

CR-Cross: A Novel Approach for Cross Domain Coral Recognitions with Reject Options

Anonymous submission

Abstract

Although coral reefs are special and vital marine ecosystems, massive coral degradation began to occur due to the increase in global temperatures and the intensification of human industrial activities. Coral reef protection requires accurate coral recognition because it is the foundation for learning the distribution, disease, and growth of coral reefs, hereby informing the proper ways for further action. Recently, CNNs have been applied in automated coral image classification. These classifier models, however, are difficult to be generalized from the trained coral images in a marine region (source domain) to the coral images in a different marine region (target domain) since the corals have significant within-species morphological variability among the different geographic location domains. In this paper, a novel coral recognition algorithm is introduced via knowledge transfer across domains and its advantages lie in the following aspects. (1) It simultaneously transfers corals' texture and structure features across domains thus providing useful knowledge to assist the coral recognition tasks in the target marine domain. (2) To overcome the difficulty that the confusing coral images (e.g., bleached corals) are prone to be misclassified and transfer useless or even negative information, our algorithm is equipped with the reject option for the confusing corals while adapting. These corals can be sent to an expert or a more expensive but accurate system, resulting in strengthened transferability and reliability. Furthermore, we develop a new cross-domain coral image dataset to enhance coral research. Without the label information from the target marine region, our method significantly reduces the distribution gap and domain shift among the different marine regions. In addition, CR-Cross goes a step further in tackling the challenges of missing coral data, maximizing the utilization of available coral datasets, and enhancing the reusability of both coral data and coral recognition models. A series of empirical studies show that our method remarkably outperforms a broad range of baselines.

Introduction

Coral reefs are special and vital marine ecosystems, which have exceptionally high biodiversity and are often referred to as the tropical rainforest of the ocean. Coral reefs directly provide abundant biological resources for human production and livelihood, including plentiful fishery resources, raw materials for anti-cancer and antibacterial drugs. However, since the 1980s, mass coral degradation began to occur due to the increase in global temperatures and the intensi-

fication of human industrial activities (Reaka-Kudla 1997). According to, less than 30% of coral reefs worldwide are currently in a healthy state (Mulhall 2008). Indeed, coral reef reservation and restoration has been an emergent issue of global interest, attracting a great deal research efforts from multiple disciplines. Coral reef protection often requires accurate coral recognition because it is the foundation for learning the distribution, disease, and growth of coral reefs, hereby informing the proper ways for further actions (Mora et al. 2006). With the latest progress in image acquisition technologies (using autonomous underwater vehicles), substantial volumes of coral reef images are currently being amassed. On the foundation of these extensive image data, deep learning-based computer vision methods (powerful yet inexpensive) have been widely adopted in automated coral recognition to alleviate the burden on ocean scientists. To achieve better generalization capability, these deep models need to collect plenty of labeled training coral images that come from the same marine region as the test images. However, in real coral protection tasks, training and test images are always from different marine regions since it is expensive to re-labeling training images by ocean scientists. Different marine regions may lead to significant changes in the structure and texture of coral images in the same class. The different image data distributions between regions (or domains) limit the generalization capability of deep models.

In coral recognition, bridging the distribution discrepancy and realizing the knowledge transfer across domains is not an easy problem, and it has several great challenges to deal with. (1) Coral-specific features. Existing transfer learning models mostly focus on extracting texture features from coral images. However, for corals, texture and structural information together determine species. If only local texture features are considered in the transfer learning process, it can affect the transfer of valuable and discriminative information. (2) Confusing-coral rejection. In an unknown marine area, even for the same coral species, corals exhibit evident intra-specific morphological variability, in terms of age, diseases, depth, light, and species competition (Raphael et al. 2020). These confusing coral images are prone to be misclassified and transfer useless or even negative information. It is important to estimate when the coral recognition system is in doubt and then allow it to abstain from target domain predicting. The system can turn over the

decision-making power to an ocean expert or a more expensive but accurate system. Consequently, the system will become smarter and more reliable. The purpose of our paper is to address the problems of bridging the distribution discrepancy, realizing the knowledge transfer across domains, maximizing the utilization of available coral datasets, and eliminating the confusing coral images. We utilize the coral’s texture and structure feature from the source marine region to achieve cross-domain coral recognition from the target marine region. By the marginal distribution alignment and the reject rule, we reduce the distribution discrepancy among the different marine regions and exclude the confusing coral images, which further strengthens the transferability and reliability of our method.

