

Representative-Discriminative Learning for Open-set Land Cover Classification of Satellite Imagery

Razieh Kaviani Baghbaderani¹, Ying Qu^{1*}, Hairong Qi¹, and Craig Stutts²

¹ The University of Tennessee, Knoxville, TN, USA
{rkavian1,yqu3,hqi}@utk.edu

² Applied Research Associates, Raleigh, NC, USA
cstutts@ara.com

Abstract. Land cover classification of satellite imagery is an important step toward analyzing the Earth’s surface. Existing models assume a closed-set setting where both the training and testing classes belong to the same label set. However, due to the unique characteristics of satellite imagery with extremely vast area of versatile cover materials, the training data are bound to be non-representative. In this paper, we study the problem of open-set land cover classification that identifies the samples belonging to unknown classes during testing, while maintaining performance on known classes. Although inherently a classification problem, both representative and discriminative aspects of data need to be exploited in order to better distinguish unknown classes from known. We propose a representative-discriminative open-set recognition (RDOSR) framework, which 1) projects data from the raw image space to the embedding feature space that facilitates differentiating similar classes, and further 2) enhances both the representative and discriminative capacity through transformation to a so-called abundance space. Experiments on multiple satellite benchmarks demonstrate effectiveness of the proposed method. We also show the generality of the proposed approach by achieving promising results on open-set classification tasks using RGB images.

Keywords: Hyperspectral image classification, open-set recognition

1 Introduction

Recent advancements in computer vision, especially the advent of Convolutional Neural Networks (CNN), have significantly improved the performance of image classification [1–3], detection [4, 5], and segmentation [6, 7] tasks, enabling their deployment in many different fields. One of such field of applications is satellite image analysis that includes resource management, urban development planning, and climate control. Land cover classification or material classification is

* Corresponding author




Fig. 1: Open-set land-cover classification: Data samples corresponding to ground truth categories are from the known class set ($\mathbf{K}$). It is likely that some categories are not known during training and will be encountered at testing, *i.e.*, samples from unknown class set ($\mathbf{U}$). The goal is to identify pixels coming from ($\mathbf{U}$), while correctly classify any pixel belonging to ($\mathbf{K}$). From left to right, a satellite image from the Pavia University dataset [8] showing unknown materials surfaces with yellow bounding boxes, the ground truth labels, and visualization of the feature space for both known and unknown classes using tSNE [9].

one of the building blocks of satellite image analysis, providing essential inputs to a series of subsequent tasks including object segmentation, 3D reconstruction and modeling, as well as texture mapping. Supervised land cover classification involves categorization of multispectral or hyperspectral image pixels into predefined material classes, e.g., asphalt, tree, concrete, water, metal, soil, etc. Note that both multispectral and hyperspectral images (MSI and HSI) try to provide additional spectral information, beyond the visible spectra, to reveal extra details and compensate for the coarse spatial resolution of these images.

Although inherently a classification problem, material classification in satellite imagery faces a unique challenge: the vast area covered by the satellite imagery makes the task of generating representative training samples almost impossible, as there are a large variety of materials existed on the Earth’s surface, especially those not well-exploited regions. Therefore, one of the most essential capabilities of land cover classification is to be able to automatically identify which test image and which area or pixel location of the image, has a higher probability of hosting new classes of materials. This would provide essential guideline to human operators in collecting training samples for the new classes.

The vast majority of existing works for land cover classification have been done under the “static closed world” assumption, meaning that the training and testing sets are drawn from the same label set. As a result, a system observing any unknown class is forced to misclassify it as one of the known classes, thus, weakening the recognition performance. A more realistic scenario is to work in a non-stationary and open environment that not all categories are known *a priori*

and testing samples from unseen classes can emerge unexpectedly. Recognition of known and unknown pixels in a given image and correctly classifying known pixels is defined as “open-set land cover classification”. Fig. 1 explains this process using a real-world satellite image.

In this paper, we present a multi-tasking representative-discriminative open-set recognition (RDOSR) framework to address the challenging land cover classification problem, where both the representative and discriminative aspects of data are exploited in order to best characterize the differences between known and unknown classes. We propose the representative and discriminative learning among three spaces, as shown in Fig. 2, including 1) the transformation from the raw image space to the embedding feature space, and 2) the transformation from the embedding feature space to a so-called abundance space. See Supplement A for an illustration of the effect in different spaces.




Fig. 2: Representative-discriminative learning through the transformation among 3 spaces: the raw image space, the embedding space, and the abundance space.

