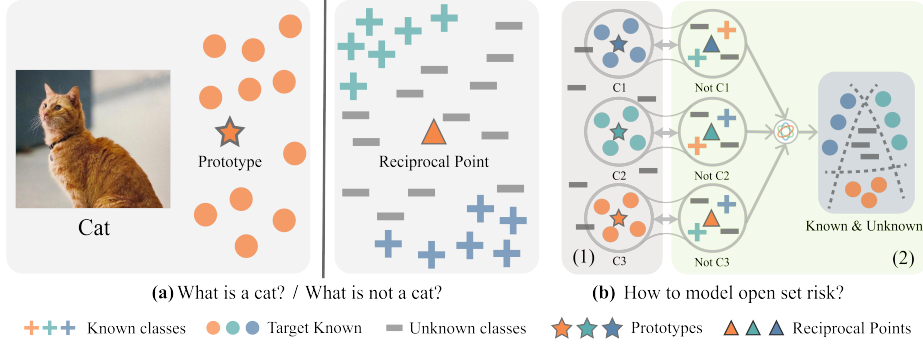


Fig. 1. (a): How to identify a cat in open set problem? Most methods focus on learning the representative potential feature of cats as prototypes. In contrast, *Reciprocal Points*, representative potential feature of *non-cat*, identify the cat by otherness. (b): (1) Focusing only on known information, it is hard to reduce open space risk without unknown information. (2) The *extra-class* space represented by *Reciprocal Points* can bring more connections with unknown.

classes or as unknown. Apart from returning the most likely category, the robust recognition system must also support the rejection of unseen/unknown. An open set scenario has classes in testing phase that are not seen during training, which brings great difficulties to solve open set recognition.

In the majority of deep neural networks, the last layer for multiclass recognition is the SoftMax function which produces a probability distribution over the known classes. However, the SoftMax layer brings a significant difficulty to open set recognition because of its closed nature. Bendale *et al* [2] improves SoftMax to OpenMax to estimate the probability of unknown classes but does not optimize for unknown classes during the training phase. Potential solutions to the open set recognition should optimize for unknown classes, as well as the known classes [23]. More knowledge of unknown information can facilitate reducing of open space risk. Many methods [8,19,21,26] utilize generative models for generating samples as unknown. However, Nalisnick *et al* [18] finds the instability of the mainstream generation models to the open set problem. How to model open space risk with only known training data is still an urgent problem to be solved. For example, as shown in Fig.1(a), how to identify a cat in open set problem. Most classification methods pay attention to "what is a cat?". They look for a more representative feature of the cat. Yang *et al* [25] proposes convolutional prototype learning (CPL) to solve the multiclass recognition problem in the open world. For neural networks, each known class k exists in its own way. Although focusing only on known samples can reduce the intra-class distance, it is still inevitable that it will bring open space risks because the dark information in the unknown space cannot be considered at all. In addition, for each known category k , most unknown samples obviously belong to the space of non- k and their features should be more similar to the representation of non- k correspondingly

for neural networks, which means that the corresponding unknown information is more implicit in each non- k embedding space. Therefore, we further focus on "what is not a cat?" to learn representations in the form of one more latent vector of non-cat data. The new representation is called *Reciprocal Point*, which can bring more connections with unknown by exploring each extra-class space.

Based on the exploration of the extra-class space, we propose a novel framework, termed as **Reciprocal Point Learning (RPL)**, which can effectively reduce the open space risk by bounded space while learning reciprocal points of each known class, and the whole framework of the proposed method is shown in Fig. 2. Specifically, we extract the feature of each sample with a deep Convolutional Neural Network model. Since the reciprocal point is the learnable representation of the extra-class space, we classify the input by the otherness between the embedding feature and the reciprocal point. Then we extend the model to capture the risk of the unknown based on reciprocal points by a novel open space risk regularization term. The extra-class embedding space of each known class is limited to a bounded range through the reciprocal point and the learnable margin. When multiple classes interact with each other in the training stage, all known classes are not only pushed to the periphery of the space by the corresponding reciprocal point but also pulled in a certain bounded range by other reciprocal points. Finally, the embedding space of the network is limited to a bounded range. All known classes are distributed around the periphery of the embedding space, and the unknown samples are limited to the internal bounded space. The bounded domain by RPL can prevent the neural network from generating arbitrarily high confidence for unknown [10]. Although only known samples are available during the training stage, the interval between known and unknown classes is separated by reciprocal points indirectly. The advantage of RPL is that it can shrink unknown space while considering the known space classification and form a good bounded feature space distribution in which the magnitudes of unknown samples are lower than those of known samples.

Moreover, in order to facilitate the research of open set recognition in the real visual world, we further construct a new large-scale aircraft dataset from the Internet: **Aircraft 300 (Air-300)**, which contains 320,000 annotated colored images from 300 different classes in total. Note that each category contains 100 images at least, and a maximum of 10,000 images, which leads to the long tail distribution of our proposed dataset. Compared with the existing benchmark datasets, the tailored Air-300 dataset maintains a long tail distribution to simulate the real visual world, and can also be utilized to perform fine-grained classification and object recognition.