Related Works

Coral Recognition

In a recent effort to automate coral recognition, Mehta et al. (2007) utilize the support vector machines (SVM) and the texture features of coral to classify coral reef images. Stokes and Deane (2009) propose a method that applies the texture, normalized color space, and the k -nearest neighbor classifier for the coral reef benthic classification. Recently, Convolutional Neural Networks (CNN) have demonstrated outstanding performance in image classification (He et al. 2016; Mahmood et al. 2016). In addition, hybrid feature and data augmentation is applied to improve the performance of the coral recognition (Mary and Dharma 2017; Gómez-Ríos et al. 2019; Asha Paul, Arockia Jansi Rani, and Liba Manopriya 2020), which enhances the model’s generalization and robustness. However, all the aforementioned studies focus on a single dataset from one marine region, while distribution discrepancy among the different marine regions has not been taken into account.

Unsupervised Domain Adaptation

Deep learning has achieved impressive performance in various classification tasks. The outstanding performance is premised on the abundant annotated training samples and training data is drawn from identical distribution with the test data. However, it is unrealistic in various real applications. Therefore, unsupervised domain adaptation is proposed to address the problems of the lack of annotated training samples and the distribution discrepancy between train data and test data. Domain adaptation can transfer the knowledge from the rich labeled source domain to the unlabeled target domain. To reduce the domain shift, on the one hand, unsupervised domain adaptation always defines a specific metric for distribution discrepancy. Zellinger et al. (2017) and Long et al. (2017) respectively propose the central moment discrepancy and the joint maximum mean discrepancy to align the distribution among the different domains. On the other hand, unsupervised domain adaptation learns a domain-invariant feature representation by the extra domain discriminator (Ganin et al. 2016; Long et al. 2018; Du et al. 2021).

Method

In this section, we present a new architecture, named Cross Domain Coral Recognitions with Reject Options (CR-Cross). As shown in Figure 1, CR-Cross is composed of four components: a feature extractor, a feature fusion model, a domain adaptation model, and the reject rule. The core idea is to utilize the texture features, structure features, marginal distribution alignment, and the reject rule to improve the cross-domain coral recognition performance.

Problem Definition and Notations

Based on the coral recognition, a sample of n labeled training data $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ from $\mathcal{X} \times \mathcal{Y}$ is fed into the model, where \mathcal{X} is the input space of the model and $\mathcal{Y} = \{1, \dots, k\}$ is the output space of the model in multi-class classification. During the testing stage, the model receives the m labeled test data $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m)\}$ from $\mathcal{X} \times \mathcal{Y}$. In the supervised coral recognition setting, the training data and test data are drawn from the same distribution P . In the cross-domain coral recognition setting, let $\{(\mathbf{x}_1^s, y_1^s), \dots, (\mathbf{x}_n^s, y_n^s)\}$ be a set of n labeled samples from $\mathcal{X} \times \mathcal{Y}$ in the source marine region. Let $\{\mathbf{x}_1^t, \dots, \mathbf{x}_m^t\}$ be m unlabeled sample from \mathcal{X} in target marine region. The source marine region and target marine region are drawn from different distributions, and we denote P_s and P_t as the marginal probability distributions of the source marine region and target marine region, respectively.

Texture Feature and Structure Feature

Almost all coral recognitions are based on the coral’s texture feature, which has led to the loss of structure information. Therefore, we design a new architecture to simultaneously transfer the coral’s domain-invariant texture feature and structure feature between the source marine region and the target marine region.