The contributions of this paper are thus summarized as follows: First, unlike other open-set recognition methods applied on the raw image space directly, we propose to first learn a classification network that would transform from the raw image space to an embedding feature space such that a more discernible input is fed into the subsequent open-set learning network. Second, we propose to use the so-called Dirichlet-net to transform data from the embedding feature space to the abundance space. Due to the resolution issue, each pixel in a satellite image covers a large area with more than one constituent material, resulting in “mixed pixel”. The mixtures are generally assumed to be a linear combination of a few spectral bases, with the corresponding mixing coefficients (or abundances). This way, instead of looking at the mixed pixel, we study the mixing coefficients of each spectral basis in making up the mixture. Thus the abundance space provides a finer-scale representation. Third, to the best of our knowledge, this work is the first attempt to address the critical open-set land cover classification problem essential for analyzing the Earth’s surface. Fourth, while the proposed method was motivated by satellite imagery analysis, it is generalizable to RGB images and achieves promising results.

2 Related Work

Conventional Land Cover Classification. These methods mainly employ a traditional classifier on spectral information, with its discriminative power further enhanced through feature engineering algorithms such as minimum noise fraction (MNF) [10], independent component analysis (ICA) [11], morphological profiles [12], and spectral unmixing [13–15]. The advent of deep learning has enabled the extraction of hierarchical features automatically and achieved unprecedented performance. [16, 17] applied a 1D-CNN framework in the spectral domain to take into account the correlation between adjacent spectral bands. Several works use a patch surrounding the desired pixels by adopting a 2D-CNN structure [18, 19] to incorporate the spatial correlation as well. More recently, integration of both spectral and spatial domains using 3D-CNN structures have been employed to further improve the classification accuracy [20, 21].

Although each approach has its own merit, all the existing land cover classification approaches work under the closed-set assumption where the training and testing sets share the same label set.

Open-set Recognition. Open-set recognition has gained considerable attention due to its handling of unknown class samples based on incomplete knowledge of the data during model training. Early studies are based on traditional classification models including Nearest Neighbor, Support Vector Machine (SVMs), Sparse Representation, etc. The open-set version of Nearest Neighbor was developed based on the distance of the testing samples to the known samples [22]. The SVM-based approaches employed different regularization terms or kernels to detect unknown samples [23, 24]. In [25], the residuals from the Sparse Representation-based Classification (SRC) algorithm were used as the score for unknown class detection.

In the context of deep networks, [26] employed a statistical model to calibrate the SoftMax scores and produced an alternative layer, called OpenMax. [27] improved upon the OpenMax layer approach by maximizing the inter-class distance and minimizing the intraclass distance on the penultimate layer. The work of [28] proposed a k-sigmoid activation-based loss function for training a neural network to be able to find an operating threshold on the final activation layer. [29] incorporated the latent representation for reconstruction along with the discriminative features obtained from a classification model to enhance the feature vector used for open-set detection. Unlike previous methods, [30] utilized reconstruction error obtained from a multi-task learning framework as a detection score. Recently, [31] proposed using self-supervision and augmented the input image to learn richer features to improve separation between classes.

More recent works try to simulate open-set classes in order to provide explicit probability estimation over unknown classes. Ge et al. [32] extended OpenMax [26] by synthesizing the unknown samples using a Generative Adversarial Network (GAN) based framework. Along the same line, [33] proposed the counterfactual image generation (OSRCI) framework which employs a GAN to generate samples placing between decision boundaries that can be treated as unknown examples. [34] proposed class-conditioned auto-encoder (C2AE) algorithm where




Fig. 3: An overview of the proposed framework: i) Closed-set embedding learning: the classifier F is trained on the spectral domain X to produce latent discriminative embedding z_F . ii) Representative-discriminative feature learning: the encoder E takes the embedding feature z_F and derives the representative features S using a Dirichlet-Net. The classifier C applied on S enhances the discriminative aspect of S , and the reconstruction error between the decoder output ($\hat{z}_F$) and input to encoder (z_F) enhances the representative aspect of S .

conditional reconstruction helps learning of both known and unknown score distributions.

It should be noted that there are related problems in the literature including outlier detection [35, 36] and anomaly detection [37, 38] that have some overlap with open-set recognition and can be treated as a relaxed version of open-set recognition. These problems assume the availability of one abnormal class during training. However, general open-set recognition problems usually do not provide information about the type or the number of the unknown classes in advance.