To summarize our contributions: (1) we propose a new concept, *Reciprocal Point*, which is the potential representation of the extra-class space corresponding to each known category. (2) Meanwhile, we offer a novel framework for learning open set network. Through introducing unknown information by reciprocal points, the neural network can learn more compact and robust feature space, and separate known space and unknown space effectively. (3) Then, we further introduce a new challenging large-scale dataset with long-tailed distribution, Air-

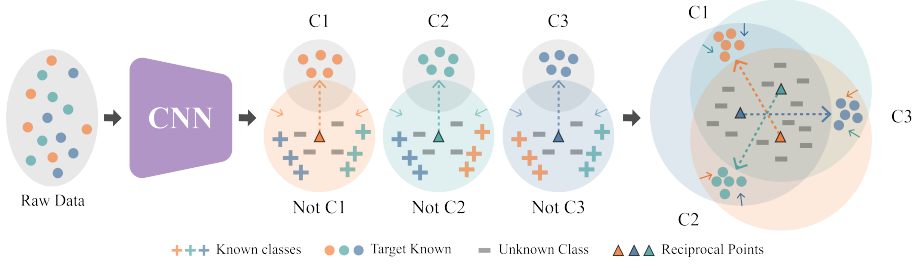


Fig. 2. An overview of the proposed Reciprocal Point Learning (RPL) approach to open set recognition. The extra-class embedding space of each known class is limited to a bounded range through the reciprocal point and the learnable margin. With interaction of multiple classes, the known classes are pushed to the periphery of the feature space, and the unknown classes are limited in the bounded space.

300, which is a new dataset to tailor for open set recognition. (4) The results on multiple benchmark datasets show that our approach significantly outperforms other existing state-of-the-art deep open set classifiers.

2 Related Work

Open Set Recognition Inspired by a classifier with rejection option [1,3,27], Scheirer *et al.* [23] define the open set recognition problem for the first time and propose a base framework to perform training and evaluation. The deep neural network has achieved great success in many areas and is introduced to open set recognition by Bendale *et al.* [2]. They prove threshold on SoftMax probability does not yield a robust model for open set recognition. Openmax [2] detects unknown classes by modeling distance of activation vectors, which are from the mean penultimate vector of each class with EVT. Ge *et al.* [8] propose G-Openmax, a direct extension of Openmax, which trains deep neural networks with unknown samples generated by generative models. However, this method does not extend to tasks that generate modeling difficulties. Yoshihashi *et al.* [26] addresses deep reconstruction-based representation learning in open set recognition by adding image reconstructed as a regularizer. However, this method introduces auxiliary decoder network for training, which brings additional overhead to training.

Out-of-Distribution Detection Based on the concern for safety of AI systems, the detection of out-of-distribution (OOD) examples, one of the concerns regarding the distribution shift, is first introduced by Hendrycks *et al.* in [11]. OOD is the detection of samples that do not belong to the training set but can appear during testing [11]. Several works [11,16,15,12] seek to address these problems by giving deep neural network classifiers a means of assigning anomaly scores to inputs, which are used for detecting OOD. Hendrycks *et al.* [11] demonstrated that anomalous samples had a lower maximum softmax probability than

in-distribution samples based on a deep, pre-trained classifier. Liang *et al.* [16] propose ODIN to allow more effective detection by using temperature scaling and adding small perturbations to the input. Similar with G-Openmax, De Vries *et al.* [12] utilize generative models for generating most effective samples from out-of-distribution and deriving a new OOD score from this branch. Hendrycks *et al.* [12] propose Outlier Exposure, which uses an auxiliary dataset in order to teach the network better representations for anomaly detection. Detecting out-of-distribution (OOD) is related to the rejection of unknown classes in open set recognition (OSR). They all are studying separating in-distribution (known) and out-of-distribution (unknown) [11,23].

Prototype Learning Prototypes are learnable representations in the form of one or more latent vector per class. Wen *et al.* [24] proposed a center loss to learn centers for deep features of each identity and used the centers to reduce intra-class variance. It can enhance the discrimination power of features. Yang *et al.* in [25] propose Convolutional Prototype Learning (CPL) for robust classification. Moreover, a Generalized CPL (GCPL) with prototype loss was proposed as a regularization to improve the intra-class compactness of the feature representation. However, center loss and GCPL only explicitly encourage intra-class compactness. Conversely, the reciprocal point limits the embedding space outside each known class to bounded space. Known and unknown classes are separated by maximizing otherness between known and corresponding reciprocal points.

3 Reciprocal Point Learning

3.1 Preliminaries

We first establish preliminaries related to open set learning, following which we formally formulate the proposed idea. Suppose we are given a set of n labeled samples $\mathcal{D}_{\mathcal{L}} = \{(x_1, y_1), \dots, (x_n, y_n)\}$ where $y_i \in \{1, \dots, N\}$ is the label of x_i and a larger amount of test data $\mathcal{D}_{\mathcal{T}} = \{t_1, \dots, t_u\}$ where the label of t_i belong to $\{1, \dots, N\} \cup \{N+1, \dots, N+U\}$ and U is the number of unknown classes in actual scenarios. We then denote the embedding space of category k as \mathcal{S}_k and its corresponding **open space** as $\mathcal{O}_k = \mathbb{R}^d - \mathcal{S}_k$, where \mathbb{R}^d is the d -dimensional full space. In order to formalize and then manage open space risk finely, we further address the positive open space from other known classes as \mathcal{O}_k^{pos} and the remaining infinite unknown space as negative open space \mathcal{O}_k^{neg} , namely $\mathcal{O}_k = \mathcal{O}_k^{pos} \cup \mathcal{O}_k^{neg}$.

