

CMIR-NET : A Deep Learning Based Model For Cross-Modal Retrieval In Remote Sensing

Ushasi Chaudhuri

Indian Institute of Technology Bombay
India
ushasi@iitb.ac.in

Avik Bhattacharya

Indian Institute of Technology Bombay
India
avikb@csre.iitb.ac.in

Biplab Banerjee

Indian Institute of Technology Bombay
India
getbiplab@gmail.com

Mihai Datcu

German Aerospace Center DLR
Germany
Mihai.Datcu@dlr.de

ABSTRACT

We address the problem of cross-modal information retrieval in the domain of remote sensing. In particular, we are interested in two application scenarios: i) cross-modal retrieval between panchromatic (PAN) and multi-spectral imagery, and ii) multi-label image retrieval between very high resolution (VHR) images and speech based label annotations. Notice that these multi-modal retrieval scenarios are more challenging than the traditional uni-modal retrieval approaches given the inherent differences in distributions between the modalities. However, with the growing availability of multi-source remote sensing data and the scarcity of enough semantic annotations, the task of multi-modal retrieval has recently become extremely important. In this regard, we propose a novel deep neural network based architecture which is considered to learn a discriminative shared feature space for all the input modalities, suitable for semantically coherent information retrieval. Extensive experiments are carried out on the benchmark large-scale PAN - multi-spectral DRSID dataset and the multi-label UC-Merced dataset. Together with the Merced dataset, we generate a corpus of speech signals corresponding to the labels. Superior performance with respect to the current state-of-the-art is observed in all the cases.

KEYWORDS

Deep learning, Unified representation, cross-modal retrieval, remote sensing, multimedia retrieval

1 INTRODUCTION

We address the problem of cross-modal information retrieval in the domain of remote sensing. In particular, we are interested in two application scenarios: i) cross-modal retrieval between panchromatic (PAN) and multi-spectral imagery, and ii) multi-label image retrieval between very high resolution (VHR) images and speech based label annotations. Notice that these multi-modal retrieval scenarios are more challenging than the traditional uni-modal retrieval approaches given the inherent differences in distributions between the modalities. However, with the growing availability of multi-source remote sensing data and the scarcity of enough semantic annotations, the task of multi-modal retrieval has recently become extremely important. In this regard, we propose a novel deep neural network based architecture which is considered to learn a discriminative shared feature space for all the input modalities, suitable for semantically coherent

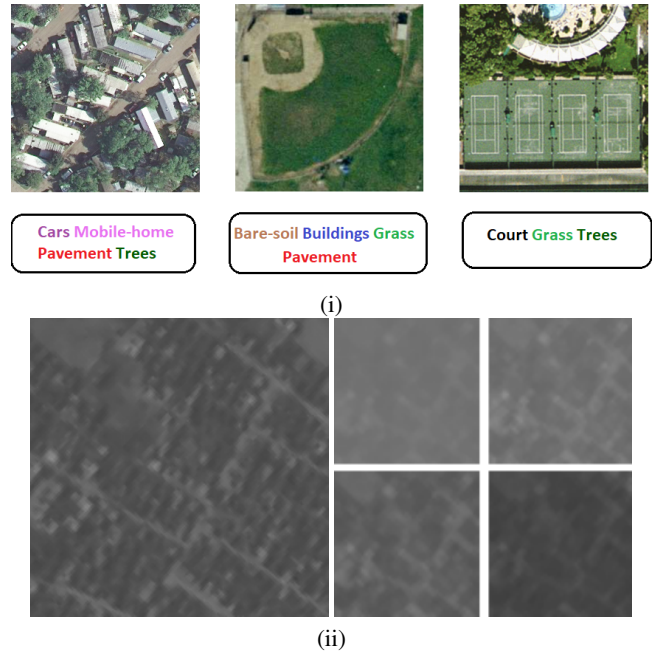


Figure 1: Examples of cross-modal information retrieval in RS: i) VHR RS images with multiple semantic labels, ii) PAN - multi-spectral bands of the same area on ground.

information retrieval. Extensive experiments are carried out on the benchmark large-scale PAN - multi-spectral DRSID dataset and the multi-label UC-Merced dataset. Together with the Merced dataset, we generate a corpus of speech signals corresponding to the labels. Superior performance with respect to the current state-of-the-art is observed in all the cases.

Although the notion of image retrieval has received extensive attention, several other cross-modal combinations have been tried in the past, specifically in the area of computer vision and natural language processing: image - text (label) pairs [21, 22], RGB image - depth image pairs [5, 26], etc. While most of the recent endeavors in this respect follow a closed-form formulation to associate the samples from different domains directly [11, 19, 21], others [13, 14, 22] depend on the conventional learning-based approaches. In

contrast to the single-label based image retrieval, few approaches extend the framework to support multiple semantic labels for each image, thus contriving the notion of multi-label image retrieval [2] (Figure 1).

In the same spirit, the notion of cross-modal information retrieval is practiced in the RS domain for many applications concerning the following modality pairs: SAR - optical, PAN - multi-spectral, image - text, to name a few. Note that these data can be made available in paired or unpaired fashion, depending upon the underlying data extraction strategy. In this respect, [14] classifies a generic cross-modal retrieval scenario into four broad categories: single-label paired, multi-label paired, single-label unpaired, and multi-label unpaired, respectively. By definition, the instances come as pairs in the *paired* case, else it is considered as *unpaired*. However, in both the scenarios, the training data are considered in pairs given the modalities to learn the similarity function while the retrieval strategy may differ. While a given paired dataset has the same number of instances in both the domains, the domains are of a different size for the unpaired dataset.