3 Proposed Approach

We propose a representative-discriminative open-set recognition (RDOSR) structure, as shown in Fig. 3. The network mainly consists of two components, 1) a closed-set embedding component to project the data from the original image domain to the embedding domain, such that different classes with similar spectral characteristics are more distinguishable, and 2) a multi-task representative-discriminative learning component to learn a better representation scheme at a finer scale in the abundance space, such that unknown classes can be better differentiated from known classes.

3.1 Network Architecture

One challenging issue of the open-set satellite land cover classification problem is that different classes may possess similar spectral characteristics. Thus, it is

likely that an unknown class, whose spectral profile is close to that of a known class, may be misclassified as the known class. To address this issue, instead of detecting unknown classes on the image domain, we detect them on the embedding domain projected by a closed-set embedding layer, as shown in Fig. 3.i. The closed-set embedding layer increases the discriminative power of network to a large extent, such that unknown classes can be better recognized even if their spectra are similar to those of the known classes. The weights of the closed-set embedding layer are trained with a classifier F , which is further elaborated in Sec. 3.2.

To recognize unknown classes in the embedding domain, we propose a multi-task representative-discriminative feature learning framework to boost both the representative and discriminative power of the extracted feature vector, such that it is more informative and effective to recognize unknown samples. This is shown in Fig. 3.ii. The network consists of an encoder-decoder architecture with the representative features S extracted using a sparse Dirichlet encoder E , and a decoder formed by the bases shared among known classes. A classifier C applied on S is also included to further increase its discriminative capability. In this way, the data from unknown classes fed into the network would produce higher reconstruction error, thus can be detected accordingly. The details of network design are further elaborated in Sec. 3.3.

3.2 Closed-set Embedding Learning

Given the set of sample pixels $X_k = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_k}\}$ from the known classes with each pixel, $\mathbf{x}_i$, being a high-dimensional vector recording the reflectance readings of different spectral bands in the hyperspectral image, the corresponding labels are denoted with $Y_k = \{y_1, y_2, \dots, y_{N_k}\}$, where N_k is the number of known pixels and $\forall y_i \in \{1, 2, \dots, L\}$, where L is the number of known classes. To distinguish classes with similar spectral distributions, we project the input data X_k from the image domain to the embedding domain Z_F . The projection is learned through a classifier, F , with parameters Θ_F and the embedded features, $\mathbf{z}_F$ is forced to be discriminative through the cross-entropy loss,

$$\mathcal{L}_f(\Theta_F) = -\frac{1}{N_k} \sum_{i=1}^{N_k} y_i \log[F(\mathbf{x}_i)], \quad (1)$$

where y_i is a one-hot encoded label and $F(\mathbf{x}_i)$ denotes the vector carrying the predicted probability score of the i^{th} known sample. Such vector is generated by applying a softmax function on the features $\mathbf{z}_F$ in the embedding domain.

This general structure is sufficient for a common classification problem where the encountered classes are known. However, our goal is to increase the discriminative power of the features from classes with similar spectral characteristics. Therefore, we further increase the discriminative capacity of the embedded features with the l_1 -norm sparse constraint defined by

$$\mathcal{L}_z(\Theta_F) = \frac{1}{N_k} \sum_{i=1}^{N_k} \|\mathbf{z}_{\mathbf{F}_i}\|, \quad (2)$$

where $\mathbf{z}_{\mathbf{F}_i}$ is the embedding feature vector learned by the classifier F . With such constraint, the embedded features from samples of different classes are more discriminative, even if their spectra are similar in the image domain.

3.3 Multi-Task Representative-Discriminative Feature Learning

With the proposed closed-set embedding layer, the samples are projected from the image domain to the embedding domain possessing more distinguishable features. In order to better identify unknown samples, both discriminative and representative nature of the samples need to be exploited. Previous approaches [30, 34] usually train a general auto-encoder to reconstruct samples from known classes. When the samples from unknown classes are fed into the network, the reconstruction error is expected to be larger than that of the known classes in the ideal case, since the network weights are optimized during the training procedure using known samples. However, the challenge lies in the scenario when the unknown classes especially the ones close to the known classes may potentially contribute to small reconstruction error too which would lead to failure in detection.

In this work, instead of adopting a general-purpose auto-encoder, we propose a multi-task representative-discriminative feature learning framework to improve the detection accuracy. The purpose of this network is to decrease the reconstruction error of the known classes while increasing the reconstruction error of unknown classes intensively.