We apply the ResNet50 as a backbone network to capture the texture features and employ the hierarchical coupling ViT as a backbone network to obtain the structure features. According to the setting of ResNet50 (He et al. 2016), the entire network is divided into five stages. The first stage involves preprocessing the input data. The other stages are composed of multiple bottlenecks including the 1×1 down-sampling convolution, 3×3 spatial convolution, and 1×1 up-sampling convolution for the extraction of texture features. Following the ViT (Dosovitskiy et al. 2020), the ViT is composed of the 12 transformer encoder blocks. The transformer encoder blocks consist of multi-head self-attention (MSA) and MLP blocks. With self-attention, the semantic dependencies among different patches can be better extracted, thereby capturing coral’s structure features more effectively. In order to obtain structure features of different hierarchical levels and achieve better generalization, we design grouped transformer encoder blocks, and the transformer encoder blocks are sequentially divided into three groups. In order to capture structure features at various hierarchical levels, the feature from the last encoder block of the different groups simultaneously outputs to the decoder block. The hierarchical feature from the decoder block is fed

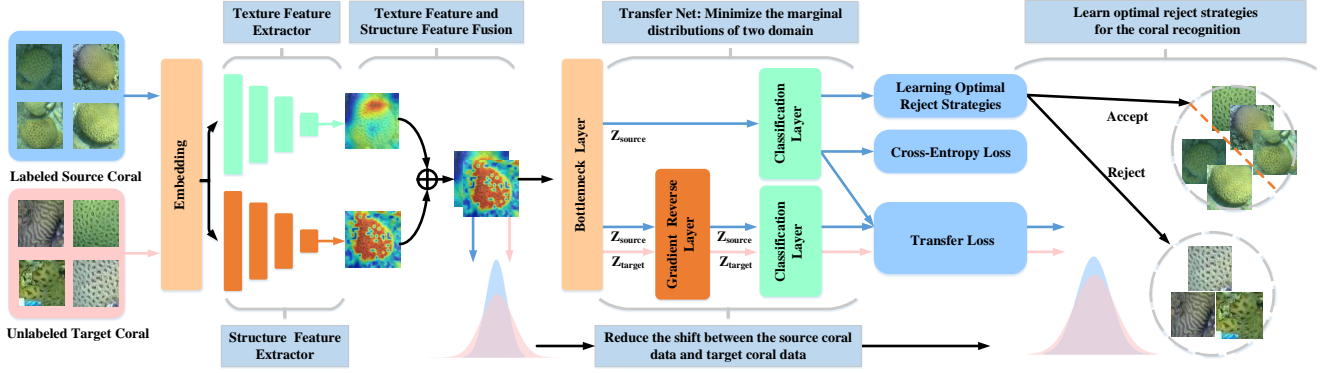


Figure 1: Framework of CR-Cross. CR-Cross is composed of four components: a feature extractor, a feature fusion model, a transfer network model, and the reject rule. The core idea is to utilize the texture features, structure features, marginal distribution alignment, and the reject rule to improve the cross-domain coral recognition performance.

into the domain adaptation layer and fused with the texture feature from ResNet50.

Marginal Distribution Alignment

To overcome the difficulty that the confusing coral images (e.g., bleached corals) are prone to be misclassified and transfer useless or even negative information, CR-Cross is equipped with the reject option for the confusing corals while adapting. CR-Cross consider a triple (s, c, g) , where $s : \mathcal{X} \rightarrow \mathcal{Z}$ is a feature extractor, $c : \mathcal{Z} \rightarrow \mathcal{Y}$ is a classifier, and $g : \mathcal{X} \rightarrow \{0, 1\}$ is a reject rule. Specifically, g serves as a binary qualifier for s and c as follows:

$$(s, c, g)(x) = \begin{cases} c(s(x)) & g(x) = 1 \\ \text{reject} & g(x) = 0 \end{cases} \quad (1)$$