Unarguably the major challenge for designing a cross-modal retrieval framework lies in learning the unified feature space for the input data streams which is expected to be both discriminative and class-wise compact simultaneously. Once learned, such a common feature space can be used to compare the data from different domains efficiently. However, since both the modalities have mostly different data distributions, learning such a feature space is inherently non-trivial. The potentials of the deep learning techniques have been proven as far as the task of uni-modal image retrieval is concerned [3]. Despite the fact, such methods cannot be directly extended to the multi-modal case without additional constraints since it fails to adequately capture the underlying correlation among the samples from all the domains.

Our contributions: Inspired by the discussions mentioned above, we propose a deep neural network based framework for cross-modal retrieval (CMIR-NET) in RS given a pair of data streams depicting the i) paired single-label, and ii) unpaired multi-label scenarios. In particular, we consider paired PAN- multi-spectral images for the *paired* case, whereas we introduce a new scenario for the multi-label *unpaired* image - speech retrieval case. Specifically, we capture the speech signals of the semantic labels in different pronunciations/ascents which constitute our speech domain whereas the images with multiple semantic categories define our image domain.

Effectively, the proposed model has simultaneous feature extractors for both the modalities which project the data into a shared latent space. Several intuitive constraints are further introduced to ensure discriminative and domain-independence of the latent space. During inference, given the query data from any of the modalities, the related data from the other modality are retrieved as a ranked list using the nearest-neighbor search in the latent space.

We consider the large-scale DRSID dataset [12], consisting of paired samples of singly-labelled PAN and multi-spectral data for the image to image retrieval. In contrast, the second set of experiments are performed on the very high resolution (VHR) multi-label UC-Merced [24] dataset and a new uncorrelated speech corpus.

2 RELATED WORKS

In this section, we primarily restrict our discussions on (i) single and multi-label image retrieval from a uni-modal RS dataset, (ii) multi-modal image retrieval in the context of RS and computer vision, and (iii) feature extraction strategies for RS images.

Retrieval from uni-modal data: The image retrieval frameworks can either be content-based or text-based, depending upon nature of the query data. Some of the notable early works regarding content-based image retrieval (CBIR) in RS include [8, 16]. These approaches utilized the physical response properties relating to surface reflectance and absorption in different spectral regions, primitive shape, texture, color based parameters, to classify and retrieve the RS images. However, the applications of CBIR in RS mainly gained its motion after the Landsat thematic mapper multi-spectral satellite data were made accessible. More advanced feature extraction strategies (both at the pixel and object level) are adopted later on for improved retrieval performance [1, 3]. Subsequently, several techniques focus on better training the retrieval models through the notion of active learning [6] by incorporating the user’s knowledge during training iterations. As opposed to the real-valued feature representations, several techniques consider the binary features for quicker response during retrieval. In this regard, [4] proposes a hashing-based approximate nearest neighbor search for a fast and scalable RS image retrieval. Visual databases targeted by CBIR applications have been extended for multi-label image retrieval [2, 3]. In most of the works in RS for multi-label retrieval, region-based feature extraction is preferred over the pixel-based ones.

Retrieval from cross-modal data: As already mentioned, cross-modal retrieval has mainly received its attention in the image to text or text to image settings [10, 11, 14, 19, 25]. There have also been works on color image for depth perception for improved scene synthesis [5, 26]. However, there exists not much prior research regarding cross-modal retrieval in RS given the scarcity of multi-source images databases suitable for the retrieval task. In this respect, [12] very recently has proposed a large-scale dual-source remote sensing image dataset (DRSID) which comprises of panchromatic and multi-spectral image pairs acquired from the Geofen-1 satellite. Another attempt for cross-modal data retrieval in RS is by [15], which introduces a deep cross-modal retrieval framework for image and audio data. To the best of our knowledge, these are the only two reported works in RS in cross-modal retrieval. However, a major limitation to these works is that both of them are modality specific and not generic, unlike ours.

Image feature embeddings: Both hand-crafted and deep learning based feature extractors are applied in conjunction with the RS image retrieval. The ad-hoc features range from the basic color moments, texture, shape, and morphological descriptors as well as combinations of them [2]. Besides, the interest points based feature descriptors (SIFT, SURF) are explored along with different encoding strategies (VLAD, Fisher’s vector, etc.) [20]. The consideration of the region based features is also prominent in this respect given fine-grained visual categories [17].

The data-driven deep learning techniques have shown excellent performance for tasks including image classification, segmentation, and retrieval. Several endeavors have directly used the Imagenet pre-trained deep Convnet models like VGG-16 [17], GoogLeNet [18],

Resnet [7] for RS image retrieval, while others depend on fine-tuning the pre-trained models, for the task at hand [14]. The use of end-to-end models is also prominent in this respect which directly trains the feature extraction modules from the RS data [10]. Apart from the supervised Convnet architectures, the self-supervised encoder-decoder based representation learning techniques are also considered for RS image retrieval [23].

How we are different: In contrast to the very few previous methods like [12, 15] in cross-modal retrieval in RS, we focus on the discriminative and domain independence of the shared latent feature space simultaneously. We find that our real-valued feature learning strategy outperforms the previous discriminative hashing based feature encoding substantially. Besides, we introduce the novel notion of cross-modal retrieval between multi-label image and speech domains, respectively. Extensive experiments are performed to showcase the efficacy of the proposed framework.

