Deep Uniformly Distributed Centers on a Hypersphere for Open Set Recognition

Hakan Cevikalp HCEVIKALP@OGU.EDU.TR
Hasan Serhan Yavuz HSYAVUZ@OGU.EDU.TR
Eskisehir Osmangazi University, Machine Learning and Computer Vision Laboratory, Eskisehir, Turkey.

Hasan Saribas HASAN.SARIBAS1@HUAWEI.COM
Huawei Turkey R&D Center, Istanbul, Turkey.

Editors: Berrin Yanıkoğlu and Wray Buntine

Abstract

This study introduces a new approach for open set recognition, wherein we propose a novel method utilizing uniformly distributed centers on a hypersphere. Each class in the proposed method is represented by a center, and these centers and features of the deep learning architecture are jointly learned from the training data in an end-to-end fashion. We ensure that the centers lie on the boundary of a hypersphere whose center is positioned at the origin. The class-specific samples are compelled by the proposed loss function to be closer to their respective centers. In open set recognition scenarios, an additional loss term is employed to separate the background samples from the known class centers. The assignment of test samples to classes is based on the Euclidean distances calculated from the learned class centers. Experimental results show that the proposed method yields the state-of-the-art accuracies on open set recognition datasets.

Keywords: Open set recognition, classification, deep learning, uniformly distributed centers.

1. Introduction

Theoretical studies indicate that high-dimensional data samples tend to cluster around the outer surface of an expanding hypersphere, (Jimenez and Landgrebe, 1998; Hall et al., 2005). When the number of samples is fixed and the dimension of the feature space is increased, the radius of the hypersphere increases as well. As a result, the hyperspherical spaces are largely used in pattern classification problems. Especially, all the recent state-of-the-art deep face recognition methods such as CosFace (Wang et al., 2018), UniformFace (Duan et al., 2019), ArcFace (Deng et al., 2019), etc. utilize the hyperspherical spaces. These methods use the revised softmax losses where the class-specific weights and feature vectors are normalized so that they all lie on the boundary of a hypersphere with a predefined radius value. Both additive and angular margins are used for separation among the classes (i.e., inter-class separation). Cosine distances between the test feature samples and class-specific weights are used during testing stage to assign the test samples to the known classes. Although these methods yield the state-of-the-art accuracies on classical face recognition problems, their performance is unsatisfactory when it comes to open-set recognition problems. In open-set recognition, the test samples may originate from unknown classes that were not encountered during the training phase.

In this study, we propose a novel method called the Deep Uniformly Distributed Centers on a Hypersphere (DUDC-HS) for both closed and open set recognition problems. In the proposed method, the classes are represented with the centers that are uniformly distributed on the boundary of a hypersphere. During the training stage, the samples belonging to the known classes are forced to lie closer to the centers that represent them. Both the class centers and features are jointly learned from the data. In contrast to the other methods using hyperspheres, we do not enforce the class samples to lie on the boundary of the hypersphere and we utilize Euclidean distances between the samples and their corresponding centers instead of cosine distances.

1.1. Related Work

We propose a novel method using hyperspherical spaces for open set recognition. Therefore, we summarize the related works on hyperspheres and open set recognition.

1.1.1. Classifiers Using Hyperspheres

The methods using hyperspherical output spaces can be broadly categorized into two groups depending on whether they utilize class-specific prototypes or classifier weights. Among the methods using class-specific prototypes, (Mettes et al., 2019) utilizes class-specific prototypes chosen from the outer boundary of a unit hypersphere. As opposed to our work, they do not learn the prototypes from the data samples as we did. Instead, they place the prototypes, that represent the classes, on the hypersphere based on the semantic relations obtained from word embeddings of the class names. There are also similar deep neural network classifiers using class representative prototypes (or centers) that are selected from the boundary of a hypersphere, (Cevikalp and Saribas, 2023; Bytyqi et al., 2023; Li et al., 2022; Graf et al., 2021). These methods use the vertices of a regular simplex enclosed within a hypersphere with the following constraint on the feature dimensionality: the dimension of the feature space, denoted as d, must be equal to or greater than the total number of classes minus one, $d \geq C-1$, where C represents the number of classes. These methods also use prefixed vectors representing classes and these vectors are not learned from data as in our proposed method. During the testing stage, the test samples are assigned to the classes based on the Cosine or Euclidean distances between the samples and class prototypes.