Due to the resolution issue, each pixel in a satellite image usually covers a large geographical area or footprint (e.g., $30 \times 30\text{m}$ for Landsat-8), resulting in the so-called “mixed pixel” (i.e., each pixel tends to cover more than one constituent materials). These mixtures are generally assumed to be a linear combination of a few spectral bases with the corresponding mixing coefficients (or abundance). The proposed method is designed based on this assumption as shown in Eq. 3. Assume that the feature vector of a sample from known classes $\mathbf{z}_{\mathbf{F}}$ is a linear combination of a few bases B , and such bases are shared among features of the known classes. Thus, each sample of known classes can be decomposed with

$$\mathbf{z}_{\mathbf{F}} = \mathbf{s}B, \quad (3)$$

where $\mathbf{s}$ denotes the proportional coefficients of the shared bases, which serves as a form of “representation” of the embedding feature, that we refer to as the abundance. The abundance vector, or representation, should satisfy two physical constraints, i.e., non-negative and sum-to-one. The samples from unknown classes are also able to be decomposed by Eq. 3 using shared bases of the known classes, B . However, since B does not include the bases of the unknown classes, the distribution of its representations $\mathbf{s}$ should deviate from that of the known




Fig. 4: The flowchart of the multi-task representative-discriminative feature learning framework.

classes. Therefore, we design a network following the model of Eq. 3, which enforces $\mathbf{s}$ from known classes to follow a certain distribution. And if the network can extract $\mathbf{s}$ from unknown classes with similar distributions, then we expect they have high reconstruction errors.

The flowchart of the proposed multi-task representative-discriminative feature learning is detailed in Fig. 4. The network performs both the reconstruction task and the classification task. The reconstruction branch consists of a sparse Dirichlet-based encoder E with weights Θ_E and a decoder D with weights Θ_D . The encoder and decoder can be defined by the functions $E : Z_F \rightarrow S$ and $D : S \rightarrow Z_F$, respectively, where Z_F is the embedding space obtained by the closed-set classifier F , and S is the abundance space of latent representations projected by the encoder E . The representations $\mathbf{s}$ in the latent space S is enforced to follow a Dirichlet distribution. And a sparse constraint is introduced to enhance the representativeness of $\mathbf{s}$. More details and justifications are provided below. In addition, S is also enforced to be discriminative by the classifier C , which can be defined by the function $C : S \rightarrow Y$ with weights Θ_C , where Y is the space of known labels.

Representative Feature Learning with Reconstruction The reconstruction branch is constructed according to Eq. 3, where the shared bases are embedded in the decoder D of the network and the corresponding representation $\mathbf{s}$ is extracted with the encoder E . Since $\mathbf{s}$ denotes the proportional coefficients of the bases, we enforce it to follow a Dirichlet distribution meeting the non-negative and sum-to-one physical constraints. Following the work of [39–41], we adopt the stick-breaking structure in the encoder to enforce the representations $\mathbf{s}$ to follow the Dirichlet distribution.

In the stick-breaking structure, a single element s_j in $\mathbf{s}$ can be expressed by

$$s_j = \begin{cases} v_1 & \text{for } j = 1 \\ v_j \prod_{o < j} (1 - v_o) & \text{for } j > 1, \end{cases} \quad (4)$$

where v_j is drawn from a Kumaraswamy distribution, i.e., $v_j \sim \text{Kuma}(u, 1, \beta)$ as shown in Eq. (5),

$$v_o \sim (1 - (1 - u^{\frac{1}{\beta}})). \quad (5)$$

Then, there are two parameters used to extract representations $\mathbf{s}$, i.e., u and β , both of which are hidden layers in the encoder of the network. A softplus activation function is adopted on the layer β due to its non-negative property, and a sigmoid is used to map u into the $(0, 1)$ range at the layer u . More details of the stick-breaking structure can be found in [40] and [41].

In addition, the entropy function [42] is adopted to reinforce the sparsity of the representation layer. Let $\hat{s}_j = \frac{|s_j|}{\|\mathbf{s}\|}$, for each pixel, the entropy function is defined as,

$$\mathcal{H}_s(\Theta_E) = - \sum_{j=1}^c \hat{s}_j \log \hat{s}_j. \quad (6)$$

where c is the dimension of the representation $\mathbf{s}$. The reconstruction loss $\mathcal{L}_r$ is adopted to reduce the reconstruction error of the known classes. It is defined by,

$$\mathcal{L}_r(\{\Theta_E, \Theta_D\}) = \frac{1}{N_k} \sum_{i=1}^{N_k} \|\mathbf{z}_{\mathbf{F}_i} - \hat{\mathbf{z}}_{\mathbf{F}}\|_2, \quad (7)$$

where $\mathbf{z}_{\mathbf{F}_i}$ is the embedding feature vector fed into the encoder E , and $\hat{\mathbf{z}}_{\mathbf{F}}$ is the reconstructed $\mathbf{z}_{\mathbf{F}_i}$ obtained from the decoder D .