3 METHODOLOGY

Preliminaries: Let \mathbf{A} and \mathbf{B} denote data from two different modalities/domains with shared labels from $\mathbf{L} \in \{1, 2, \dots, C\}$. In the paired case, both \mathbf{A} and \mathbf{B} have the same number of instances which come as pairs while the sizes of the domains are different for the unpaired case. In this regard, we define the training labeled data triplets as $\mathbf{X} = \{(a_i, b_i, l_i)\}$ where $a_i \in \mathbf{A}$, $b_i \in \mathbf{B}$, and $l_i \in \mathbf{L}$ is the semantic label for both a_i and b_i . In the proposed scenario, \mathbf{A} represents the image data while \mathbf{B} can be image or speech. However, they can also be realized in terms of features extracted from the instances in which case $a_i \in \mathbb{R}^{d_a}$ and $b_i \in \mathbb{R}^{d_b}$, respectively.

Under this setup, we aim for learning a unified feature representation space V with dimensions d_v . The instances from \mathbf{A} and \mathbf{B} are projected in such a way that i) the samples with same labels should be close to each other (irrespective of their domains) while samples from different classes should be distinctively placed, ii) given a triplet (a_i, b_i, l_i) from \mathbf{X} , the projected a_i and b_i should have highly similar feature representations. In this regard, let \mathbf{V}_{a_i} and \mathbf{V}_{b_i} represent the projected representations corresponding to a_i and b_i , respectively. During inference, the query image is projected onto V and subsequently the corresponding ranked list is obtained using a nearest-neighbor based approach.

To realize V , the proposed framework broadly consists of two parallel feature extractor networks for \mathbf{A} and \mathbf{B} , which are further compared in the shared V space. All the concerned loss measures are evaluated in the V space for ensuring its discriminative and domain-independence. A detailed depiction of the proposed framework can be found in Figure 2. In the following, both the training and inferences stages are discussed in greater details.

3.1 Training: modeling the unified feature space V

In order to learn $\mathbf{V}_a = \{\mathbf{V}_{a_i}\}$ and $\mathbf{V}_b = \{\mathbf{V}_{b_i}\}$ given \mathbf{X} , we propose a two-stage training process as follows:

- First, we train two separate domain specific classification networks for $\{(a_k, l_k)\}_{k=1}^{|\mathbf{A}|}$ and $\{(b_j, l_j)\}_{j=1}^{|\mathbf{B}|}$ where both the networks consist of a feature extraction module followed by the classification module ($|\cdot|$ denotes the number of samples in the respective domain). It is further ensured that the extracted features for given a_k and b_j , to be represented by \mathbf{Z}_{a_k}

and \mathbf{Z}_{b_j} henceforth, are highly non-redundant as far as the dependencies among the individual feature dimensions are concerned.

- $\{\mathbf{Z}_{a_i}\}$ and $\{\mathbf{Z}_{b_i}\}$ (corresponding to \mathbf{X}) are further considered as the inputs to the network designed for obtaining V .

The intuition behind obtaining the intermediate representations $\{\mathbf{Z}_{a_k/b_j}\}$ from the original data is to ensure that the $\{\mathbf{Z}_{a_i/b_i}\}$ now contains enough class information explicitly, which subsequently helps in learning a better V . Although we considered combining both the stages for an end-to-end training, however, we find that this setup leads to a sub-optimal V .

Generation of \mathbf{Z}_{a_k/b_j} : Given $\{(a_k, l_k)\}$ and $\{(b_j, l_j)\}$, separate neural network based architectures are deployed for both the modalities where the penultimate network layers (prior to the classification layer) represent the \mathbf{Z}_{a_k/b_j} features. To further ensure non-redundancy in the \mathbf{Z}_{a_k/b_j} features, a soft orthogonality constraint is included in the cost function along with the standard cross-entropy based classification loss as follows:

$$\mathcal{L}_{\mathbf{A}/\mathbf{B}} = \text{CE}(\tilde{\mathbf{Z}}_{a/b}) + \|\tilde{\mathbf{Z}}_{a/b}^T \tilde{\mathbf{Z}}_{a/b} - \mathbf{I}\|_{\text{F}}^2 \quad (1)$$

where $\tilde{\mathbf{Z}}_a = \{\mathbf{Z}_{a_k}\}$, $\text{CE}(\tilde{\mathbf{Z}}_a)$ denotes the cross-entropy loss for the labeled samples from \mathbf{A} , and \mathbf{I} is a vector of ones, respectively. The network for obtaining $\{\mathbf{V}_{a_i/b_i}\}$ is trained henceforth given the learned $\{\mathbf{Z}_{a_i/b_i}\}$.

Construction of V : To construct the unified representations $\{\mathbf{V}_{a_i}\}$ and $\{\mathbf{V}_{b_i}\}$ from $\{\mathbf{Z}_{a_i}\}$ and $\{\mathbf{Z}_{b_i}\}$, we use a neural network based discriminative encoder-decoder architecture which minimizes the following four loss measures:

- the difference between each pair of corresponding i^{th} samples in \mathbf{V}_a and \mathbf{V}_b : (\mathcal{L}_2).
- a classification loss on $\mathbf{V}_{ab} = [\mathbf{V}_a, \mathbf{V}_b]$: (\mathcal{L}_3).
- separate feature norm loss measures on both \mathbf{V}_a and \mathbf{V}_b , respectively: (\mathcal{L}_4).
- a decoder loss which is deemed to reconstruct cross-domain samples given the latent representations: (\mathcal{L}_5).

\mathcal{L}_2 ensures that the latent feature embeddings for similar samples from both the domain turn out to be highly analogous in the V space, while \mathcal{L}_3 imposes an inter-class separation constraint by minimizing the cross-entropy loss on all the samples from both the domains. In addition, we penalize any unbounded modulation of the latent representations $\mathbf{V}_{a/b}$ by explicitly minimizing their ℓ_2 norm in \mathcal{L}_4 . Finally, the decoder loss \mathcal{L}_5 ensures the domain invariance of the space V by reconstructing the cross-domain samples given the latent $\mathbf{V}_{a/b}$.