The Open Set Recognition Problem. For simplicity, we will first introduce the open set recognition of a single class and then extend the entire learning process to a multiclass form. Given the labeled data $\mathcal{D}_{\mathcal{L}}$ with N known classes, let samples from category k be positive training data $\mathcal{D}_{\mathcal{L}}^k$ (space \mathcal{S}_k), samples from other known classes be negative training data $\mathcal{D}_{\mathcal{L}}^{\neq k}$ (space \mathcal{O}_k^{pos}), and samples from \mathbb{R}^d but except $\mathcal{D}_{\mathcal{L}}$ be potential unknown data $\mathcal{D}_{\mathcal{U}}$ (space \mathcal{O}_k^{neg}). Let

$\psi^k : \mathbb{R}^d \mapsto \{0, 1\}$ be a binary measurable prediction function, mapping embedding x to the label y . For 1-class open set recognition problem, the overall goal is to optimize a discriminant binary function ψ_k by minimizing the expected error \mathcal{R}^k as:

$$\arg \min_{\psi_k} \{ \mathcal{R}^k \mid \mathcal{R}_\epsilon(\psi_k, \mathcal{S}_k \cup \mathcal{O}_k^{pos}) + \alpha \cdot \mathcal{R}_o(\psi_k, \mathcal{O}_k^{neg}) \} \quad (1)$$

where α is a positive regularization parameter, \mathcal{R}_ϵ is the empirical classification risk on known data and \mathcal{R}_o is the **open space risk** function to measure the uncertainty of labeling the unknown samples as the class of interest or as unknown, with further formulating as a non-zero integral function on space \mathcal{O}_k^{neg} .

Afterwards we identify open set recognition problem by integrating multiple binary classification tasks (*one vs. rest*) into a multiclass recognition problem (see Fig. 2), by summarizing the expected risk defined in the Eq.(1) category by category, i.e. $\sum_{k=1}^N \mathcal{R}^k$, which leads to the following formulation as:

$$\arg \min_{f \in \mathcal{H}} \{ \mathcal{R}_\epsilon(f, \mathcal{D}_L) + \alpha \cdot \sum_{k=1}^N \mathcal{R}_o(f, \mathcal{D}_U) \} \quad (2)$$

where $f : \mathbb{R}^d \mapsto \mathbb{N}$ is a measurable multiclass recognition function (see more details of derivation in supplementary material). According to the Eq. (2), solving the open set recognition problem is equivalent to minimize the combination of the empirical classification risk on labeled known data and the open space risk on potential unknown data simultaneously, over the space of allowable recognition functions, which leads to be more distinguishable between known and unknown spaces. Further, we propose reciprocal points for both closed set classification and reducing the open space risk on \mathcal{D}_U during training.

3.2 Reciprocal Points for Classification

The **reciprocal points** for category k are denoted as $\mathcal{P}^k = \{p_i^k \mid i = 1, \dots, M\}$, where M is the number of reciprocal points for each class and \mathcal{P}^k can be regarded as the set of latent representations for the sub-dataset $\mathcal{D}_L^{\neq k} \cup \mathcal{D}_U$ (space \mathcal{O}_k). Unlike the learning driven by prototype or center loss, we agree the fact that most samples from the classes in \mathcal{O}_k are more similar to the classificatory reciprocal points \mathcal{P}^k than samples from \mathcal{S}_k , which can be formulated as:

$$\forall d \in \zeta(\mathcal{P}^k, \mathcal{D}_L^k), \max(\zeta(\mathcal{P}^k, \mathcal{D}_L^{\neq k} \cup \mathcal{D}_U)) \leq d \quad (3)$$

where $\zeta(\cdot, \cdot)$ calculates a set of distances of all samples between two sets and $\max(\cdot)$ is a maximal function.

In fact, the most intuitive goal of open set learning is usually to separate known and unknown spaces as much as possible. To achieve this, for any category k , we propose to separate the two mutually exclusive space by maximizing the distance between the reciprocal points of the category and its corresponding known samples. Specifically, we optimize a set of d -dimensional representations \mathcal{P}^k , or reciprocal points, of each class through an embedding function f_θ with

learnable parameters θ . Given the sample x and classificatory reciprocal points \mathcal{P}_k , we calculate the distance between them with the formulation as:

$$d(f_\theta(x), \mathcal{P}^k) = \frac{1}{M} \sum_{i=1}^M \|f_\theta(x) - p_i^k\|_2^2 \quad (4)$$