The methods using classifier weights modify the conventional softmax loss function to maximize the angular margins within the hyperspherical output spaces. To this end, (Wang et al., 2018) introduced the CosFace method which imposes an additive angular margin on the learned deep CNN (Convolutional Neural Network) features. Both the features and the learned class weights are normalized so that they lie on the surface of a hypersphere with a radius, s, that is set to larger values. An additive margin is used for separation of the classes. A similar method called the ArcFace using additive margins is proposed in (Deng et al., 2019). As opposed to the using additive margins, SphereFace method (Liu et al., 2017, 2016) used multiplicative margins for inter-class separation. Also, this method did not normalize the sample feature vectors as in CosFace and ArcFace methods, but a new variant of the SphereFace method (Liu et al., 2023) normalizes both feature vectors and class weights so that they lie on the surface of a hypersphere. A method called UniformFace which is similar to our proposed one is introduced in (Duan et al., 2019). This method aims to learn equidistributed representations on hyperspherical spaces. To this end, the authors adopt the angular softmax used in SphereFace method and they combine it with another loss term that encourages uniformly distributed CNN features on the hypersphere. All face recognition methods using classifier weights utilize the angles between test samples and class weights for class label assignments.

In addition to these methods, there are also methods that introduce loss functions for learning uniformly distributed representations on the hypersphere manifold through potential energy minimization, (Liu et al., 2018; Lin et al., 2020; Liu et al., 2021). However, these studies primarily focus on addressing the issue of layer regularization rather than the direct classification problem and apply hyperspherical uniformity to the learned weights. The key concept is to train deep neural network weights that are uniformly distributed on a hypersphere, with the aim of minimizing redundancy among the learned weights. Some of these methods (e.g., (Liu et al., 2021)) also require that the weights must be orthogonal to each other. This yields a constraint that the dimension must exceed the total number of learned weight vectors, similar to the restriction observed in methods utilizing regular simplex vertices for approximating the classes.

There are some classifiers that use hyperspheres for approximating classes as in (Cevikalp et al., 2023; Yang et al., 2020; Uzun et al., 2023; Cho and Choo, 2022). However, these methods do not use hyperspherical output spaces, therefore they are quite different than our proposed method here. Instead, each class is represented with a class-specific hypersphere. The hypersphere centers are learned from data as in (Cevikalp et al., 2023; Yang et al., 2020; Uzun et al., 2023) or randomly chosen from some distributions and fixed to some specific vectors, (Cho and Choo, 2022). These hyperspheres can arbitrary lie indifferent regions of the feature space. Test samples are classified based on the closest distances from the samples to the hypersphere centers.

1.1.2. Open Set Recognition

In open set recognition problems, there is a possibility of encountering novel classes (that were not used in the training phase) during the testing phase. The objective is to accurately classify the samples belonging to the known classes while rejecting those from unknown classes, (Scheirer et al., 2013). Recent approaches aim to address this open set recognition problem by restricting the acceptance regions of the known classes. In this regard, the Convolutional Prototype Network was proposed (Yang et al., 2020), which learns multiple prototype vectors to represent each class. It enforces the known class samples to cluster tightly around these learned prototypes, and novel test samples are rejected based on their distances to the prototype vectors. In a similar manner, (Cevikalp et al., 2023) uses hypersphere centers for approximation of known classes in the context of open set recognition. Both methods jointly learn the class centers (prototypes) and feature representations from data in an end-to-end manner, but they do not use hyperspherical output spaces since the hypersphere centers can arbitrary lie indifferent regions of the feature space. The Deep Polyhedral Conic Classifier method of (Cevikalp et al., 2021) also returns bounded and compact acceptance regions for the known classes. It achieves this by simultaneously maximizing the separation between different classes and minimizing the variations within each