For \mathcal{L}_2 , let w_a and w_b be the learnable parameters such that $\mathbf{V}_{a_i} = \mathbf{Z}_{a_i} w_a$ and $\mathbf{V}_{b_i} = \mathbf{Z}_{b_i} w_b$, and w_c denotes the parameters for the classifier block. Given that, we define \mathcal{L}_2 , \mathcal{L}_3 , and \mathcal{L}_4 as follows:

$$\mathcal{L}_2 = \|\mathbf{V}_a - \mathbf{V}_b\|_{\text{F}}^2 \quad (2)$$

$$\mathcal{L}_3 = \text{CE}(\mathbf{V}_{ab}) \quad (3)$$

$$\mathcal{L}_4 = \|\mathbf{V}_a\|_{\text{F}}^2 + \|\mathbf{V}_b\|_{\text{F}}^2 \quad (4)$$

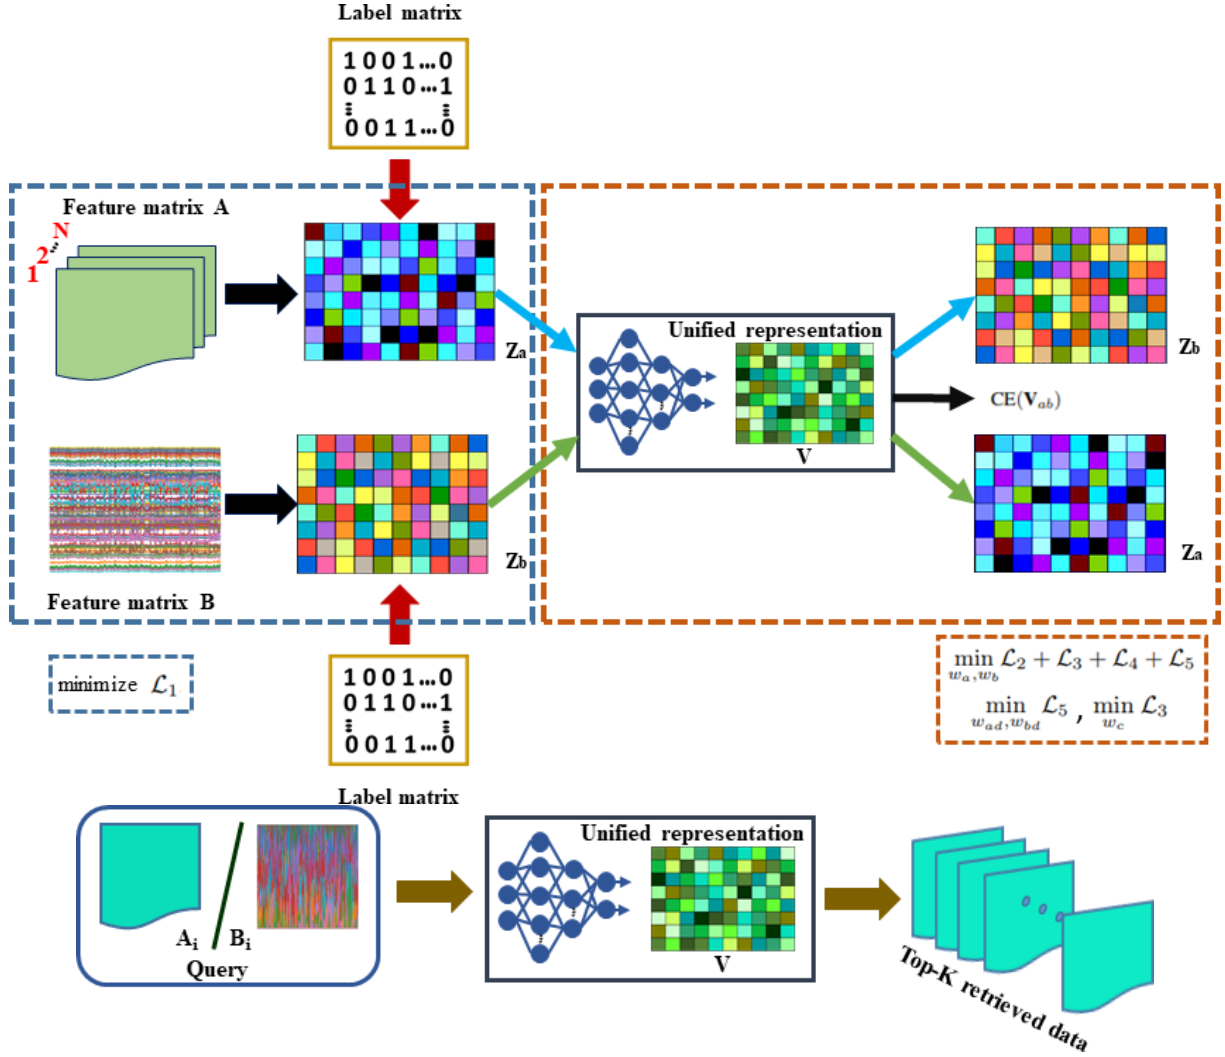


Figure 2: Overall pipeline of the proposed CMIR-NET architecture for learning the unified embedding space from two different modality of signals. The loss functions are defined and discussed in detail in section 3.1.

Similarly for \mathcal{L}_5 , let w_{ab} and w_{ba} denote the learnable weights for the decoder blocks. Given the latent V_a , we try to reconstruct Z_b in the decoder and vice-versa by as follows:

$$\mathcal{L}_5 = \|w_{ab}V_a - Z_b\|_F^2 + \|w_{ba}V_b - Z_a\|_F^2 \quad (5)$$

where $Z_{a/b} = \{Z_{a_i/b_i}\}$.