Afterwards, based on the distance we propose above, our framework estimates the otherness between the embedding feature $f_\theta(x)$ and the reciprocal points \mathcal{P}^k of all known classes to determine which category it belongs to. Actually, following the nature of reciprocal points, the probability of sample x belongs to category k is proportional to the otherness between $f_\theta(x)$ and the reciprocal points of category k , i.e. $p(y = k|x) \propto d(f_\theta(x), \mathcal{P}^k)$, which means the greater the distance between $f_\theta(x)$ and \mathcal{P}^k , the greater the probability that x will be assigned with label k . According to the sum-to-one property, we embody the final probability as:

$$p(y = k|x, f_\theta, \mathcal{P}) = \frac{e^{\gamma d(f_\theta(x), \mathcal{P}^k)}}{\sum_{i=1}^N e^{\gamma d(f_\theta(x), \mathcal{P}^i)}} \quad (5)$$

where γ is a hyper-parameter that controls the hardness of probability assignment. Learning proceeds by minimizing the reciprocal points classification loss based negative log-probability of the true class k via SGD as:

$$\mathcal{L}_c(x; \theta, \mathcal{P}) = -\log p(y = k|x, f_\theta, \mathcal{P}) \quad (6)$$

which can be further seen as integrating multiple binary classification tasks to solve the multiclass open set recognition problem. Through minimizing Eq. (6), on the one hand, maximizing the otherness between known data \mathcal{D}_L and reciprocal points set \mathcal{P} has a facilitating effect on maximizing the interval between closed space \mathcal{S}_k and open space \mathcal{O}_k , which is consistent with our initial goal. On the other hand, corresponding to $\mathcal{R}_\epsilon(f, \mathcal{D}_L)$ in the Eq. 2, the reciprocal points classification loss reduces the empirical classification risk through the reciprocal points.

3.3 Reducing Open Space Risk

In solving the open set recognition problem, in addition to utilizing the reciprocal points loss to manage the empirical risk, we also further reduce the open space risk $\mathcal{R}_o(f, \mathcal{D}_U)$ in Eq. 2 with introducing a regularization term. According to the notion of open set in Sec. 3.1, each particular category k has positive open space \mathcal{O}_k^{pos} and infinite negative open space \mathcal{O}_k^{neg} respectively. For multiclass open-set recognition scenarios, we further union multiple class-wise open spaces into a global open space \mathcal{O}_G as:

$$\mathcal{O}_G = \bigcap_{k=1}^N (\mathcal{O}_k^{pos} \cup \mathcal{O}_k^{neg}) \quad (7)$$

where we would restrict the total open space risk for all known classes. Moreover, based on the principle of maximum entropy, for an unknown sample x_u without

Algorithm 1 The reciprocal point learning algorithm.

Input: Training data $\{x_i\}$. Initialized parameters θ in convolutional layers. Parameters \mathcal{P} and R in loss layers, respectively. Hyperparameter λ, γ and learning rate μ . The number of iteration $t \leftarrow 0$.

Output: The parameters θ , \mathcal{P} and R .

- 1: while *not converge* do
 - 2: $t \leftarrow t + 1$.
 - 3: Compute the joint loss by $\mathcal{L}^t = \mathcal{L}_c^t + \lambda \cdot \mathcal{L}_o^t$.
 - 4: Compute the backpropagation error $\frac{\partial \mathcal{L}^t}{\partial x^t}$ for each i by $\frac{\partial \mathcal{L}^t}{\partial x^t} = \frac{\partial \mathcal{L}_c^t}{\partial x^t} + \lambda \cdot \frac{\partial \mathcal{L}_o^t}{\partial x^t}$.
 - 5: Update the parameters \mathcal{P} by $\mathcal{P}^{t+1} = \mathcal{P}^t - \mu^t \cdot \frac{\partial \mathcal{L}^t}{\partial \mathcal{P}^t} = \mathcal{P}^t - \mu^t \cdot (\frac{\partial \mathcal{L}_c^t}{\partial \mathcal{P}^t} + \lambda \cdot \frac{\partial \mathcal{L}_o^t}{\partial \mathcal{P}^t})$.
 - 6: Update the parameters R by $R^{t+1} = R^t - \mu^t \cdot \frac{\partial \mathcal{L}^t}{\partial R^t} = R^t - \lambda \cdot \mu^t \cdot \frac{\partial \mathcal{L}_o^t}{\partial R^t}$.
 - 7: Update the parameters θ by $\theta^{t+1} = \theta^t - \mu^t \cdot \frac{\partial \mathcal{L}^t}{\partial \theta^t} = \theta^t - \mu^t \cdot (\frac{\partial \mathcal{L}_c^t}{\partial \theta^t} + \lambda \cdot \frac{\partial \mathcal{L}_o^t}{\partial \theta^t})$.
 - 8: end while
-

any prior information, a well-trained closed-set discriminant function tends to assign known labels to x_u with equal probability, which leads to the deep neural networks usually embed unknown samples into the inside of known spaces rather than random positions in the full space. This phenomenon is also consistent with the observation of [6] and our visualization results shown in Fig. 4.

