Evaluating domain adaptation for underwater image analysis

Varshani Brabaharan, University of Ruhuna, Sri Lanka

Mentors: Kakani Katija, Eric Orenstein, Joost Daniels, Paul Roberts, and Kevin Barnard

Summer 2023

Keywords: Distribution shift, domain adaptation; ADAPT; unsupervised;

ABSTRACT

In the realm of underwater image analysis, handling out-of-distribution data poses significant challenges. This study delves into the application of domain adaptation techniques to enhance model performance across varied underwater imaging scenarios. Data collected from Remotely Operated Vehicles (ROVs) and Autonomous Underwater Vehicles (AUVs) served as the foundation. The research employed the ADAPT toolkit, emphasizing techniques like MCD, MDD, and DeepCORAL in the context of a ResNet50 model. While the model achieved high accuracy on in-distribution ROV data, its performance suffered on out-of-distribution AUV datasets. However, with domain adaptation, considerable improvements were observed, with DeepCORAL emerging as the most impactful. The findings underscore the pivotal role of domain adaptation in achieving robust underwater image analysis, emphasizing the complexities introduced by environmental and technological variations.

1. INTRODUCTION

1.1. UNDERWATER IMAGE ANALYSIS

Underwater image analysis has become an increasingly important field of study, driven by the need for comprehensive marine research, resource management, and environmental conservation (Han et al., 2020). The unique challenges of this domain stem from the inherent complexities of underwater environments. Factors such as variable lighting conditions, water turbidity, and the diverse range of marine organisms contribute to the intricate nature of image data collected from underwater settings (Han et al., 2020). Advanced machine learning models have been employed to automate the identification and classification of marine species (Beijbom et al., 2012), assess coral reef health (González-Rivero et al., 2020), and even map underwater topographies (Salas & Argialas, 2022).

However, the effectiveness of these machine learning models is heavily influenced by the quality and representativeness of the training data. Often, models are trained on data collected from controlled or relatively stable environments, such as shallow waters with good visibility. When these models are deployed in more challenging or less-studied environments, their performance can be significantly impacted (Er et al., 2023).

1.2. DISTRIBUTION SHIFT

Distribution shift refers to the phenomenon where the statistical properties of the data on which a machine learning model is trained (known as the source distribution) differ from those of the data on which the model is deployed or tested (known as the target distribution) (Quionero-Candela et al., 2009). This discrepancy can manifest in various ways, such as changes in the underlying features, class imbalances, or even alterations in the relationships between features and labels. Distribution shift is a significant concern in machine learning because models are generally optimized to perform well on the source distribution; when exposed to a different target distribution, their performance can degrade, sometimes dramatically (Quionero-Candela et al., 2009).

In the realm of deep-water image analysis, distribution shifts can have a profound impact on the performance of machine learning models. For instance, consider a model trained on images of marine organisms captured at relatively shallow depths, where natural light is abundant. This model is optimized to recognize features such as color patterns, shapes, and textures under well-lit conditions. However, when the same model is deployed to analyze images from greater depths, where the environment is darker and visibility is reduced, its performance may deteriorate significantly. At these depths, the color of organisms can appear distorted, and the lack of light

may result in less distinguishable textures and shapes. Additionally, the presence of bioluminescent organisms or floating sediment particles could introduce new, unexpected features that the model is not trained to handle. In such scenarios, the distribution shift between the source (shallow water images) and target (deep water images) can lead to high rates of misclassification or false detections, thereby compromising the reliability of the analysis (Jian et al., 2021).
Therefore, understanding and addressing distribution shifts is crucial for developing machine learning models that are robust and generalizable, especially when deploying models in real-world scenarios of the deep ocean where the data distribution can be dynamic and unpredictable.

1.3. TRANSFER LEARNING AND DOMAIN ADAPTATION

Transfer learning and domain adaptation (DA) are foundational techniques in machine learning that aim to bridge the gap between the training (source) and deployment (target) data distributions. These methodologies have gained considerable traction and undergone significant advancements, particularly because they address the ubiquitous issue of "domain shifts" in real-world applications. In the specialized realm of deep-water organism image classification, the utility of these techniques becomes even more pronounced (de Mathelin et al., 2021).