Discriminative Feature Learning with Classification Branch. To further increase the discriminative capacity of the representations, a classifier is adopted on the representations $\mathbf{s}$ with the classification loss $\mathcal{L}_c$ defined as,

$$\mathcal{L}_c(\{\Theta_E, \Theta_C\}) = \frac{1}{N_k} \sum_{i=1}^{N_k} y_i \log[E(\mathbf{z}_{\mathbf{F}_i})], \quad (8)$$

where y_i is the ground truth label and $E(\mathbf{z}_{\mathbf{F}_i})$ denotes the representative feature vector of the i^{th} known sample. Note that the weights of both the reconstruction branch and classifier C are updated together, such that the learned representations can be both representative and discriminative.

3.4 Training Procedure and Network Settings

We first learn the embedding projection by optimizing the weights Θ_F of the classifier F with the loss function,

$$\min_{\Theta_F} \lambda_f \mathcal{L}_f + \lambda_z \mathcal{L}_z, \quad (9)$$

where λ_f and λ_z are two parameters to balance the trade-off between the cross-entropy loss and the sparsity loss.

Then, having the learned embedding layer, the multi-task representative-discriminative feature learning network is trained to minimize both the reconstruction loss and the classification error of the known classes with loss function,

$$\min_{\Theta_E, \Theta_D, \Theta_C} \lambda_r \mathcal{L}_r + \lambda_s \mathcal{H}_s + \lambda_c \mathcal{L}_c, \quad (10)$$

where λ_r , λ_s , and λ_c are parameters balancing the trade-off between the reconstruction loss, the sparsity loss, and the classification loss.

The structures of the four networks, F , E , D , and C are listed in Table 1.

Table 1: The nodes in the proposed network

Networks	F	E	D	C
nodes	[512,1024,512,32, L]	[3,3,3,3, 10]	[10,10, L]	[L]

4 Experiments and Results

In this section, the effectiveness of the proposed RDOSR method is evaluated on several widely used benchmark hyperspectral image datasets. In addition, we demonstrate the generalization capacity of the proposed approach on RGB image datasets. Furthermore, the contribution from each component of the proposed framework is analyzed through ablation study.

4.1 Implementation Details

We train the network, described in Sec. 3.1, using an Adam optimizer [43], with a learning rate of 10^{-3} . The classifier F and the joint structure of encoder-decoder-classifier (E - D - C) are trained separately for a total number of 15K epochs. However, the other methods which do not have two separate components were trained for 6K epochs.

For training the classifier F , λ_f and λ_z are set equal to 1 and 0.1, respectively. The weights for the reconstruction λ_r , sparsity λ_s , and classification λ_c losses in training the E - D - C structure are set equal to 0.5, 10^{-3} , and 0.5, respectively. The sparsity weight λ_s is decayed with a weight decay of 0.9977. The classifier F is trained until its accuracy reaches 0.9988. It should be noted that all input data of the datasets are normalized to their mean value and unit variance. In addition, the feature vector obtained from the classifier F is divided by 10 to avoid divergence.

One of the factors that affects the performance of the open-set recognition algorithm is Openness [44] of the problem, defined as,

$$Openness = 1 - \sqrt{\frac{2 \times N_{train}}{N_{test} + N_{target}}}, \quad (11)$$

where N_{train} , N_{test} and N_{target} are the number of classes known during training, the number of classes given during testing, and the number of classes that need to be recognized correctly during testing phase, respectively. In the experiments, classes of each dataset is partitioned into known and unknown sets according to the Openness.

The code is written in TensorFlow, and all the experiments are performed on a desktop computer having GeForce GPU of 10 GB Memory. The code is available at <https://github.com/raziehkaviani/rdosr>.

4.2 Metrics

To compare performance of different methods, there are several metrics including overall accuracy or F-score on a combination of known and unknown classes, and Receiver Operating Characteristic (ROC). The first two metrics do not characterize the performance of the model well due to their sensitivity not only to the performance of the model in classifying the known classes, but also an arbitrary operating threshold for detecting unknown samples.