Obviously, the idea of which wants to manage open space risk by restricting the open space to a bounded space directly is almost impossible because the open space contains a large number of potentially unknown samples. However, considering space \mathcal{S}_k and \mathcal{O}_k are complementary with each other, we indirectly bound the open space risk by constraining the distance between samples from \mathcal{S}_k and reciprocal points \mathcal{P}^k to a certain extent as:

$$\mathcal{L}_o(x; \theta, \mathcal{P}^k, R^k) = \frac{1}{M} \cdot \sum_{i=1}^M \|d(f_\theta(x), p_i^k) - R^k\|_2^2 \quad (8)$$

where R^k is the learnable margin in our framework. Specifically, minimizing the Eq. (8) is equivalent to making R^k and $\zeta(\mathcal{D}_L^k, \mathcal{P}^k)$ in Eq. (3) as close as possible, so that the following formula is established as:

$$\max(\zeta(\mathcal{D}_L^{\neq k} \cup \mathcal{D}_U, \mathcal{P}^k)) \leq R^k \quad (9)$$

Considering bounded space $\mathcal{B}(\mathcal{P}^k, R^k)$ with the reciprocal points \mathcal{P}^k as centers and R^k as its corresponding intervals, in order to separate known and unknown space, we further utilize these bounded spaces to approximate the global unknown space \mathcal{O}_G as much as possible. As a result, calculating the regularization loss with Eq. (8) can be viewed as managing open space risk $\mathcal{R}_o(f, \mathcal{D}_U)$ in Eq. (6).

3.4 Learning Open Set Network

Finally, the overall loss function of reciprocal points learning combines the empirical classification risk and reducing open space risk as:

$$\mathcal{L}(x; \theta, \mathcal{P}, R) = \mathcal{L}_c(x; \theta, \mathcal{P}) + \lambda \mathcal{L}_o(x; \theta, \mathcal{P}, R) \quad (10)$$

where λ is a hyper-parameter of controlling the weight of reducing open space risk module and θ, \mathcal{P}, R represent the learnable parameters of our framework. As shown in Fig. 2, learning open set recognition problem with Eq. (10) results in pushing the known spaces to the periphery of global open space and then separating the two spaces as much as possible. In Algorithm 1, we summarize the learning details of the open set network with joint supervision.

How to detect unknown classes? Samples from a class on which the network was not trained would have probability scores distributed across all the known classes [6]. Similar to thresholding softmax scores in [11], we agree that unknown samples have a closer distance with all reciprocal points than samples of known training classes based on a deep classifier. Therefore, the probability that test instance x belongs to known classes is proportional to the distance between the test instance and the farthest reciprocal point corresponding to category k :

$$p(\text{known}|x) \propto \max_{k \in \{1, \dots, N\}} d(f(x), \mathcal{P}^k) \quad (11)$$

The models for open set problem focus on two key questions, **Q1**) *what is a good score for open-set identification?* (i.e., identifying a class as known or unknown), and given a score, **Q2**) *what is a good operating threshold for the model?* [21]. Furthermore, since how rare or common samples from unknown space is not known in the actual scenario, the approaches to open set recognition which require an arbitrary threshold or sensitivity for comparison is unreasonable [19]. Thus, we utilize the difference between known and unknown distribution to measure the learned models' ability to detect unknown, which provides a calibration-free measure of detection performance.

4 Experiments

4.1 Experiments for Open Set Identification

Evaluation Metrics. The evaluation protocol defined in [19] is considered and Area Under the Receiver Operating Characteristic (AUROC) curve is used as evaluation metric. AUROC is a threshold-independent metric [4]. It can be interpreted as the probability that a positive example is assigned a higher detection score than a negative example [7]. Following the protocol in [19], we report the AUROC averaged over five randomized trials.

Network Architecture. Except for RPL-WRN, the encoder and decoder architecture for this experiment is same to the architecture used in [19]. For a fair comparison, we trained RPL with a Wide-ResNet [28] with depth 40, width 4

Table 1. The AUROC results of open set identification. The second column shows the backbone for each method, including Encoder (E) or Decoder (D). The architecture used to infer the unknown class are highlighted in bold. Values other than RPL are taken from [21]. RPL++ means using GCPL to assist RPL training. Best performing methods are highlighted in bold.

Method	Backbone	MNIST	SVHN	CIFAR10	CIFAR+10	CIFAR+50	TinyImageNet
Softmax	E	97.8	88.6	67.7	81.6	80.5	57.7
Openmax [2]	E	98.1	89.4	69.5	81.7	79.6	57.6
RPL	E	98.9	93.4	82.7	84.2	83.2	68.8
RPL++	E	99.3	95.1	86.1	85.6	85.0	70.2
G-OpenMax [8]	E + E + D	98.4	89.6	67.5	82.7	81.9	58.0
OSRCI [19]	E + E + D	98.8	91.0	69.9	83.8	82.7	58.6
C2AE [21]	E + D	98.9	92.2	89.5	95.5	93.7	74.8
RPL-WRN	E(WRN)	99.6	96.8	90.1	97.6	96.8	80.9

and dropout 0, WRN-40-4. The parameters of WRN-40-4 are about 9M, which is less than the networks of C2AE [21] and OSRCI [19]. λ for RPL is set to 0.1 in the training stage. More details in the supplementary material.

Datasets. Similar to [21], we provide a simple summary of these protocols for each dataset.