In such complex and variable underwater conditions, the integration of domain adaptation algorithms becomes not just advantageous but essential. These algorithms enable the development of robust, accurate, and adaptable classification models that can effectively navigate the multifaceted challenges presented by deep-sea ecosystems. By addressing these domain shifts, we can build models that are not only theoretically sound but also practically applicable for real-world deep-water research and conservation efforts (Fujimura et al., 2020).

1.4. ROVS AND AUVS IN UNDERWATER IMAGE COLLECTION

Remotely Operated Vehicles (ROVs) and Autonomous Underwater Vehicles (AUVs) have revolutionized the field of underwater image analysis, particularly in the study of marine animals.

These advanced instruments provide a means to explore and capture high-resolution images in environments that are often inaccessible or hazardous for human divers.

1.4.1. ROVs

ROVs are typically tethered vehicles controlled by operators on a surface vessel. They are often equipped with powerful lighting and high-definition cameras, allowing for the capture of detailed images even in the darker regions of the ocean. ROVs are especially useful for targeted studies where real-time decision-making is required, such as close-up inspections of specific marine habitats or the tracking of individual animals (Christ & Wernli, 2014).

1.4.2. AUVs

Unlike ROVs, AUVs operate without a tether and are pre-programmed to navigate through underwater environments autonomously. They are often deployed for broader, large-scale surveys and can cover extensive areas in a single mission. AUVs are equipped with a range of sensors and imaging devices, enabling them to capture a diverse array of data types, from simple photographs to complex 3D reconstructions of underwater landscapes (Bellingham & Rajan, 2007).

1.4.3. DOMAIN SHIFTS IN ROV AND AUV IMAGES

While both ROVs and AUVs are invaluable tools for underwater image collection, the images they capture can exhibit significant domain shifts, complicating the task of image analysis. For instance, ROVs, being manually operated, often capture images under varying lighting conditions depending on operator control, leading to inconsistencies in image quality. AUVs, on the other hand, may capture images at different times of day or at varying depths, resulting in natural lighting changes that affect image characteristics (Jian et al., 2021).

Moreover, the types of cameras and lenses equipped on these vehicles can also differ, leading to variations in image resolution, field of view, and color representation. Such discrepancies introduce domain shifts that can severely impact the performance of machine learning models trained on one type of data but deployed on another (Fujimura et al., 2020).

Given these challenges, it becomes crucial to employ domain adaptation techniques to ensure that the image analysis models can generalize well across different types of data collected by ROVs and AUVs. This is essential for the reliable study and conservation of marine animals, as well as for the broader applicability of these advanced underwater imaging technologies.

1.5. ADAPT - AWESOME DOMAIN ADAPTATION PYTHON TOOLBOX

ADAPT is a robust, open-source Python package engineered to seamlessly incorporate domain adaptation strategies into computational models (de Mathelin et al., 2021). This package is organized into three user-friendly modules, each targeting a key domain adaptation strategy: feature-based, instance-based, and parameter-based. These modules are highly versatile, designed to seamlessly integrate with both supervised and unsupervised learning paradigms. As such, they provide a comprehensive and adaptable toolkit for tackling a wide array of domain adaptation challenges. This includes addressing the complexities introduced by distribution shifts commonly encountered in deep-water image analysis models.

1.6. OBJECTIVES

- Investigating the implementation of domain adaptation strategies using the ADAPT toolkit for enhancing underwater image analysis.
- Assessing the efficacy of these fine-tuned models in classifying marine organisms (*Eusergestes similis*) under conditions of distribution shift.