Denote the classifier $h = s \circ c$ as the joint variables of s and c . The learning problem of CR-Cross consists of selecting h and g such that the following error (risk) with respect to the target marine region is minimized:

$$\text{err}_{P_t}(h, g) = \mathbb{E}_{x \sim P_t} [\mathbb{I}(h(x) \neq y) \cdot g(x)] \quad (2)$$

where \mathbb{I} is the indicator function. The risk of the target marine region can be minimized by the reject rule, but the distribution discrepancy is still present. To reduce the distribution discrepancy among the different marine regions, a bound of cross-domain coral recognitions is proposed. The bound suggests that the target error is minimized by matching the marginal distributions among the different marine regions and minimizing the error on the source marine region simultaneously. With the incorporation of reject rule g in the cross-domain coral recognitions objective, the new loss function is defined:

$$\text{err}_{P_s}(h, g) + \text{disc}(P_s, P_t | g) \quad (3)$$

where err_{P_s} is the cross-entropy loss of the source marine region, $\text{disc}(P_s, P_t | g)$ is the discrepancy between source marine region and target marine region. The classifier h in Eq. 3 is trained by minimizing the loss function. Meanwhile, the loss function is determined by the rule g which

also needs to be trained. The reject rule plays a crucial role in training h : while abstaining the confusing coral images (e.g., bleached corals and coral samples with significant within-species morphological variations) as needed during the transfer process, both discriminability and transferability can be enhanced. That is, different categories are more easily to be separated, and distribution discrepancy is bridged more positively. Meanwhile, the classifier helps to find g . We propose the following learning procedure:

Step 1: Train classifier h using the labeled samples from the source marine region and unlabeled samples from the target marine region.

Step 2: Find the optimal reject rule g based on classifier h .

To further facilitate the marginal distribution alignment, we integrate the MDD into the $\text{disc}(P_s, P_t | g)$. MDD considers $f : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ with hypothesis space \mathcal{F} , where $f(x, y)$ indicates the component of $f(x)$ corresponding to the label y . Thus $f(x)$ induces a classifier h_f from \mathcal{X} to \mathcal{Y} in a hypothesis space \mathcal{H} :

$$h_f : x \rightarrow \arg \max_{y \in \mathcal{Y}} f(x, y). \quad (4)$$

The margin of the hypothesis f at a labeled example (x, y) is defined as

$$\rho_f(x, y) = \frac{1}{2} \left(f(x, y) - \max_{y' \neq y} f(x, y') \right) \quad (5)$$

The corresponding margin loss of f is

$$\text{err}_P^\rho(f) = \mathbb{E}_{x \sim P} \Phi_\rho(\rho_f(x, y)) \quad (6)$$

where Φ_ρ is

$$\Phi_\rho(x) = \begin{cases} 0 & \rho \leq x \\ 1 - x/\rho & 0 \leq x \leq \rho \\ 1 & x \leq 0 \end{cases} \quad (7)$$

Based on the margin loss, the expression of MDD is

$$d_{f, \mathcal{F}}^\rho(P_s, P_t) = \sup_{f' \in \mathcal{F}} (\text{disp}_{P_t}^\rho(f', f) - \text{disp}_{P_s}^\rho(f', f)), \quad (8)$$

where

$$\text{disp}_{P'}^{\rho}(f', f) = \mathbb{E}_{x \sim P} \Phi_{\rho}(\rho_{f'}(x, h_f))$$