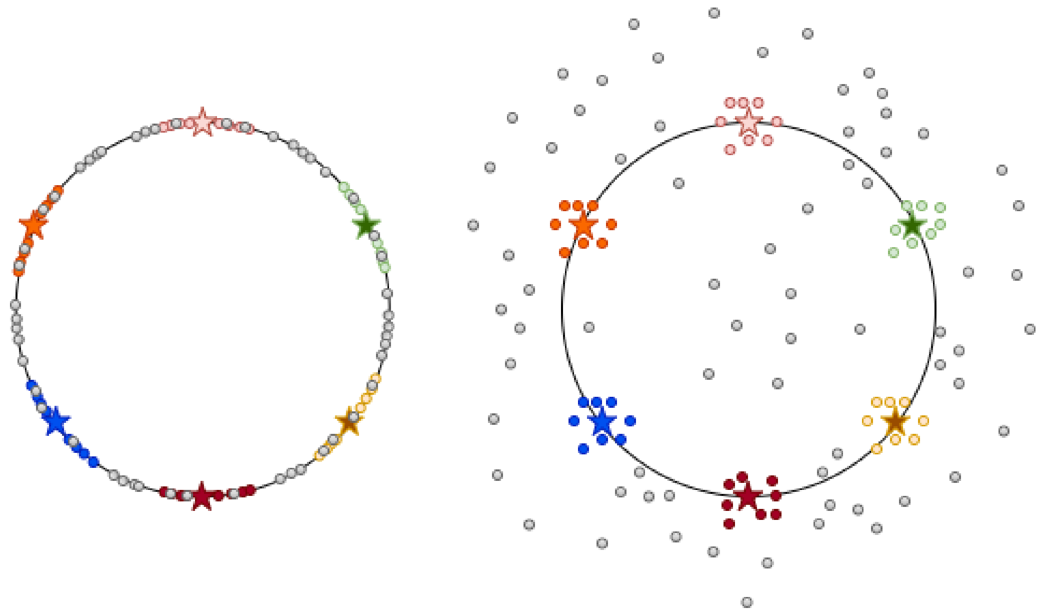

Figure 1: Comparison of the embedding spaces of the proposed method and other state-of-the-art face recognition methods maximizing the margin in angular spaces in the context of open set recognition: The known class centers are shown with different colored star symbols, and corresponding class samples are represented with the circle symbols having the same color. The unknown class samples are represented with gray colored circles. When both the features and classifier weights are normalized as in face recognition methods (e.g., CosFace and ArcFace), the known and unknown class samples overlap with each other and it becomes harder to identify and reject the unknonw class samples. In contrast, in the proposed method, we can easily identify and reject the unknown class samples based on the Euclidean distances to the known class centers as seen on the right since the features of samples are not normalized.

class using polyhedral conic functions. Another method described in (Miller et al., 2021) encourages the known class training data to tightly cluster around class-specific centers in the logit space. This method operates in the logit space and aims to cluster the classification scores around anchored class centers, positioned on the respective class coordinate axes. Each anchored class center is represented as a scaled one-hot vector. Consequently, this ensures equal distances between all classes, restricting the learning of semantically meaningful features, as mentioned by the authors of (Miller et al., 2021). To address open set recognition, (Dhamija et al., 2018) introduced entropic open set and object-to-sphere loss functions. These functions yield feature embedding spaces where known class samples exhibit a large feature magnitude and low entropy, while unknown samples have a smaller magnitude.

1.2. Motivation and Contributions

State-of-the-art face recognition methods maximizing the margin in angular spaces yield excellent accuracies for closed set recognition problems, yet their accuracies are not satisfactory in open set recognition settings. The reason behind this is demonstrated in Fig. 1, where the known class centers are denoted by star symbols and the corresponding class samples are represented by circle symbols of the same color. In open set recognition settings, samples coming from the unknown classes are used, and these samples are represented with the gray circles in the figure. Popular face recognition methods such as CosFace, ArcFace etc. normalize both the features and classifier weights so that they all lie on the boundary of a hypersphere as seen left part of the figure. Because of this, known and unknown class samples greatly overlap as seen in the figure despite the known class samples are easily separated from each other with a large angular margin. In this case, it is impossible to identify and reject many unknown class samples, therefore the accuracy drops significantly in open set recognition settings. In contrast, when the all features space is used as seen on the right part of the figure, we can still successfully identify and reject the unknown class samples based on the Euclidean distances to the known class centers. Therefore, we propose a novel method that uses uniformly distributed centers lying on the boundary of a hypersphere for approximation of the classes. In contrast to the face recognition methods maximizing the margin in angular spaces, we do not normalize the features of samples to lie on the boundary of the hypersphere in order to prevent overlap with the unknown class sample features during the testing stage. We also utilize the Euclidean distances rather than cosine distances for assigning labels to the test samples. The class centers are completely learned from the training data in end-to-end manner, therefore our proposed method differs from other methods using predetermined centers chosen from the vertices of a regular simplex inscribed in a hypersphere or determined based on semantic word embeddings.

Our contributions can be summarized as follows:

- The proposed method learns uniformly distributed centers that approximate the known classes based on the training data in end-to-end manner. Therefore, there is no need to set the centers in advance, which is a tedious task. The proposed method also learns semantically related features as shown in the experiments.
- As opposed to the class centers, the learned CNN features are not inherently normalized to reside on the boundary of the hypersphere. As a result, it becomes relatively simpler to identify and reject the unknown samples by evaluating the Euclidean distances to the known class centers.
- The proposed method effectively handles imbalanced datasets since it independently minimizes the distances between the samples and their respective centers, without being influenced by class imbalances.

2. Deep Uniformly Distributed Centers on a Hypersphere Method

In this paper, we propose a novel method that uses uniformly distributed centers on a hypersphere for classification. In the proposed method, each class is approximated with a center lying on the boundary of a hypersphere, and the test samples are classified based on the smallest Euclidean distances between the test samples and the class-specific centers. The problem of uniformly distributing points on a hypersphere is largely studied in the literature in different contexts, e.g., for finding an equilibrium state with the minimum potential energy that distributes C electrons on a unit hypersphere as evenly as possible known as Thomson problem, (Thomson, 1904), or best-packing problem on hypersphere known as Tammes problem, (Tammes, 1930). As theoretically proved in (Liu et al., 2021), when the dimension of the feature space is larger than the number of centers, vertices of a regular simplex inscribed in the hypersphere offer the optimal solution in terms of margin maximization. However, placing the centers uniformly on the hypersphere is an open mathematical problem when this criterion is not satisfied. Instead of solving this tedious problem and placing the centers to fixed positions, we learn the centers' positions from the training data and learn both the features and centers jointly in an end-to-end manner in the proposed methodology.