The overall objective function and optimization: We consider a weighted combination of the loss mentioned above measures for obtaining the latent space V . In particular, the overall loss function to be minimized can be mentioned as:

$$\mathcal{L} = \lambda_1\mathcal{L}_2 + \lambda_2\mathcal{L}_3 + \lambda_3\mathcal{L}_4 + \lambda_4\mathcal{L}_5 + \lambda_5\mathcal{R} \quad (6)$$

given the non-negative weights λ_{1-5} and \mathcal{R} defines the standard ℓ_2 regularizer on w_a and w_b , respectively. In general, \mathcal{R} takes the form of:

$$\mathcal{R} = \|w_a - \alpha\|_F^2 + \|w_b - \alpha\|_F^2 \quad (7)$$

for $\alpha \geq 0$.

We follow the standard alternate stochastic mini-batch gradient-descent based optimization strategy for minimizing \mathcal{L} . Algorithm 1 mentions the stages for minimizing \mathcal{L} .

3.2 Uni-modal and Cross-modal Retrieval

During inference, given a query $(a/b)^{query}$ from any of the modalities, it is possible to perform i) uni-modal, ii) cross-modal, and iii) mixed-modal information retrieval in the V space. The retrieval is based on the nearest-neighbor search in V which outputs the top- K data according to their similarity with $(a/b)^{query}$ in terms of the Euclidean distance measure (Algorithm 1).

Algorithm 1 The proposed training and inference stage

Input: $\{(a_k, l_k)\}, \{(b_j, l_j)\}$, and \mathbf{X} **Output:** Unified representations $\mathbf{V}_{a/b}$ ($w_a \mathbf{Z}_a$ and $w_b \mathbf{Z}_b$).

- 1: Normalize \mathbf{A} and \mathbf{B} .
- 2: Generate intermediate representations $\{\mathbf{Z}_{a_k}\}$ and $\{\mathbf{Z}_{b_j}\}$ by minimizing $\mathcal{L}_{\mathbf{A}/\mathbf{B}}$.
- 3: Train the network to obtain V by optimizing \mathcal{L} . The optimization follows the following stages:

4: **do**

5:

$$\min_{w_a, w_b} \lambda_1 \mathcal{L}_2 + \lambda_2 \mathcal{L}_3 + \lambda_3 \mathcal{L}_4 + \lambda_4 \mathcal{L}_5 \quad (8)$$

6:

$$\min_{w_{ad}, w_{bd}} \mathcal{L}_5 \quad (9)$$

7:

$$\min_{w_c} \mathcal{L}_3 \quad (10)$$

8: **while** until convergence9: **return** w_a and w_b (for projecting data onto V)

Input: $a \in \mathbf{A}$ or $b \in \mathbf{B}$ **Output:** Top- K retrieved data.

- 10: Uni-modal retrieval using K -NN from $w_a \mathbf{Z}_a$ **or** $w_b \mathbf{Z}_b$.
 - 11: Cross-modal retrieval using K -NN from $w_a \mathbf{Z}_a$ **and** $w_b \mathbf{Z}_b$.
-

4 EXPERIMENTS

We discuss the performance analysis of the proposed method in the following.

Datasets and model architecture: For the image to image retrieval case, we consider the large-scale DSRSID dataset proposed in [12]. The dataset comprises of 80,000 pairs of panchromatic and multi-spectral images from 8 major land-cover classes where both the images of a given pair focus on the same geographical area on the ground (Figure 3). The images are acquired by the GF-1 panchromatic and GF-1 multi-spectral sensors, respectively. The panchromatic data samples are of image size 256×256 , with a spatial resolution of 2m and single spectral channel ($\mathbf{A} \in \mathbb{R}^{256 \times 256 \times 80,000}$), while the multi-spectral images are of dimensions 64×64 , with a resolution of 8m, and 4 spectral channels ($\mathbf{B} \in \mathbb{R}^{64 \times 64 \times 4 \times 80,000}$). In order to learn $\tilde{\mathbf{Z}}_{a/b}$ from the raw image data, we consider two separate convolutional network architectures. For each of the networks, we use three convolution - pooling - non-linearity blocks as: $(3 \times 3 \times \text{channel} \times 32)$, $(3 \times 3 \times 32 \times 32)$, and $(3 \times 3 \times 64 \times 128)$ where *leaky_ReLU*(\cdot) is considered as the non-linear activation. Note that the batch-normalization layer is also appended after each convolution block. After the convolution modules, we consider a fully-connected layer of size $((4 \times 4 \times 128) \times 128)$, which is activated by a *ReLU*(\cdot) function and a drop-out layer with probability 0.5. Hence, the dimensions of $\mathbf{Z}_{a/b}$ is 128-d.

For the second set of experiments, we consider the multi-label VHR UC-Merced dataset [24] (domain \mathbf{A}), where the \mathbf{B} domain corresponds to speech signals. In particular, we construct a corpus of spoken speech samples for each of the land-cover semantic labels in .wav format. In order to increase the diversity of the speech samples, it is ensured that the labels are pronounced with different

English accents. In this way, We gather 15 speech samples for each label, leading to 255 speech samples in total. Also note that the multi-label UC-Merced dataset consists of 2100 VHR images of size 256×256 where each image has multiple associated semantic labels from a set of 17 land-cover categories [2]. As opposed to the DSRSID dataset where features are directly learned from the images, we consider the features extracted from pre-trained Convnet networks for initial feature extraction for this data. Essentially, the \mathbf{B} space of speech signals is constructed using the *vgg-ish* [9] model which is pre-trained on the large-scale AudioSet dataset and outputs a $d_b = 128$ -d vector corresponding to each input signal. On the other hand for \mathbf{A} , the images are represented using $d_a = 4096$ -d Imagenet pre-trained VGG-16 [17] features. In order to obtain $\tilde{\mathbf{Z}}_a$, we train a multi-label classifier for the images while the standard 17-class multi-class classifier is trained for the speech signals in order to obtain $\tilde{\mathbf{Z}}_b$. The size of the final $\mathbf{Z}_{a/b}$ space is 128-d, similar to DSRSID.