Combining the Eq. 3 and Eq. 8, the optimization problem of step 1 can be written as

$$\min_f \text{err}_{P_s}^{\rho}(f, g) + d_{f, \mathcal{F}}^{\rho}(P_s, P_t | g),$$

where

$$\begin{aligned} \text{err}_{P_s}^{\rho}(f, g) &= \mathbb{E}_{x \sim P_s} [\Phi_{\rho}(\rho_f(x, y)) \cdot g(x)], \\ d_{f, \mathcal{F}}^{\rho}(P_s, P_t, g) &= \sup_{f' \in \mathcal{F}} (\text{disp}_{P_t}^{\rho}(f', f, g) - \text{disp}_{P_s}^{\rho}(f', f, g)) \\ &= \sup_{f' \in \mathcal{F}} (\mathbb{E}_{x \sim P_t} [\Phi_{\rho}(\rho_{f'}(x, h_f)) \cdot g(x)] - \\ &\quad \mathbb{E}_{x \sim P_s} [\Phi_{\rho}(\rho_{f'}(x, h_f)) \cdot g(x)]). \end{aligned} \quad (9)$$

Eq. 9 can be easily trained end-to-end by a minimax optimization procedure through standard back-propagation.

Optimal Reject Rule

The problems of target risk and distribution discrepancy among the different marine regions have been addressed, we need to search for the optimal reject rule g in step 2. That is, we need to find g_{na} using the non-abstained samples ($g^r = g^{r-1} \cdot g_{na}^r$). In information theoretic learning, the normalized mutual information (NMI) (Finn 1993) measures the degree of dependence between the targets T and the decision outputs Y :

$$NMI(T, Y) = \frac{I(T, Y)}{H(T)}, \quad (10)$$

where $I(T, Y)$ is the mutual information of T and Y , and $H(T)$ is the Shannon's entropy of T . NMI in Eq. 10 restricts the mutual information score to the range $[0, 1]$. The cost information about the error and reject can be derived from NMI in abstaining classification, which is determined by the data distributions (Hu 2013). Therefore, NMI provides objectivity in getting the optimal error-reject trade-off. In a k -class abstaining classification (the rejected class is denoted as 0), the empirical estimation of NMI is defined as:

$$\begin{aligned} NMI &= \frac{\sum_{i=1}^k \sum_{j=1}^k P(T=i, Y=j) \log \frac{P(T=i, Y=j)}{P(T=i)P(Y=j)}}{\sum_{i=1}^k P(T=i) \log P(T=i)} \\ &= \frac{\sum_{i=1}^k \sum_{j=1}^k c_{ij} \log \left(\frac{c_{ij}}{C_i \sum_{i=1}^k \frac{c_{ij}}{n}} \right)}{\sum_{i=1}^k C_i \log \frac{C_i}{n}} \end{aligned} \quad (11)$$

where c_{ij} is the number of the instances that belongs to the i th class classified as the j th class, $C_i = \sum_{j=0}^k c_{ij}$ and $n = \sum_{i=1}^k \sum_{j=0}^k c_{ij}$. In the cross-domain coral recognitions setting, denote z as the number of remaining non-abstained coral samples from the source marine region, and the output label vector with respect to the coral sample is defined as $\mathbf{h}^s = [h_1^s, h_2^s, \dots, h_z^s]$:

$$h_i^s = \left(\arg \max_l f(\mathbf{x}_i^s, l) \right) \cdot g_{na}(\mathbf{x}_i^s) \quad (12)$$

where $h_i^s \in \{0, 1, \dots, k\}$, $f(\mathbf{x}_i^s, l)$ is the probabilistic output of class $l \in \{1, \dots, k\}$ for x_i^s from the source marine regions, i.e., $\sum_{l=1}^k f(\mathbf{x}_i^s, l) = 1$ and $0 \leq f(\mathbf{x}_i^s, l) \leq 1$. The optimization problem of the reject rule g_{na} is formulated as:

$$\begin{aligned} &\max_{\tau} NMI(\mathbf{y}^s, \mathbf{h}^s) \\ g_{na}(\mathbf{x}_i^s) &= \begin{cases} 1 & \max \left(\frac{f(\mathbf{x}_i^s, l)}{\tau_l} \right) \geq 1 \\ 0 & \text{otherwise,} \end{cases} \quad (13) \\ &0 < \tau_l \leq 1, l = 1, 2, \dots, k, \end{aligned}$$

Finally, the optimal g_{na} learned based on real source labels is utilized to find newly rejected data from two marine regions in step 1. In the two-step procedure, we alternatively bridge the domain gap and improve the reject quality until convergence.