On the other hand, the ROC curve would illustrate the ability of a binary classification system (here, known vs. unknown detection) as a discrimination threshold is varied from the minimum to the maximum value of the given detection measure (here, reconstruction error). Thus, it provides a measure free from calibration. To have a quantitative comparison, the area under the ROC (AUC) is computed in the experiments.

4.3 Open-set Recognition for Hyperspectral Data

The experiments are conducted on three hyperspectral image datasets:

Pavia University (PU) and Pavia Center (PC). Both PU and PC datasets were gathered over Northern Italy in 2011 by the Reflective Optics Systems Imaging Spectrometer which has a resolution of 1.3 m. The dimension of the PU dataset is 1096×715 pixels with 103 spectral bands, ranging from 430 to 860 nm. The PC dataset has 610×340 pixels with 102 spectral bands. The PU and PC datasets both include nine land cover categories.

Indian Pines (IN). The IN dataset was collected over Northwest Indiana in 1912 by Airborne Visible/Infrared Imaging Spectrometer (AVIRIS). It has a dimension of 145×145 with a resolution of 20 m by pixel and 200 spectral bands. Its ground truth consists of 16 land cover classes.

We compare the performance of the proposed approach with three methods described in the following:

SoftMax: In a neural network classifier, a common confidence-based approach to detect open-set examples is thresholding the SoftMax scores. We use network structure of the classifier F without considering sparsity constraint.

OpenMax [26]: This approach calibrates the SoftMax scores in a classifier and augments them with a $N_k + 1$ class for an unknown category. The replaced SoftMax layer with an OpenMax layer is used for open-set recognition. We adopt

the classifier mentioned earlier in the SoftMax method, and use the Weibull fitting approach with parameter *Weibull tail size* = 10 to generate OpenMax layer values.

AE+CLS: Fully connected version of MLOS [30] which utilizes a multi-task learning framework, composed of a classifier and a decoder with a shared feature extraction part, to detect open-set examples. To have a fair comparison with our approach, the encoder, decoder, and classifier are designed as our *E* (without the Dirichlet-Net), *D*, *C*, and trained with $\mathcal{L}_r$ and $\mathcal{L}_c$ loss, with weights of 0.5.

First, each of the L classes is assumed to be unknown which equates to an openness of 2.99%, 2.99%, and 1.63% for PU, PC, IN, respectively. The AUC values corresponding to choosing each of the L labels as unknown are averaged and reported for each method in Table 2. It can be observed that the proposed method outperforms other methods on all three datasets. See Supplement B for detailed comparison on the PU dataset. The minor improvement on the PC dataset can be justified by its differentiated spectrum of different classes which diminishes the effect of the classifier *F*.

Second, for the Openness equal to 6.46%, the ROC curves of different methods for the PU and IN datasets are illustrated in Fig. 5. As seen from the results on both datasets, the *AE+CLS+Dirichlet* method which adopts the Dirichlet net to the *AE+CLS* framework and the proposed method lie above all other methods. It should be noticed that our proposed method is able to detect unknown classes with 60% and above 90% accuracy and almost zero false detection for the PU and PC datasets, respectively.

Third, the histograms of reconstruction error for both known and unknown sets with Openness= 2.99% are shown in Fig. 6. It can be observed that the reconstruction errors corresponding to the known set have small values. However, the unknown set produces larger error due to mismatches in terms of representative and discriminative features learned from the known classes examples.

4.4 Open-set Recognition for RGB Images

To show the generalization capacity of the proposed method, we evaluate the performance of the proposed approach on two RGB datasets and compare with

Table 2: Area under the ROC curve for open-set detection. Results are averaged over L partitioning of the selected dataset to $L - 1$ known and 1 unknown classes.

Method	PU	PC	IN
SoftMax	0.385	0.816	0.555
OpenMax [26]	0.441	0.884	0.415
AE+CLS [30]	0.586	0.757	0.669
AE+CLS+Dirichlet	0.714	0.927	0.681
RDOSR (Ours)	0.773	0.963	0.802




Fig. 5: Receiver Operating Curve curves for open-set recognition for PU and PC datasets, for $L = 7$ (openness=6.46%).




Fig. 6: Reconstruction error distribution of known and unknown classes using the proposed method for PU and PC datasets, for $L = 8$.

several state-of-the-art methods. For this purpose, the classifier F performing pixel-wise classification is substituted with a DenseNet structure which takes a 2D image as input.