Subsequently, the network responsible for obtaining V consists of three fully-connected neural network layers each coupled with *ReLU*(\cdot) non-linearity both in the encoder as well as the decoder branches. The latent space with dimensions d_v (the third encoder layer) is used as V for both the dataset.

Training protocol and evaluation: We consider the Adam optimizer with a learning rate of 0.01 and batch size of 64 for training both the networks (corresponding to \mathbf{Z} and V). For performance evaluation, we report the mAP and P@10 scores. Also, we compare the performance on the DSRSID data with three techniques from the literature: SCM [25], DCHM [10], and SIDGCNN [12], all of which are based on discriminative hash-code learning. In addition, we perform extensive ablation study to showcase the importance of i) the individual loss terms, ii) the hyper-parameter α , iii) comparison between $\mathbf{V}_{a/b}$ and $\mathbf{Z}_{a/b}$ features, and iv) separate pre-training of \mathbf{Z} features.

4.1 Results on DSRSID dataset

Following the protocol mentioned in [12], we consider a split of 75000-5000 for constructing the training and test sets for this dataset. The same set of samples are used for all similar techniques to avoid any bias. The λ_1 , to λ_4 parameters (Equation 6) are set to 1 since we did not find the training to be sensitive to the choice while we consider several values for λ_5 in the range $[1, 0.001]$ and consider the one giving the best performance on a validation set (considering unimodal retrieval). Similarly, the weight-norm parameter α of Equation 7 is set to 0 as a higher value is found to degrade the subsequent retrieval precision.

We report the performance on both uni-modal as well as cross-modal retrieval for this dataset on the V space. In this regard, we consider different dimensions of the embedding space V and report a sensitivity analysis in Table 1. It can be found that the retrieval accuracy for this dataset is mainly invariant to the dimensionality of v , although we obtain the best performance for $d_v = 64$. On the other hand, we compare the performance of our technique with the literature only for the cross-modal scenarios. It can be observed from Table 2 that the proposed technique outperforms SCM and DCMH substantially by a large margin ($\geq 13\%$) in terms of the mAP score for both the cases when the query image comes from \mathbf{A}




Figure 3: Average precision for top-50 retrieval for both PAN→MUL and MUL→PAN cases.

Table 1: Performance of the proposed CMIR-NET framework on the DRSID dataset in terms of mAP (%) and precision at top-10 (P@10) (%) values, under different embedding vector code lengths (d_v).

Task	$d_v=16$		$d_v=32$		$d_v=64$	
	mAP	P@10	mAP	P@10	mAP	P@10
Pan→Mul	95.52	97.10	98.96	98.99	99.05	99.40
Mul→Pan	98.77	99.00	97.95	97.99	98.93	98.60
Pan→Pan	99.41	99.82	98.11	98.40	98.69	99.40
Mul→Mul	99.55	99.69	98.18	98.60	98.25	98.40

Table 2: Comparison of mAP values of our proposed framework and other comparative algorithms based on the code length of the embedding space.

Task	Model	$d_v=16$	$d_v=32$
		mAP(%)	mAP(%)
Pan→Mul	SCM [25]	34.72	37.67
	DCMH [10]	80.76	85.09
	SIDHCNN [12]	95.53	96.43
	Proposed CMIR-NET	95.52	98.96
Mul→Pan	SCM [25]	36.71	38.71
	DCMH [10]	80.23	84.45
	SIDHCNN [12]	97.25	97.89
	Proposed CMIR-NET	98.77	97.95

or **B** (PAN ↔ MUL). Likewise for 32-dimensional V , we observe that our method outperforms SIDHCNN for both PAN→MUL and MUL→PAN. On the other hand for 16-dimensional V space, we surpass the performance of SIDHCNN for MUL→PAN by 2%, while we report comparable retrieval performance to SIDHCNN for the PAN→MUL case. Overall for $d_v = 32$ and 64, we report the new state-of-the-art performance for this dataset. Additionally, Figure 3 shows the average class-wise precision both the PAN→MUL and MUL→PAN retrieval cases for the top-50 retrieval scenario. It can be observed that the proposed method produces very high precision measures for all the classes.

4.2 Results on UC Merced-Audio dataset

As already mentioned, this dataset contains multi-label images with a varied number of labels for each of the classes ranging between 100 and 1300. Hence, while training the network for V , we feed each

Table 3: Performance of the CMIR-NET framework on UC Merced-Audio dataset in terms of mAP and P@10 values, with variation in embedding vector code length (d_v).

Model	$d_v=32$		$d_v=64$		$d_v=128$	
	mAP	P@10	mAP	P@10	mAP	P@10
Img→Aud	29.67	60.91	41.60	63.15	62.11	64.81
Aud→Img	21.60	40.11	42.36	51.29	54.21	56.00




Figure 4: Ablation study with different losses.

image with a single-label speech signal with the same label at random. As a result, the size of \mathbf{Z}_a becomes $\mathbb{R}^{7004 \times 128}$, owing to 7004 distinct labels corresponding to the multi-label images. We associate each image with a randomly selected speech signal representing one of its categories which construct the features \mathbf{Z}_b .