Experiments

In this section, we empirically demonstrate the effectiveness of the proposed method in cross-domain coral recognition.

Cross-domain coral recognition datasets

Due to the lack of a cross-domain underwater coral dataset, we propose a new cross-domain benchmark dataset from openly available EILAT datasets (Shihavuddin 2017) and CoralNet's website. Owing to the datasets collected from three different marine regions, we have named this dataset EILAT3. EILAT3 consists of three sub-datasets: E1, E2, and E3, which contain 1094 image patches of size 64×64 . Each sub-dataset has the same species, including the Urchin, Sand, Favid Coral, Branching Coral, and Brain Coral. E1 is composed of image patches cropped from coral images. These coral patches obtain the texture feature of different parts of the coral image and do not capture the global structure feature of the coral image. E2 is collected from different marine regions near Eilat in the Red Sea and consists of coral images with the global structure. E3 is a dataset obtained by scraping data from CoralNet's website, which has coral images with different marine regions and various growth cycles. We construct 6 cross-domain underwater coral recognition tasks by using the three different marine region combinations.

Compared Methods

In cross-domain coral recognition, we systematically compare the proposed method with the following state-of-the-art coral recognition methods. The compared methods can be roughly categorized into two groups. The first group is composed of coral recognition methods without any transfer learning, including ResNet50 and ResNet101 (He et al. 2016). Due to the outstanding performance of the ResNet, most previous work (Beijbom et al. 2012; Gómez-Ríos et al. 2019; Raphael et al. 2020) utilize the different variations of ResNet to demonstrate the state-of-the-art coral recognition accuracies. The second group consists of previous transfer learning methods including DAN (Long et al. 2015), DANN (Ganin and Lempitsky 2015), JAN (Long et al. 2017), CDAN (Long et al. 2018), MCD (Saito et al. 2018), MDD (Zhang et al. 2019) and CGDM (Du et al. 2021).

Table 1: Accuracy(%) on coral dataset for cross-domain coral recognition

Method	E1→E2	E1→E3	E2→E1	E2→E3	E3→E1	E3→E2
ResNet50	78.9	43.8	97.6	84.2	90.4	93.8
ResNet101	79.9	44.2	97.5	82.9	94.0	91.1
DANN	80.9	56.2	96.9	88.3	74.9	90.5
DAN	84.5	67.1	97.6	86.7	91.3	94.7
JAN	69.1	52.1	93.6	75.0	90.5	88.5
MCD	78.0	52.9	97.1	80.8	65.5	94.7
CDAN	84.2	58.3	96.2	87.9	80.0	91.8
CGDM	79.4	64.8	90.6	87.8	58.3	91.5
MDD	83.9	57.1	97.5	87.5	75.5	95.1
CR-Cross(ours)	88.2	77.8	98.4	98.6	95.7	98.7

Results and Discussion

During the process of cross-domain coral recognition, we repeat the learning process until the maximum reject rate p for the target coral data is reached and set p is 10%. From table 1, several observations can be made. (1) Due to the domain shift among the different marine regions, the coral recognition methods without transfer learning including the ResNet50 and ResNet101 have lower accuracy. Specifically, for the transfer task from E1 to E3, the accuracy of ResNet50 is only 43.8%. The larger models also can not reduce the domain shift, and the cross-domain coral recognition accuracy of ResNet101 is only 44.2%. (2) On the previous transfer learning methods, the domain shift can be slightly reduced in some transfer tasks. For the transfer learning from E2 to E3, DANN and CDAN gained 4.1% and 3.7% improvement, respectively. However, JAN and CGDM perform even worse than the non-transfer learning model impacted by the negative transfer. (3) Depending on the fusion of the texture feature and the structure feature and the optimal reject strategy, our method achieves the best accuracy on all transfer tasks. Compared to ResNet50, CR-Cross achieved improvements of 9.3%, 34%, 0.8%, 14.4%, 5.3%, and 4.9% on six cross-domain coral recognition tasks. Different from the previous transfer learning methods which rely on general information to complete the cross-domain coral recognition, CR-Cross transfers the structure and texture information of corals from one marine region to another. In addition, the optimal reject rule can further eliminate the transfer of useless or even negative information.