Let us assume that the deep neural network features of training samples are given in the form $(\mathbf{f}_i, y_i)$, $i = 1, \ldots, n$, $\mathbf{f}_i \in \mathbb{R}^d$, $y_i \in \{j\}$ where $j = 1, ..., C$. Here, C is the total number of known classes. Let $\mathbf{s}_j$ denote the center of the $j-$th class which will be used for approximation of that class. The class centers must lie on the boundary of a hypersphere whose center is the origin with a radius u, thus the lengths of the center vectors must satisfy the restriction, $\|\mathbf{s}_j\| = u$ for $j = 1, ..., C$. In this case, the loss function of the proposed deep neural network classifier can be written as,

$$\mathcal{L} = \frac{1}{n}\sum_{i=1}^{n} \|\mathbf{f}_i - \mathbf{s}_{y_i}\|^2 + \lambda \sum_{i=1}^{n} \sum_{j=1, j\neq y_i}^{C} \max\left(0, m + \|\mathbf{f}_i - \mathbf{s}_{y_i}\|^2 - \|\mathbf{f}_i - \mathbf{s}_j\|^2\right) + \kappa \sum_{j=1}^{C} \sum_{\widetilde{j}=1, \widetilde{j}\neq j}^{C} \frac{1}{\left\|\mathbf{s}_j - \mathbf{s}_{\widetilde{j}}\right\|_2 + 1}, \quad (1)$$

under the constraint, $\|\mathbf{s}_j\| = u$ for $j = 1, ..., C$.

The first term in the loss function minimizes the Euclidean distances between the class samples and their corresponding centers. The second term is used for the inter-class separation, and it ensures that the distances between the training samples and their corresponding class centers must be smaller than the distances between these samples and rival class centers by at least a selected margin, m. The last term targets to place the centers uniformly on the hypersphere. To this end, we adopt the same loss term used in the UniformFace method (Duan et al., 2019) for this purpose. In this loss term, uniformity is defined as the potential energy of all the centers, and the repulsion between two class centers is represented by the inverse of the distance between the centers. In order to learn equidistributed representations, the potential energy of all class centers is minimized via so-called uniform loss function which is set to the average of all pairwise repulsions. It should be noted that our proposed method is quite different than the UniformFace method despite we utilize a common loss term since we use the Euclidean distances for decision making and employ a completely different loss term than the Angular Softmax (A-Softmax) used in the UniformFace. The margin term m, and the weight parameters λ and κ must be set by the user.

2.1. Including Background Samples for Open Set Recognition

In open set recognition scenarios, the training of classifiers commences by exclusively utilizing samples of known classes. Subsequently, both known and unknown class samples are employed in testing the resulting classifiers. The primary objective in this task is to ensure accurate classification of known class samples, while also detecting and rejecting samples from unknown classes, (Scheirer et al., 2013). Prior methods for open set recognition relied solely on the use of the known class samples during training. However, recent investigations (Dhamija et al., 2018; Miller et al., 2021; Cevikalp et al., 2023; Geng et al., 2021) have shown that augmenting the training dataset with the background dataset with samples coming from the classes that differ from the known classes can greatly enhances accuracy. The background samples can come various classes, thus they do not form a compact and coherent class group. Instead, they are scattered over all feature space similar to the unknown class samples depicted in Fig. 1. Let us represent the deep neural network features of the background samples by $\mathbf{f}_k \in \mathbb{R}^d$, $k = 1, ..., K$. In order to incorporate the background samples, we add an additional loss term that pushes the background samples away from the known class centers as follows:

$$\begin{aligned}\mathcal{L} = \frac{1}{n}\sum_{i=1}^{n} \|\mathbf{f}_i - \mathbf{s}_{y_i}\|^2 + \lambda \sum_{i=1}^{n} \sum_{j=1, j\neq y_i}^{C} \max\left(0, m + \|\mathbf{f}_i - \mathbf{s}_{y_i}\|^2 - \|\mathbf{f}_i - \mathbf{s}_j\|^2\right) \\ + \kappa \sum_{j=1}^{C} \sum_{\widetilde{j}=1, \widetilde{j}\neq j}^{C} \frac{1}{\left\|\mathbf{s}_j - \mathbf{s}_{\widetilde{j}}\right\|_2 + 1} + \eta \sum_{i=1}^{n}\sum_{k=1}^{K} \max\left(0, m + \|\mathbf{f}_i - \mathbf{s}_{y_i}\|^2 - \|\mathbf{f}_k - \mathbf{s}_{y_i}\|^2\right), \quad (2)\end{aligned}$$

under the constraint, $\|\mathbf{s}_j\| = u$ for $j = 1, ..., C$.