We consider a 70%: 30% training - test split for this dataset. The parameters $\lambda_1, \lambda_3, \lambda_4$, and λ_5 are set to 0.00001, 0.01, 0.01, 1, respectively, upon grid-search on a validation set (similar to DRSID). α is set to 1 for this case (more about α in Section 4.3). We report the mAP and the P@10 values for both the cross-modal scenarios where we vary the dimensions of V as 32, 64, and 128, respectively (Table 3). We observe that the performance increases substantially with increasing size of the V space. On the other hand, We find that the performance of the Image to Audio retrieval case is substantially superior to the Audio to Image retrieval case. This can be attributed to the complexity of the multi-label annotations of the image space and the high intra-class variability of the speech signals.

4.3 Critical analysis

Advantage of pre-training $\mathbf{Z}_{a/b}$ over end-to-end training of the entire model: We find that the training error abruptly modulates when only one network is considered to learn the V space directly from the input data (bypassing the modeling of $\mathbf{Z}_{a/b}$), which further leads to a somewhat less discriminative V space. We observe a steady enhancement of at least 3–5% in mAP values for all the experimental cases when $\mathbf{Z}_{a/b}$ are separately learned.

Ablation on the loss terms: Here we analyze the effect of the individual loss terms. First, we consider the cross-modal retrieval

Table 4: Sensitivity to critical parameter α for the UC-Merced-Audio case.

Model	$\alpha = 0$		$\alpha = 1$		$\alpha = 2$	
	mAP	P@10	mAP	P@10	mAP	P@10
Img→Aud	0	0	62.11	64.81	32.09	52.77
Aud→Img	0	0	54.21	56.00	33.64	45.01

experiments directly from the feature matrices \mathbf{Z}_a and \mathbf{Z}_b . We find that the mAP and P@10 scores are less than 15 % for all the four experimental scenarios (for both the dataset) with $\mathbf{Z}_{a/b}$ features. We refer to this case as the base model, on top of which we add the proposed loss terms incrementally and showcase their effectiveness. On the other hand, the model with all the loss measures (Equation 6) is henceforth termed as the full model. In particular, Fig. 4 shows a bar graph based analysis, wherein we show the mAP values obtained on the two experimental cases: Image→Audio (UC-Merced) and PAN→MUL (DSRSID).

After studying the base case, we consider the effect of the latent loss (\mathcal{L}_2) by examining the performance of the full model excluding the \mathcal{L}_2 loss in Equation 6. It is found that \mathcal{L}_2 boosts the performance of the model quite meagerly. We then investigate the effect of decoder loss (\mathcal{L}_5) by excluding this loss from the full model. The decoder loss function is found to be essential to boost the performance of the framework, and it restricts any trivial solution in the V space. It is observed that without the consideration of \mathcal{L}_5 , most of the latent feature vectors rapidly approach to zero values in V . We furthermore follow a similar protocol to study the effect of the classification loss function (\mathcal{L}_3). The classification loss help in fine-tuning the performance of the model. This is achieved by encoding the respective class information labels in the unified embedding vector. By keeping the entire loss function as in Equation 6, we see that we can achieve a very high retrieval performance.

Sensitivity to weight norm parameter α : We show the effect of the critical parameter α which is used in the weight regularizing term \mathcal{R} (Equation 7). Setting α to 0 for both the cross-modal retrieval scenarios of the UC-Merced↔Audio dataset leads to the learning of a trivial V where all the projected samples collapse to a single point, and hence one is unable to perform meaningful retrieval. The performance of the model on the UC Merced↔Audio dataset with different α values can be obtained in Table 4. It shows that the best retrieval performance can be achieved for $\alpha = 1$ for both the combinations in this case. On the other hand, we observe that $\alpha = 0$ gives the best results for the DSRSID dataset.

5 CONCLUSIONS

We propose a novel framework for cross-modal information retrieval and evaluate the same in conjunction with remote sensing data. The framework focuses on learning a unified and discriminative embedding space from different input modalities. The proposed model is generic enough to handle both uni-modal and cross-modal information retrieval scenarios. We further introduce a novel experimental scheme of cross-modal retrieval between the multi-label image and audio domains where a variety of speech signals are used to represent the semantic labels. As a whole, we showcase the performance

of the proposed model on the large-scale DSRSID and the image-audio dataset where improved performance can be observed. We are currently interested in extending the model to support more than two modalities and also in exploring the possibility of using compact hash codes to learn the shared embedding space.