- **MNIST, SVHN, CIFAR10.** For MNIST [14], SVHN [20] and CIFAR10 [13], by randomly sampling 6 known classes and 4 unknown classes.
- **CIFAR+10, CIFAR+50.** For CIFAR+ N experiments, 4 classes are sampled from CIFAR10 for training. N non-overlapping classes are used as unknown, which are sampled from the CIFAR100 dataset [13].
- **TinyImageNet.** For experiments with TinyImageNet [22], 20 known classes and 180 unknown classes are randomly sampled for evaluation.

Result Comparisons. As shown in Table 1, RPL outperforms other methods based on encoder in [19] in all datasets. Noted that RPL with only using known training samples and a single encoder outperforms methods based on encoder and decoder in MNIST and SVHN. With less information and simple training method, RPL achieves better performance than OSRCI [19] and G-OpenMax [8] using generation-based model on all datasets. Furthermore, we need fewer network parameters and a more simple strategy for training to identify unknown classes. RPL-WRN has similar parameters to C2AE model and performs significantly better than other recent state-of-the-art methods in SVHN, CIFAR, and TinyImageNet. It shows that open set network trained by RPL learns better scores for identifying known and unknown classes.

4.2 Experiments for Open Long-Tailed Recognition

Datasets. In real visual world, the frequency distribution of visual classes is long-tailed, with a few common classes and many more rare classes [17]. To simulate natural data distribution, we adopt ImageNet-LT [17] and a new image datasets, Air-300, to evaluate our algorithm for open set identification.

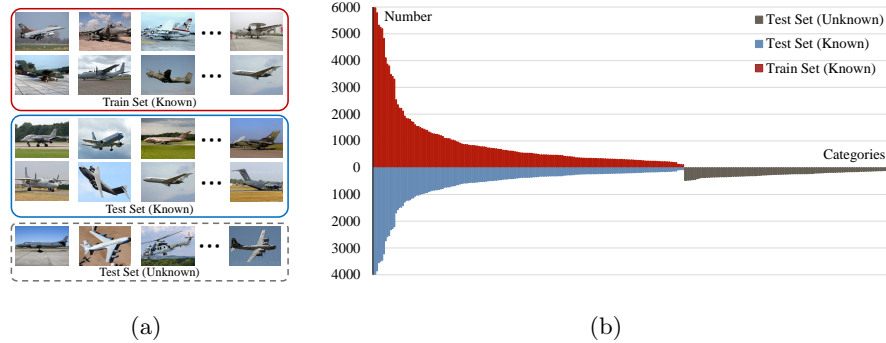


Fig. 3. (a) The overview of Air-300. (b) The number of images with classes.

ImageNet-LT. ImageNet-LT is a long-tailed version of the original ImageNet-2012 [5]. It has 115.8K images from 1000 categories, with maximally 1280 images per class and minimally 5 images per class. The test set contains the open set based on the additional classes of images in ImageNet-2010.

Aircraft 300 Dataset. To facilitate research on open set recognition and evaluate the performance of different approaches, we further construct a new large-scale dataset from the Internet, **Aircraft 300** (Air-300), which contains 320,000 annotated colour images from 300 different classes in total, see thumbnails in Figure 3(a) briefly. Each category contains 100 images at least, and a maximum of 10,000 images, which leads to the long tail distribution as shown in Figure 3(b). According to the number of images in each category, we then divide all classes into two parts with 180 known classes for training and 120 novel unknown classes for testing respectively. Furthermore, we aggregate 300 different aircraft models into 10 super classes according to their roles, such as a bomber, fighter, helicopter and so on. As such, each image is annotated with a fine-grain label (the sub-class category) and a coarse-grain label (the superclass category) by human. Note we also split these 10 superclasses into two parts for coarse-grain open set recognition. Compared with the existing benchmark datasets, the tailored Air-300 dataset is not only much larger and brings new challenges to open set recognition, but also can be utilized to perform fine-grain classification and object recognition respectively.

Evaluation Metrics. The AUROC and AUPR are adopted for evaluation, AUPR-Known and AUPR-Unknown denote the area under the precision-recall curve where known and unknown are specified as positives, respectively.

Network Architecture. All methods are trained based on WRN-40-4 for Air-300 and ResNet50 for ImageNet-LT [17]. We optimize GCPL performance to best by adjusting its parameters. More details in the supplementary material.

Result Comparisons. Table. 2 shows the test accuracy of known classes and the performance of open set identification on Air-300 and ImageNet-LT. Under the $\lambda = 0.1$ settings, RPL has greatly improved in all evaluations compared

Table 2. Test accuracy and open set identification of different methods on Air-300.

Method	Air-300 (WRN-40-4)				ImageNet-LT (ResNet50)			
	Acc	AUROC	AUPR-K	AUPR-U	Acc	AUROC	AUPR-K	AUPR-U
Softmax	85.7	77.1	85.8	48.0	37.8	53.3	45.0	60.5
GCPL	84.5	79.1	88.0	61.3	37.1	54.5	45.7	61.9
$\lambda = 0$	84.6	64.6	78.5	43.6	38.7	55.0	46.0	62.5
RPL $\lambda = 0.01$	87.9	82.5	89.3	68.9	39.0	55.0	46.0	62.5
RPL $\lambda = 0.1$	88.8	83.1	89.3	70.2	39.0	55.1	46.1	62.7
$\lambda = 1$	86.7	81.7	87.5	69.7	38.2	54.4	45.7	61.6
RPL++	89.0	84.1	90.1	72.7	39.7	55.2	46.2	62.7

to softmax and GCPL. For air-300, RPL not only improves the classification accuracy of closed set, but also is about 6% higher than softmax in AUROC. ImageNet-LT has more categories and fewer training samples, which makes it very challenging for open set recognition. It is difficult to detect unknown when the accuracy of closed set is not so well, but RPL still improves the performance of classification and open set identification by nearly 2%. It shows that RPL is effective for the open set problem of long tail distribution.