The first three terms of this new loss function are same as before. But, we add another loss term to ensure that the distances between the known class samples and their corresponding centers must be smaller than the distances between the background samples and known class centers by at least a selected margin, m. As a result, the background class samples are pushed away from the known class centers so that the trained model will form compact acceptance regions for the known classes. This way, the trained deep neural network classifier model will reject the unknown class samples more correctly.

3. Experiments

We tested the proposed method on both closed and open set recognition problems. We also conducted experiments to visualize the learned CNN features.

3.0.1. Illustrations

We first conducted some experiments to visualize the learned CNN features to verify that the class samples cluster in the vicinity of the learned centers that are uniformly distributed on a hypersphere. To this end, we used a small deep neural network that yields 2-dimensional CNN features. We trained the network by using 5 and 10 classes chosen from the Cifar10 dataset by using the proposed loss function given in (1). The hypersphere radius is set to 5 for 5 classes, and it is set to 12 for 10 classes. The CNN features of test samples returned

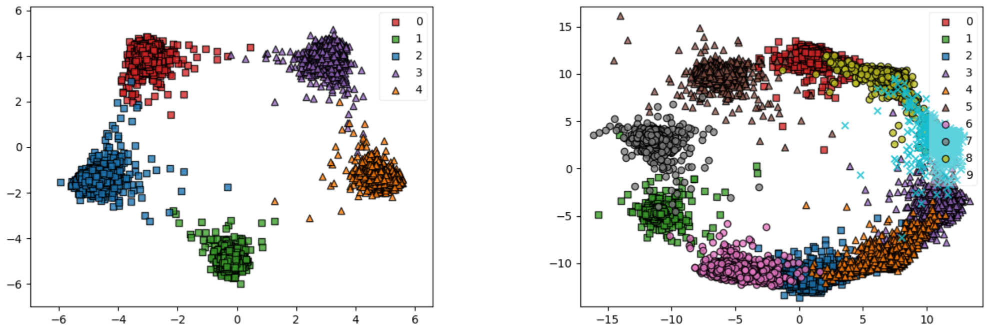

Figure 2: The outputs of the deep neural network classifiers trained by using the proposed loss function for 5 and 10 classes. For both cases, the class centers are uniformly distributed on the boundary of the hypersphere, and the class-specific samples compactly cluster in the vicinity of their class centers.

by the trained networks are illustrated in Fig. 2. As seen in the figure, the class centers are uniformly distributed on the hypersphere. In addition, the class specific samples lie near their corresponding centers as expected.

We also conducted experiments to verify that the proposed method returns semantically related embeddings. To verify this, we utilized a ResNet-18 architecture returning 512-dimensional CNN features. We trained the network by using Cifar10 classes and computed the learned centers for each class. Then, we computed the pair-wise distances between the class centers. Fig. 3 shows the distance matrix. As seen in the figure, the *automobile* class is closest to its semantically related *truck* class, the *cat* class is closest to the *dog* class, and *horse* class is closest to *deer* class. The distances also reflect appearance information as well. For example, *airplane* class is visually similar to *ship* class since the backgrounds are similar. As a result, the distance between these two class centers is the minimum among other pair-wise distances. In a similar manner, *horse* class center is also closer to *dog* class since the appearances of some samples of these classes are quite similar. Overall, these computed distances verify that the proposed methodology yields semantically related CNN features.

3.1. Closed Set Recognition Experiments

We conducted experiments to assess the closed set recognition accuracy of the proposed method. To this end, we tested the proposed method on three datasets: Mnist, Cifar10, and Cifar100 datasets. We compared our results to the methods that maximize the margin in Euclidean or angular spaces. For all methods we used the same ResNet-18 architecture as backbone. Therefore, all results are directly comparable. The hypersphere radius is set to 64. The accuracies are given in Table 1. As seen in the results, our proposed method achieves the best accuracies in all cases, but the performance improvement is significant only for the Cifar100 dataset.

Figure 3: The distance matrix computed by using the centers of the training classes. The hypersphere centers are placed on the hypersphere based on appearances. Most of the time, semantically related classes lie close to each other, e.g., semantically related cat-dog classes or automobile-truck classes. In some cases, visually similar classes such as airplane-ship and horse-dog class centers are close to each other.

Table 1: Classification accuracies (%) on closed set recognition datasets.