Following the protocol in [33], we sample 4 known classes from CIFAR10 [45] to have Openness=13.39% and 20 known classes out of 200 categories of TinyImageNet [46] resulting in an Openness of 57.35%. Table 3 summarizes the results where the values other than the proposed RDOSR are taken from [31]. It can be observed that the proposed method has better performance over the compared methods, except for GDOSR [31], on CIFAR10. However, it achieves significant improvement on TinyImageNet. It may be due to the similarity between the classes in TinyImageNet which hinders detecting unknown samples in the image space while RDOSR addresses this issue by operating in the embedding space.

Table 3: Area under the ROC curve for Open-set recognition.

Method	CIFAR10	TinyImageNet
SoftMax	0.677	0.577
OpenMax [26]	0.695	0.576
OSRCI [33]	0.699	0.586
C2AE [34]	0.711	0.581
GDOSR [31]	0.807	0.608
RDOSR (Ours)	0.744	0.752

4.5 Ablation Study

Starting with a baseline, $AE+CLS$, each component is gradually added to the framework to show its effectiveness. The results corresponding to the ablation study are shown in Fig. 7. It can be seen that employing the baseline structure applied on the spectra domain has the worst performance. However, adding the Dirichlet-based network makes a major improvement due to applying physical constraints on the latent space learned by the encoder E .

Directly performing open-set recognition on an embedding space causes instability problem which is confirmed by a performance drop compared to the $AE+CLS+Dirichlet$ method. Our proposed method addresses the instability issue by adopting a sparsity constraint on the embedding feature vector $\mathbf{z}_F$. As seen from Fig. 7, our proposed method achieves the highest AUC value compared to three other baseline methods.




Fig. 7: Ablation study of the proposed method on PU dataset

5 Conclusions

We studied the challenging problem of open-set land cover recognition in satellite images. Although inherently a classification problem, both representative and discriminative features need to be learned in order to best characterize the difference between known and unknown classes. We presented the transformation among three spaces, that is, the original image space, the embedding feature space, and the abundance space, where features with both representative and discriminative capacity can be learned to maximize success rate. The proposed multi-tasking representative-discriminative learning structure was evaluated on three hyperspectral and two RGB image datasets and exhibited significant improvement over state-of-the-art open-set recognition algorithms.

References

1. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*. (2012) 1097–1105
2. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2016) 770–778
3. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
4. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2016) 779–788
5. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: *Advances in neural information processing systems*. (2015) 91–99
6. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2015) 3431–3440
7. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence* **40**(4) (2017) 834–848
8. <http://lesun.weebly.com/hyperspectral-data-set.html>.
9. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(Nov) (2008) 2579–2605
10. Green, A.A., Berman, M., Switzer, P., Craig, M.D.: A transformation for ordering multispectral data in terms of image quality with implications for noise removal. *IEEE Transactions on geoscience and remote sensing* **26**(1) (1988) 65–74
11. Comon, P.: Independent component analysis, a new concept? *Signal processing* **36**(3) (1994) 287–314
12. Plaza, A., Martinez, P., Plaza, J., Perez, R.: Dimensionality reduction and classification of hyperspectral image data using sequences of extended morphological transformations. *IEEE Transactions on Geoscience and remote sensing* **43**(3) (2005) 466–479
13. Luo, B., Chanussot, J.: Unsupervised classification of hyperspectral images by using linear unmixing algorithm. In: *2009 16th IEEE International Conference on Image Processing (ICIP)*, IEEE (2009) 2877–2880
14. Dópido, I., Zortea, M., Villa, A., Plaza, A., Gamba, P.: Unmixing prior to supervised classification of remotely sensed hyperspectral images. *IEEE Geoscience and Remote Sensing Letters* **8**(4) (2011) 760–764
15. Baghbaderani, R.K., Wang, F., Stutts, C., Qu, Y., Qi, H.: Hybrid spectral unmixing in land-cover classification. In: *International Geoscience and Remote Sensing Symposium (IGARSS)*. (2019)
16. Hu, W., Huang, Y., Wei, L., Zhang, F., Li, H.: Deep convolutional neural networks for hyperspectral image classification. *Journal of Sensors* **2015** (2015)
17. Song, Y., Zhang, Z., Baghbaderani, R.K., Wang, F., Stutts, C., Qi, H.: Land cover classification for satellite images through 1d cnn. In: *Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*. (2019)