4.3 Further Analysis

Analysis of Reducing Open Space Risk.

RPL vs. Softmax. Reciprocal points classification loss term (the first part of Eq. 2) defines constraints for reciprocal point. However, similar to softmax, the learned representation is still linear separable only with this classification loss. See Fig.4(a) and 4(b), reciprocal points without \mathcal{L}_o are learned to the origin, while RPL with \mathcal{L}_o can achieve better distribution (shown in Fig. 4(d)). For softmax and RPL without \mathcal{L}_o , there is a significant overlap between the known and the unknown classes, so as to the neural network views unknown samples as some known classes with high confidence. In contrast, the whole feature space learned by RPL with \mathcal{L}_o is in a limited range (1-6 in abscissa for unknowns and 5-8 in abscissa for knowns in Fig.4(d)), which prevents high confidence in unknown classes.

RPL vs. GCPL. As shown in Fig.4(c), GCPL used the prototypes to reduce intra-class variance. However, without considering the unknown, GCPL extends unknown classes to the whole feature space, resulting in a significant overlap with known. By introducing \mathcal{L}_o , different known categories are spread to the periphery of the space, while unknown categories are restricted to the interior. It can be seen that a clear gap is maintained between the two types of samples (known vs. unknown) in Fig.4(d). RPL improves the robustness of neural networks by preventing the misjudgment of the unknown class through the bounded restriction, thereby enhancing and stabilizing the classification of known categories.

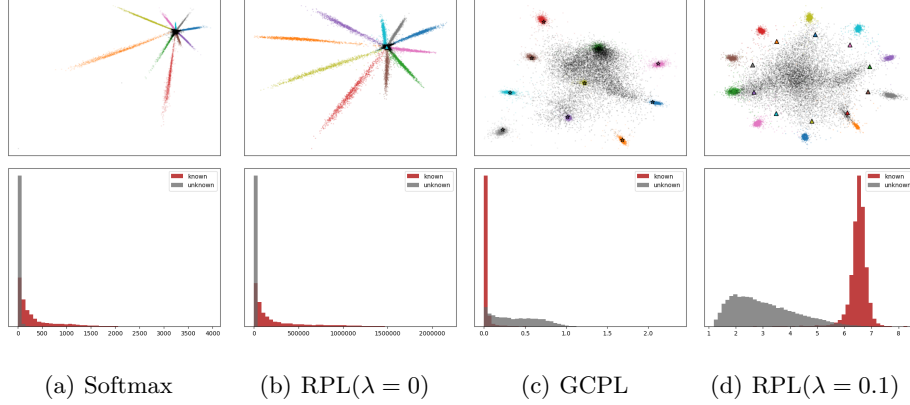


Fig. 4. (1) The first row is visualization in the learned feature space of MNIST as known and FashionMNIST as unknown. Gray shapes are data of unknown, and circles in color are data of known samples. Different colors represent different classes. Colored stars and triangles represent the prototypes and the reciprocal points learned of different known categories, respectively. (2) The second row is the maximum distance distribution between features and prototypes or reciprocal points.

Margin & λ . See Table. 2, the performance of RPL is related to the setting of λ . With the increase of λ , RPL has greatly improved on open set identification. However, when λ is too big, the accuracy of known classes will decrease slightly, which is related to \mathcal{L}_o limiting the size of the whole feature space, while the learning margin in RPL can show the size of the whole feature space. Fig. 5(a) shows the variation trend of margin with λ . With the increase of λ , the learned margin becomes smaller, and the limitation of \mathcal{L}_o on feature space also tends larger. Different training data need different reasonable margins and space sizes to ensure that known and unknown can be classified correctly. Fig. 5(a) proves that the margin increases with the number of known classes with fixed lambda. Reasonable λ can effectively control the interaction among known classes, by learning the more appropriate embedding space size.

Experiments with Multiple Reciprocal Points. Fig. 5(b) shows the impact of increasing the number of reciprocal points on open set identification to detect rare unknown classes. We can observe that the performance of open set identification is relatively stable. The performance is slightly improved when employing 8 reciprocal points. Based on a good feature space distribution, increasing the number of reciprocal points for each known category can slightly improve the ability of the neural networks to distinguish unknown. However, too many reciprocal points will bring more training overhead and affect performance.