REFERENCES

- [1] E. Aptoula. 2014. Remote Sensing Image Retrieval With Global Morphological Texture Descriptors. *IEEE Transactions on Geoscience and Remote Sensing* 52, 5 (May 2014), 3023–3034.
- [2] B. Chaudhuri, B. Demir, S. Chaudhuri, and L. Bruzzone. 2018. Multilabel Remote Sensing Image Retrieval Using a Semisupervised Graph-Theoretic Method. *IEEE Transactions on Geoscience and Remote Sensing* 56, 2 (Feb 2018), 1144–1158.
- [3] O. E. Dai, B. Demir, B. Sankur, and L. Bruzzone. 2017. A novel system for content based retrieval of multi-label remote sensing images. In *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*. 1744–1747. <https://doi.org/10.1109/IGARSS.2017.8127311>
- [4] B. Demir and L. Bruzzone. 2016. Hashing-Based Scalable Remote Sensing Image Search and Retrieval in Large Archives. *IEEE Transactions on Geoscience and Remote Sensing* 54, 2 (Feb 2016), 892–904. <https://doi.org/10.1109/TGRS.2015.2469138>
- [5] David Eigen, Christian Puhrsch, and Rob Fergus. 2014. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*. 2366–2374.
- [6] Marin Ferecatu, Michel Crucianu, and Nozha Boujemaa. 2004. Retrieval of difficult image classes using svd-based relevance feedback. In *Proceedings of the 6th ACM SIGMM international workshop on Multimedia information retrieval*. ACM, 23–30.
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [8] G. Healey and A. Jain. 1996. Retrieving multispectral satellite images using physics-based invariant representations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 18, 8 (Aug 1996), 842–848. <https://doi.org/10.1109/34.531804>
- [9] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. 2017. CNN architectures for large-scale audio classification. In *2017 IEEE international conference on acoustics, speech and signal processing (icassp)*. IEEE, 131–135.
- [10] Qing-Yuan Jiang and Wu-Jun Li. 2017. Deep cross-modal hashing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3232–3240.
- [11] Kai Li, Guo-Jun Qi, Jun Ye, and Kien A Hua. 2017. Linear subspace ranking hashing for cross-modal retrieval. *IEEE transactions on pattern analysis and machine intelligence* 39, 9 (2017), 1825–1838.
- [12] Yansheng Li, Yongjun Zhang, Xin Huang, and Jiayi Ma. 2018. Learning source-invariant deep hashing convolutional neural networks for cross-source remote sensing image retrieval. *IEEE Transactions on Geoscience and Remote Sensing* 99 (2018), 1–16.
- [13] Hong Liu, Rongrong Ji, Yongjian Wu, Feiyue Huang, and Baochang Zhang. 2017. Cross-modality binary code learning via fusion similarity hashing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7380–7388.
- [14] Devraj Mandal, Kunal N Chaudhury, and Soma Biswas. 2017. Generalized semantic preserving hashing for n-label cross-modal retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4076–4084.
- [15] Gou Mao, Yuan Yuan, and Lu Xiaoqiang. 2018. Deep Cross-Modal Retrieval for Remote Sensing Image and Audio. In *2018 10th IAPR Workshop on Pattern Recognition in Remote Sensing (PRRS)*. IEEE, 1–7.
- [16] J. F. Moreno and S. S. Saatchi. 1996. Surface information retrieval from optical/microwave data: potentials and limits of synergistic approaches. In *IGARSS '96. 1996 International Geoscience and Remote Sensing Symposium*, Vol. 3. 1861–1863 vol.3. <https://doi.org/10.1109/IGARSS.1996.516821>
- [17] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [18] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1–9.
- [19] Jun Tang, Ke Wang, and Ling Shao. 2016. Supervised matrix factorization hashing for cross-modal retrieval. *IEEE Transactions on Image Processing* 25, 7 (2016), 3157–3166.
- [20] Andrea Vedaldi and Brian Fulkerson. 2010. VLFeat: An open and portable library of computer vision algorithms. In *Proceedings of the 18th ACM international conference on Multimedia*. ACM, 1469–1472.

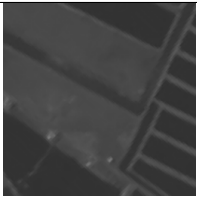
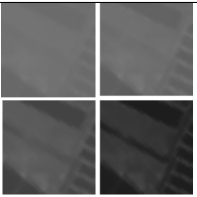

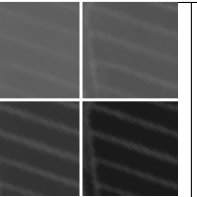
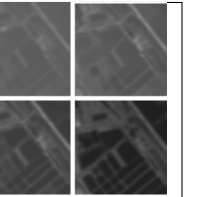
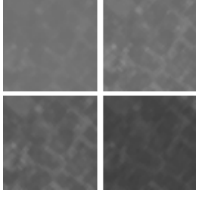






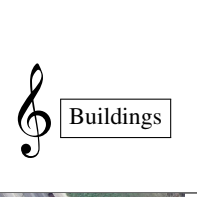


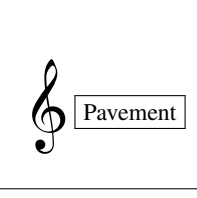




Dataset: Query	Retrieved 1	Retrieved 2	Retrieved 3	Retrieved 4
				
				
				
				

Figure 5: Top-4 retrieval instances from all the four combinations of cross-modal query samples (first column) are shown. The first two rows correspond to PAN→MUL (Aquafarm class), and MUL→PAN (High buildings class) respectively. The next two rows show the UC-Merced↔audio multi-label retrieval instances.

[21] Di Wang, X-B Gao, Xiumei Wang, and Lihuo He. 2018. Label Consistent Matrix Factorization Hashing for Large-Scale Cross-Modal Similarity Search. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2018).

[22] Liwei Wang, Yin Li, Jing Huang, and Svetlana Lazebnik. 2019. Learning two-branch neural networks for image-text matching tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 2 (2019), 394–407.

[23] Yuebin Wang, Jie Mei, Liqiang Zhang, Bing Zhang, Panpan Zhu, Yang Li, and Xingang Li. 2018. Self-Supervised Feature Learning With CRF Embedding for Hyperspectral Image Classification. *IEEE Transactions on Geoscience and Remote Sensing* (2018).

[24] Yi Yang and Shawn Newsam. 2010. Bag-of-visual-words and Spatial Extensions for Land-use Classification. In *Proceedings of the 18th SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS '10)*. ACM, New York, NY, USA, 270–279. <https://doi.org/10.1145/1869790.1869829>

[25] Dongqing Zhang and Wu-Jun Li. 2014. Large-scale supervised multimodal hashing with semantic correlation maximization. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.

[26] Yinda Zhang and Thomas Funkhouser. 2018. Deep depth completion of a single rgb-d image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 175–185.