Qualitative Analysis

Feature Visualization: To provide an intuitive understanding of the aligned feature of the source marine region and target marine region, we utilize the t-SNE to visualize the feature of different methods. We conduct the experiment on the transfer task from E2 to E3. As shown in Figure 2, we can derive three conclusions. As shown in Figure 2(a) and (b), the source coral samples have a clear classification boundary and obvious discriminative structure. Due to the existence of domain shifts, different marine regions have a huge distribution gap. Figure 2 (c), 2 (d), and 2 (e) show that DAN, DANN, and MDD can slightly reduce the domain discrepancy among the different marine regions. However, al-

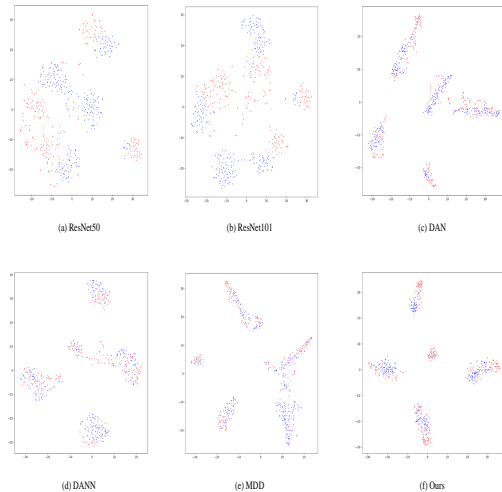


Figure 2: The t-SNE visualization of features generated by ResNet50, ResNet101, DAN, DANN, MDD, and Ours from different marine regions. Red and blue dots represent the source domain and the target domain, respectively.

though the distribution discrepancy of the DAN, DANN, and MDD is reduced, the classification boundary and discriminative structure of the target marine region are relatively poor. Relying on the fusion of the texture feature and the structure feature and the optimal reject strategy, our method has a clear classification bound and nearly similar feature distributions among the source marine region and target marine region. In addition, our method maximally aligns the discriminative structure and realizes the intra-class compactness between the source marine region and the target marine region. The results demonstrate the effectiveness and feasibility of our method.

Conclusion

In this paper, we propose a novel CR-Cross framework, which can utilize both coral’s texture features, coral’s structure features, and the optimal reject rule to achieve cross-domain coral recognitions. CR-Cross applies the texture and

structure feature extractor to obtain transferable knowledge from the source marine region to realize the coral recognition in the target marine region. To overcome the difficulty that the confusing coral images (e.g., bleached corals) are prone to be misclassified and transfer useless or even negative information, our algorithm is equipped with the reject option for the confusing corals while adapting, which further reduces the distribution discrepancy among the different marine regions. Cross-domain coral recognition goes a step further in tackling the challenges of missing coral data, maximizing the utilization of available coral datasets, and enhancing the reusability of both coral data and coral recognition models. In addition, CR-Cross can be easily integrated into almost all coral recognitions. A series of empirical studies show that our method remarkably outperforms a broad range of baselines, which offers exciting potential for large-scale cross-domain coral conservation.