Methods	**Mnist**	**Cifar10**	**Cifar100**
DUDC-HS (Ours)	**99.7**	**95.1**	**77.4**
Softmax	99.4	94.4	75.3
Center Loss	**99.7**	94.2	76.1
ArcFace	**99.7**	94.8	75.7
CosFace	**99.7**	95.0	75.8
SphereFace	**99.7**	94.7	75.1

3.2. Open Set Recognition Experiments

The datasets are split into *known* and *unknown* classes in open set recognition settings. By following the standard settings, we split the the datasets into known and unknown classes five times, trained our classifiers and computed the accuracies. The final accuracies are

Table 2: AUC Scores (%) of open set recognition methods on tested datasets ($n.r.$ stands for not reported).

Methods	Mnist	Cifar10	SVHN	Cifar+10	Cifar+50	TinyImageNet
DUDC-HS (Ours)	$\mathbf{99.6} \pm 0.1$	90.6 ± 0.5	95.2 ± 0.3	$\mathbf{97.5} \pm 0.3$	$\mathbf{96.2} \pm 0.4$	$\mathbf{81.4} \pm 0.7$
Softmax	97.8 ± 0.2	67.7 ± 3.2	88.6 ± 0.6	$81.6 \pm n.r.$	$80.5 \pm \pm n.r.$	$57.7 \pm n.r.$
OpenMax	98.1 ± 0.2	69.5 ± 3.2	89.4 ± 0.8	$81.7 \pm n.r.$	$79.6 \pm n.r.$	$57.6 \pm n.r.$
G-OpenMax	98.4 ± 0.1	67.5 ± 3.5	89.6 ± 0.6	$82.7 \pm n.r.$	$81.9 \pm n.r.$	$58.0 \pm n.r.$
C2AE	98.9 ± 0.2	89.5 ± 0.9	92.2 ± 0.9	95.5 ± 0.6	93.7 ± 0.4	74.8 ± 0.5
CAC	99.1 ± 0.5	80.1 ± 3.0	94.1 ± 0.7	87.7 ± 1.2	87.0 ± 0.0	76.0 ± 1.5
CPN	99.0 ± 0.2	82.8 ± 2.1	92.6 ± 0.6	$88.1 \pm n.r.$	$87.9 \pm n.r.$	$63.9 \pm n.r.$
OSRCI	98.8 ± 0.1	69.9 ± 2.9	91.0 ± 0.6	$83.8 \pm n.r.$	$82.7 \pm -$	$58.6 \pm n.r.$
CROSR	$99.1 \pm n.r.$	$88.3 \pm n.r.$	$89.9 \pm n.r.$	$91.2 \pm n.r.$	$90.5 \pm n.r.$	$58.9 \pm n.r.$
RPL	98.9 ± 0.1	82.7 ± 1.4	93.4 ± 0.5	84.2 ± 1.0	83.2 ± 0.7	68.8 ± 1.4
GDFRs	$n.r.$	83.1 ± 3.9	$\mathbf{95.5} \pm 1.8$	92.8 ± 0.2	92.6 ± 0.0	64.7 ± 1.2
Objecttosphere	$n.r.$	$\mathbf{94.2} \pm n.r.$	$91.4 \pm n.r.$	$94.5 \pm n.r.$	$94.4 \pm n.r.$	$75.5 \pm n.r.$

obtained by averaging the accuracies obtained in each trial. The details of the each dataset are given below:

3.2.1. Datasets

Mnist, Cifar10, SVHN: These datasets are split randomly into six known and four unknown classes by using the common testing setting. The 80 Million Tiny Images dataset (Torralba et al., 2008) is used as the background class.

Cifar+10, Cifar+50: For Cifar+N experiments, four randomly chosen classes from the Cifar10 dataset are used for training, and N non-overlapping classes chosen from the Cifar100 dataset are used as unknown classes as in (Yang et al., 2020; Miller et al., 2021; Yoshihashi et al., 2019; Chen et al., 2020). The 80 Million Tiny Images dataset (Torralba et al., 2008) is used as the background class.

TinyImageNet: For TinyImageNet (Russakovsky et al., 2015) experiments, twenty classes are randomly chosen as known classes and one hundred eighty classes as unknown classes by following the standard setting. The 80 Million Tiny Images dataset (Torralba et al., 2008) is used as the background class.