Combination of Reciprocal Point and Prototype. The reciprocal point can push known classes to the periphery in the limited feature space, thereby improving the performance of the open set identification. We deem that the reciprocal point and the prototype are opposite and complementary. Here we combine the




Fig. 5. (a) The trend of margin with λ and the number of known classes. (b) is the variation trend of testing performance with multiple reciprocal points for open set identification in Air-300. The abscissa represents the number of reciprocal points.

reciprocal point and the prototype, as *RPL++*. In the training phase of RPL, we add prototypes and further train reciprocal points and prototypes by RPL and prototype loss [25]. As shown in Table. 1, RPL++ achieves better performance than RPL. Moreover, the prototype further improves the performance of RPL on open long-tailed recognition in Table. 2. The reciprocal points form an excellent embedding space structure, then the prototype further narrows intra-class variance, to further divide the known classes and the unknown classes.

5 Conclusion

This paper proposes a new concept, *Reciprocal Point*, which is the potential representation of the extra-class space corresponding to each known category. We also propose a novel learning framework towards open set learning. The approach introduces unknown information by reciprocal points, to optimize a better feature space to separate known and unknown. Comprehensive experiments conducted on multiple datasets demonstrate that our method outperforms previous state-of-the-art open set classifiers in all cases. We also publish a open long-tailed dataset, the Air-300, which is a challenging dataset to simulate natural data distribution for open set recognition and other visual tasks.

Acknowledgments

This work is partially supported by grants from the National Key R&D Program of China under grant 2017YFB1002400, the Key-Area Research and Development Program of Guangdong Province under Grant 2019B010153002, and the National Natural Science Foundation of China under contract No.61825101, No.61702515 and No.U1611461.

References

1. Bartlett, P.L., Wegkamp, M.H.: Classification with a reject option using a hinge loss. *Journal of Machine Learning Research* **9**(Aug), 1823–1840 (2008)
2. Bendale, A., Boulton, T.E.: Towards open set deep networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1563–1572 (2016)
3. Da, Q., Yu, Y., Zhou, Z.H.: Learning with augmented class by exploiting unlabeled data. In: *Twenty-Eighth AAAI Conference on Artificial Intelligence* (2014)
4. Davis, J., Goadrich, M.: The relationship between precision-recall and roc curves. In: *Proceedings of the 23rd international conference on Machine learning*. pp. 233–240. ACM (2006)
5. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 248–255. Ieee (2009)
6. Dhamija, A.R., Günther, M., Boulton, T.: Reducing network agnostophobia. In: *Advances in Neural Information Processing Systems*. pp. 9157–9168 (2018)
7. Fawcett, T.: An introduction to roc analysis. *Pattern recognition letters* **27**(8), 861–874 (2006)
8. Ge, Z., Demyanov, S., Chen, Z., Garnavi, R.: Generative openmax for multi-class open set classification. *arXiv preprint arXiv:1707.07418* (2017)
9. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1026–1034 (2015)
10. Hein, M., Andriushchenko, M., Bitterwolf, J.: Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 41–50 (2019)
11. Hendrycks, D., Gimpel, K.: A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136* (2016)
12. Hendrycks, D., Mazeika, M., Dietterich, T.G.: Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:1812.04606* (2018)
13. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. *Tech. rep., Citeseer* (2009)
14. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
15. Lee, K., Lee, H., Lee, K., Shin, J.: Training confidence-calibrated classifiers for detecting out-of-distribution samples. *arXiv preprint arXiv:1711.09325* (2017)
16. Liang, S., Li, Y., Srikant, R.: Enhancing the reliability of out-of-distribution image detection in neural networks. *arXiv preprint arXiv:1706.02690* (2017)
17. Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., Yu, S.X.: Large-scale long-tailed recognition in an open world. *arXiv preprint arXiv:1904.05160* (2019)
18. Nalisnick, E., Matsukawa, A., Teh, Y.W., Gorur, D., Lakshminarayanan, B.: Do deep generative models know what they don’t know? *arXiv preprint arXiv:1810.09136* (2018)
19. Neal, L., Olson, M., Fern, X., Wong, W.K., Li, F.: Open set learning with counterfactual images. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 613–628 (2018)
20. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning (2011)

21. Oza, P., Patel, V.M.: C2ae: Class conditioned auto-encoder for open-set recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2307–2316 (2019)
22. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* **115**(3), 211–252 (2015)
23. Scheirer, W.J., de Rezende Rocha, A., Sapkota, A., Boulton, T.E.: Toward open set recognition. *IEEE transactions on pattern analysis and machine intelligence* **35**(7), 1757–1772 (2013)
24. Wen, Y., Zhang, K., Li, Z., Qiao, Y.: A discriminative feature learning approach for deep face recognition. In: European conference on computer vision. pp. 499–515. Springer (2016)
25. Yang, H.M., Zhang, X.Y., Yin, F., Liu, C.L.: Robust classification with convolutional prototype learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3474–3482 (2018)
26. Yoshihashi, R., Shao, W., Kawakami, R., You, S., Iida, M., Naemura, T.: Classification-reconstruction learning for open-set recognition. *arXiv preprint arXiv:1812.04246* (2018)
27. Yuan, M., Wegkamp, M.: Classification methods with reject option based on convex risk minimization. *Journal of Machine Learning Research* **11**(Jan), 111–130 (2010)
28. Zagoruyko, S., Komodakis, N.: Wide residual networks. *arXiv preprint arXiv:1605.07146* (2016)