References

- Asha Paul, M.; Arockia Jansi Rani, P.; and Liba Manopriya, J. 2020. Gradient based aura feature extraction for coral reef classification. *Wireless Personal Communications*, 114: 149–166.
- Beijbom, O.; Edmunds, P. J.; Kline, D. I.; Mitchell, B. G.; and Kriegman, D. 2012. Automated annotation of coral reef survey images. In *2012 IEEE conference on computer vision and pattern recognition*, 1170–1177. IEEE.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Du, Z.; Li, J.; Su, H.; Zhu, L.; and Lu, K. 2021. Cross-domain gradient discrepancy minimization for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3937–3946.
- Finn, J. T. 1993. Use of the average mutual information index in evaluating classification error and consistency. *International Journal of Geographical Information Science*, 7(4): 349–366.
- Ganin, Y.; and Lempitsky, V. 2015. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, 1180–1189. PMLR.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1): 2096–2030.
- Gómez-Ríos, A.; Tabik, S.; Luengo, J.; Shihavuddin, A.; Krawczyk, B.; and Herrera, F. 2019. Towards highly accurate coral texture images classification using deep convolutional neural networks and data augmentation. *Expert Systems with Applications*, 118: 315–328.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hu, B.-G. 2013. What are the differences between Bayesian classifiers and mutual-information classifiers? *IEEE transactions on neural networks and learning systems*, 25(2): 249–264.
- Long, M.; Cao, Y.; Wang, J.; and Jordan, M. 2015. Learning transferable features with deep adaptation networks. In *International conference on machine learning*, 97–105. PMLR.
- Long, M.; Cao, Z.; Wang, J.; and Jordan, M. I. 2018. Conditional adversarial domain adaptation. *Advances in neural information processing systems*, 31.
- Long, M.; Zhu, H.; Wang, J.; and Jordan, M. I. 2017. Deep transfer learning with joint adaptation networks. In *International conference on machine learning*, 2208–2217. PMLR.
- Mahmood, A.; Bennamoun, M.; An, S.; Sohel, F.; Boussaid, F.; Hovey, R.; Kendrick, G.; and Fisher, R. B. 2016. Coral classification with hybrid feature representations. In *2016 IEEE International Conference on Image Processing (ICIP)*, 519–523. IEEE.
- Mary, N. A. B.; and Dharma, D. 2017. Coral reef image classification employing improved LDP for feature extraction. *Journal of Visual Communication and Image Representation*, 49: 225–242.
- Mehta, A.; Ribeiro, E.; Gilner, J.; and van Woesik, R. 2007. Coral reef texture classification using support vector machines. In *VISAPP (2)*, 302–310.
- Mora, C.; Andréfouët, S.; Costello, M. J.; Kranenburg, C.; Rollo, A.; Veron, J.; Gaston, K. J.; and Myers, R. A. 2006. Coral reefs and the global network of marine protected areas.
- Mulhall, M. 2008. Saving the rainforest of the sea: An analysis of international efforts to conserve coral reefs. *Duke Envtl. L. & Pol’y F.*, 19: 321.
- Raphael, A.; Dubinsky, Z.; Iluz, D.; Benichou, J. I.; and Netanyahu, N. S. 2020. Deep neural network recognition of shallow water corals in the Gulf of Eilat (Aqaba). *Scientific reports*, 10(1): 12959.
- Reaka-Kudla, M. L. 1997. The global biodiversity of coral reefs: a comparison with rain forests. *Biodiversity II: Understanding and protecting our biological resources*, 2: 551.
- Saito, K.; Watanabe, K.; Ushiku, Y.; and Harada, T. 2018. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3723–3732.
- Shihavuddin, A. 2017. Coral reef dataset. v2. *Mendeley data* <https://data.mendeley.com/datasets/86y667257h/2>.
- Stokes, M. D.; and Deane, G. B. 2009. Automated processing of coral reef benthic images. *Limnology and Oceanography: Methods*, 7(2): 157–168.
- Zellinger, W.; Grubinger, T.; Lughofer, E.; Natschläger, T.; and Saminger-Platz, S. 2017. Central moment discrepancy (cmd) for domain-invariant representation learning. *arXiv preprint arXiv:1702.08811*.
- Zhang, Y.; Liu, T.; Long, M.; and Jordan, M. 2019. Bridging theory and algorithm for domain adaptation. In *International conference on machine learning*, 7404–7413. PMLR.