3.2.2. Results

The main goal of open set recognition is to detect and reject the samples that come from the novel classes. The performance of open set recognition is often measured using the Area Under the ROC curve (AUC) scores. Additionally, the closed set accuracy is also reported to evaluate classification performance on known data by disregarding unknown samples, as demonstrated in previous works such as (Yang et al., 2020) and (Neal et al., 2018). Our proposed method, Deep Uniformly Distributed Centers (DUDC-HS), is compared against to other state-of-the-art open set recognition methods including C2AE (Oza and Patel, 2019), Softmax, OpenMax (Scheirer et al., 2013), OSRCI (Neal et al., 2018), CAC (Miller et al., 2021), RPL (Chen et al., 2020), CROSR (Yoshihashi et al., 2019), ROSR (Yoshihashi et al., 2019), Generative-Discriminative Feature Representations (GDFRs) (Perera et al., 2020), and Objecttosphere (Dhamija et al., 2018) methods. Except for the TinyImageNet dataset, we employed the identical network backbone as in (Neal et al., 2018) for all datasets. To achieve higher accuracies for the TinyImageNet dataset, we utilized a deeper Resnet-50 architecture. The hypersphere radius is set to $u = 64$ as in the ArcFace method. The proposed method returned accuracies that are directly comparable to those reported in (Neal et al., 2018) for most of the tested datasets, as the network weights were randomly initialized during the training stage. AUC scores were summarized in Table 2, which showed that the proposed method achieved the best accuracies across all datasets except for the Cifar10 and SVHN. Notably, there were significant performance differences observed for the Cifar+10, Cifar+50, and TinyImageNet datasets. Closed set accuracies for open set recognition methods were reported in Table 3, where the proposed method achieved the best accuracies among the tested methods, with the exception of the Mnist and SVHN datasets. Obtaining typically the best accuracies in terms of AUC scores and closed set accuracies indicates that our proposed method can easily identify and reject the novel class samples and correctly classify the known class samples as expected.

Table 3: Closed Set accuracies (%) of open set recognition methods on tested datasets.

Methods	**Mnist**	**Cifar10**	**SVHN**	**Cifar+10**	**Cifar+50**	**TinyImageNet**
DUDC-HS (Ours)	99.6 ± 0.1	$\mathbf{94.3} \pm 1.4$	96.1 ± 0.6	$\mathbf{94.6} \pm 0.8$	$\mathbf{95.0} \pm 0.3$	$\mathbf{83.3} \pm 1.3$
Softmax	99.5 ± 0.2	80.1 ± 3.2	94.7 ± 0.6	$n.r.$	$n.r.$	$n.r.$
OpenMax	99.5 ± 0.2	80.1 ± 3.2	94.7 ± 0.6	$n.r.$	$n.r.$	$n.r.$
G-OpenMax	99.6 ± 0.1	81.6 ± 3.5	94.8 ± 0.8	$n.r.$	$n.r.$	$n.r.$
CPN	$\mathbf{99.7} \pm 0.1$	92.9 ± 1.2	$\mathbf{96.7} \pm 0.4$	$n.r.$	$n.r.$	$n.r.$
OSRCI	99.6 ± 0.1	82.1 ± 2.9	95.1 ± 0.6	$n.r.$	$n.r.$	$n.r.$
CROSR	99.2 ± 0.1	93.0 ± 2.5	94.5 ± 0.5	$n.r.$	$n.r.$	$n.r.$

4. Conclusion

In this study, we proposed a novel method that uses uniformly distributed centers on a hypersphere for open set recognition. In the proposed method, each class is represented by a center, and the centers are learned from training data in an end-to-end manner. The centers are enforced to lie on the boundary of a hypersphere with radius u, and the center of the hypersphere is located at the origin. The proposed loss function enforces the class specific samples to lie closer their corresponding centers. For open set recognition settings, there is an additional loss term that pushes the background samples away from the known class centers. The test samples are assigned to classes based on the Euclidean distances from the learned class centers. Experimental results verify that the proposed method generally achieves the best accuracies for both open and closed set recognition settings. The performance difference is significant especially on open set recognition datasets.

Acknowledgments

This work was supported by the Scientific and Technological Research Council of Turkey (TUBİTAK) under Grant number EEEAG-121E390.

References

Q. Bytyqi, N. Wolpert, E. Schomer, and U. Schwanecke. Prototype softmax cross entropy: a new perspective on softmax cross entropy. In *Scandinavian Conference on Image Analysis (SCIA)*, 2023.

H. Cevikalp and H. Saribas. Deep simplex classifier for maximizing the margin in both euclidean and angular spaces. In *Scandinavian Conference on Image Analysis (SCIA)*, 2023.

H. Cevikalp, B. Uzun, O. Kopuklu, and G. Ozturk. Deep compact polyhedral conic classifier for open and closed set recognition. *Pattern Recognition*, 119(108080):1–12, 2021.

H. Cevikalp, B. Uzun, Y. Salk, H. Saribas, and O. Kopuklu. From anomaly detection to open set recognition: Bridging the gap. *Pattern Recognition*, 138:109385, 2023.

G. Chen, L. Qiao, Y. Shi, P. Peng, J. Li, T. Huang, S. Pu, and Y. Tian. Learning open set network with discriminative reciprocal points. In *ECCV*, 2020.

W. Cho and J. Choo. Towards accurate open-set recognition via background-class regularization. In *European Conference on Computer Vision*, 2022.