18. Cao, X., Zhou, F., Xu, L., Meng, D., Xu, Z., Paisley, J.: Hyperspectral image classification with markov random fields and a convolutional neural network. *IEEE Transactions on Image Processing* **27**(5) (2018) 2354–2367
19. Sharma, V., Diba, A., Tuytelaars, T., Van Gool, L.: Hyperspectral cnn for image classification & band selection, with application to face recognition. Technical report KUL/ESAT/PSI/1604, KU Leuven, ESAT, Leuven, Belgium (2016)
20. Hamida, A.B., Benoit, A., Lambert, P., Amar, C.B.: 3-d deep learning approach for remote sensing image classification. *IEEE Trans. on geoscience and remote sensing* **56**(8) (2018) 4420–4434
21. Zhong, Z., Li, J., Luo, Z., Chapman, M.: Spectral-spatial residual network for hyperspectral image classification: A 3-d deep learning framework. *IEEE Transactions on Geoscience and Remote Sensing* **56**(2) (2017) 847–858
22. Júnior, P.R.M., De Souza, R.M., Werneck, R.d.O., Stein, B.V., Pazinato, D.V., de Almeida, W.R., Penatti, O.A., Torres, R.d.S., Rocha, A.: Nearest neighbors distance ratio open-set classifier. *Machine Learning* **106**(3) (2017) 359–386
23. Scheirer, W.J., Jain, L.P., Boulton, T.E.: Probability models for open set recognition. *IEEE transactions on pattern analysis and machine intelligence* **36**(11) (2014) 2317–2324
24. Júnior, P.R.M., Boulton, T.E., Wainer, J., Rocha, A.: Specialized support vector machines for open-set recognition. *arXiv preprint arXiv:1606.03802* (2016)
25. Zhang, H., Patel, V.M.: Sparse representation-based open set recognition. *IEEE transactions on pattern analysis and machine intelligence* **39**(8) (2016) 1690–1696
26. Bendale, A., Boulton, T.E.: Towards open set deep networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2016) 1563–1572
27. Hassen, M., Chan, P.K.: Learning a neural-network-based representation for open set recognition. *arXiv preprint arXiv:1802.04365* (2018)
28. Shu, L., Xu, H., Liu, B.: Doc: Deep open classification of text documents. *arXiv preprint arXiv:1709.08716* (2017)
29. Yoshihashi, R., Shao, W., Kawakami, R., You, S., Iida, M., Naemura, T.: Classification-reconstruction learning for open-set recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019) 4016–4025
30. Oza, P., Patel, V.M.: Deep cnn-based multi-task learning for open-set recognition. *arXiv preprint arXiv:1903.03161* (2019)
31. Perera, P., Morariu, V.I., Jain, R., Manjunatha, V., Wigington, C., Ordonez, V., Patel, V.M.: Generative-discriminative feature representations for open-set recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. (2020) 11814–11823
32. Ge, Z., Demyanov, S., Chen, Z., Garnavi, R.: Generative openmax for multi-class open set classification. *arXiv preprint arXiv:1707.07418* (2017)
33. 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)*. (2018) 613–628
34. 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*. (2019) 2307–2316
35. Xia, Y., Cao, X., Wen, F., Hua, G., Sun, J.: Learning discriminative reconstructions for unsupervised outlier removal. In: *Proceedings of the IEEE International Conference on Computer Vision*. (2015) 1511–1519

36. You, C., Robinson, D.P., Vidal, R.: Provable self-representation based outlier detection in a union of subspaces. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2017) 3395–3404
37. Chalapathy, R., Menon, A.K., Chawla, S.: Robust, deep and inductive anomaly detection. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases, Springer (2017) 36–51
38. Golan, I., El-Yaniv, R.: Deep anomaly detection using geometric transformations. In: Advances in Neural Information Processing Systems. (2018) 9758–9769
39. Sethuraman, J.: A constructive definition of dirichlet priors. *Statistica Sinica* (1994) 639–650
40. Nalisnick, E., Smyth, P.: Stick-breaking variational autoencoders. *International Conference on Learning Representations (ICLR)* (2017)
41. Qu, Y., Qi, H., Kwan, C.: Unsupervised sparse dirichlet-net for hyperspectral image super-resolution. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018) 2511–2520
42. Huang, S., Tran, T.D.: Sparse signal recovery via generalized entropy functions minimization. *IEEE Transactions on Signal Processing* **67**(5) (2018) 1322–1337
43. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
44. 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) (2012) 1757–1772
45. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. (2009)
46. Le, Y., Yang, X.: Tiny imagenet visual recognition challenge. *CS 231N* **7** (2015)