J. Deng, J. Guo, N. Xue, and S. Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *IEEE Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

A. R. Dhamija, M. Gunther, and T. E. Boult. Reducing network agnostophobia. In *Neural Information Processing Systems (NeurIPS)*, 2018.

Yueqi Duan, Jiwen Lu, and Jie Zhou. Uniformface: Learning deep equidistributed representations for face recognition. In *IEEE Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

Chuanxing Geng, Sheng-Jun Huang, and Songcan Chen. Recent advances in open set recognition: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(10):3614–3631, 2021.

F. Graf, C. D. Hofer, M. Niethammer, and R. Kwitt. Dissecting supervised constrastive learning. In *International Conference on Machine Learning (ICML)*, 2021.

P. Hall, J. S. Marron, and A. Neeman. Geometric representation of high dimension, low sample size data. *Journal of the Royal Statistical Society Series B*, 67:427–444, 2005.

L. O. Jimenez and D. A. Landgrebe. Supervised classification in high dimensional space: geometrical, statistical, and asymptotical properties of multivariate data. *IEEE Transactions on Systems, Man, and Cybernetics-Part C: Applications and Reviews*, 28(1):39–54, 1998.

T. Li, P. Cao, Y. Yuan, L. Fan, Y. Yang, R. Feris, P. Indyk, and D. Katabi. Targeted supervised contrastive learning for long-tailed recognition. In *IEEE Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

Rongmei Lin, Weiyang Liu, Zhen Liu, Chen Feng, Zhiding Yu, James M. Rehg, Li Xiong, and Le Song. Regularizing neural networks via minimizing hyperspherical energy. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6916–6925, 2020.

W. Liu, Y. Wen, Z. Yu, and M. Yang. Large-margin softmax loss for convolutional neural networks. In *International Conference on Machine Learning (ICML)*, 2016.

W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song. Sphereface: Deep hypersphere embedding for face recognition. In *IEEE Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

W. Liu, R. Lin, Z. Liu, L. Liu, Z. Yu, B. Dai, and L. Song. Learning towards minimum hyperspherical energy. In *Neural Information Processing Systems (NeurIPS)*, 2018.

W. Liu, R. Lin, Z. Liu, L. Xiong, B. Scholkopf, and A. Weller. Learning with hyperspherical uniformity. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021.

W. Liu, Y. Wen, B. Raj, R. Singh, and A. Weller. Sphereface revived: Unifying hyperspherical face recognition. *IEEE Transactions on PAMI*, 45:2458–2474, 2023.

P. Mettes, E. van der Pol, and C. G. M. Snoek. Hyperspherical prototype networks. In *Neural Information Processing Systems (NeurIPS)*, 2019.

D. Miller, N. Sunderhauf, M. Milford, and F. Dayoub. Class anchor clustering: A loss for distance-based open set recognition. In *WACV*, 2021.

Lawrence Neal, Matthew Olson, Xiaoli Fern, Weng-Keen Wong, and Fuxin Li. Open set learning with counterfactual images. In *ECCV*, 2018.

Poojan Oza and Vishal M. Patel. C2ae: Class conditioned auto-encoder for open-set recognition. In *CVPR*, 2019.

P. Perera, V. I. Morariu, R. Jain, V. Manjunatha, C. Wigington, V. Ordonez, and V. M. Patel. Generative-discriminative feature representations for open-set recognition. In *CVPR*, 2020.

O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, and M. Bernstein. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115:201–252, 2015.

W. J. Scheirer, A. Rocha, A. Sapkota, and T. E. Boult. Towards open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35:1757–1772, 2013.

P. M. L. Tammes. On the origin of number and arrangement of the places of exit on the surface of pollen-grains. *Recueil des travaux botaniques neerlandais*, 27(1):1–84, 1930.

J. J. Thomson. On the structure of the atom: an investigation of the stability and periods of oscillation of a number of corpuscles arranged at equal intervals around the circumference of a circle; with application of the results to the theory of atomic structure. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 7(39):237–265, 1904.

Antonio Torralba, Rob Fergus, and William T. Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(11):1958–1970, 2008.

Bedirhan Uzun, Hakan Cevikalp, and Hasan Saribas. Deep discriminative feature models (ddfms) for set based face recognition and distance metric learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5):5594–5608, 2023.

H. Wang, Y. Wang Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, and W. Liu. Cosface: Large margin cosine loss for deep face recognition. In *IEEE Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

Hong-Ming Yang, Xu-Yao Zhang, Fei Yin, Qing Yang, and Cheng-Lin Liu. Convolutional prototype network for open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2020. doi: 10.1109/TPAMI.2020.3045079.

R. Yoshihashi, W. Shao, R. Kawakami, S. You, M. Iida, and T. Naemura. Classification-reconstruction learning for open-set recognition. In *CVPR*, 2019.