2. METHODOLOGY

2.1. DATA COLLECTION

The data for this project were gathered by different types of underwater exploration vehicles: Remotely Operated Vehicles (ROVs) and Autonomous Underwater Vehicles (AUVs). The ROV data were collected during missions led by the Doc Ricketts ROV. This dataset is quite extensive, featuring 23,736 images along with 36,079 expertly annotated bounding boxes that identify 21 unique marine species (Figure 1, 2). Nevertheless, only less than 3% of these annotated instances were of the target organism *Eusergestes similis* (Table 1). Conversely, the AUV datasets were acquired through horizontal video transects executed by the i2map AUV. Each transect was conducted over a 10-minute span at a near constant speed of approximately 0.5 m/s. These transects were carried out at two specific depths— ca. 200 meters and ca. 400 meters, respectively. Unlike the multi-species focus of the ROV dataset, the AUV transects were exclusively aimed at the identification and localization of the target species, *Eusergestes similis* (Table 2, Figure 3).

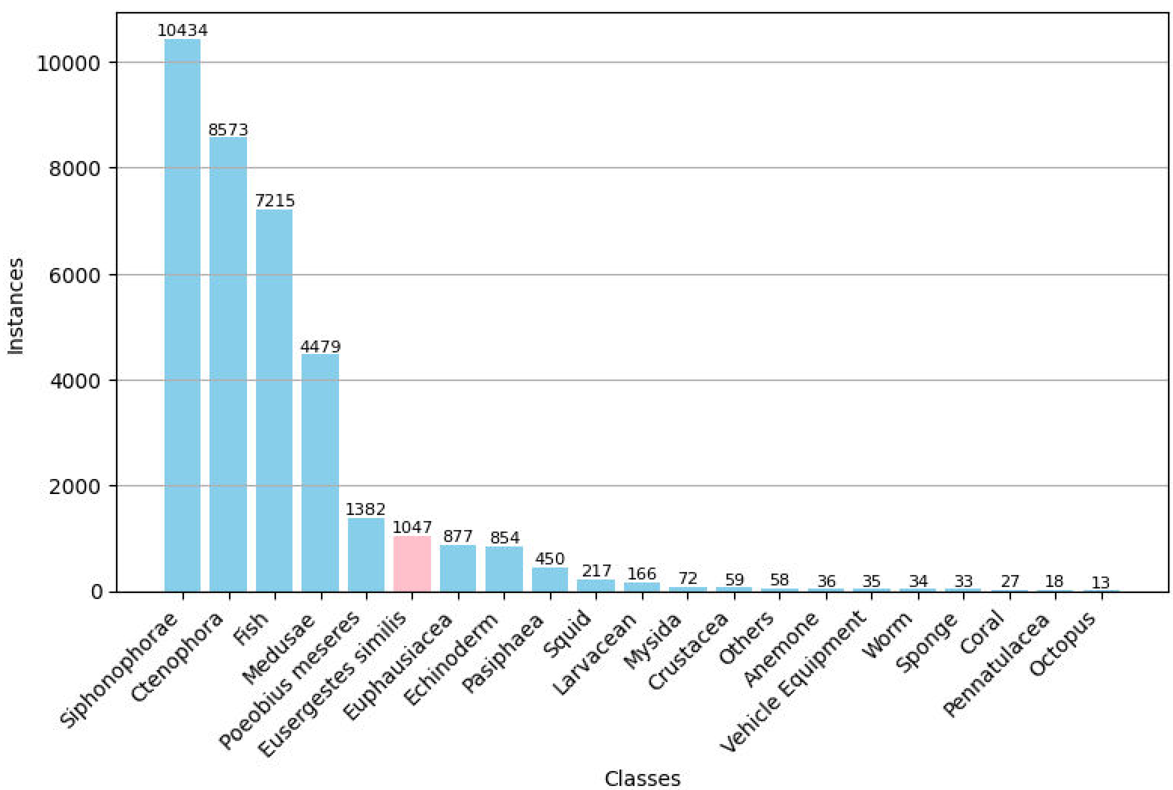

Figure 1: Distribution of data instances across 21 classes in the ROV dataset

Figure 2: The images shown above were collected during missions conducted by the Doc Ricketts ROV, providing a snapshot of few of the unique marine taxa identified.

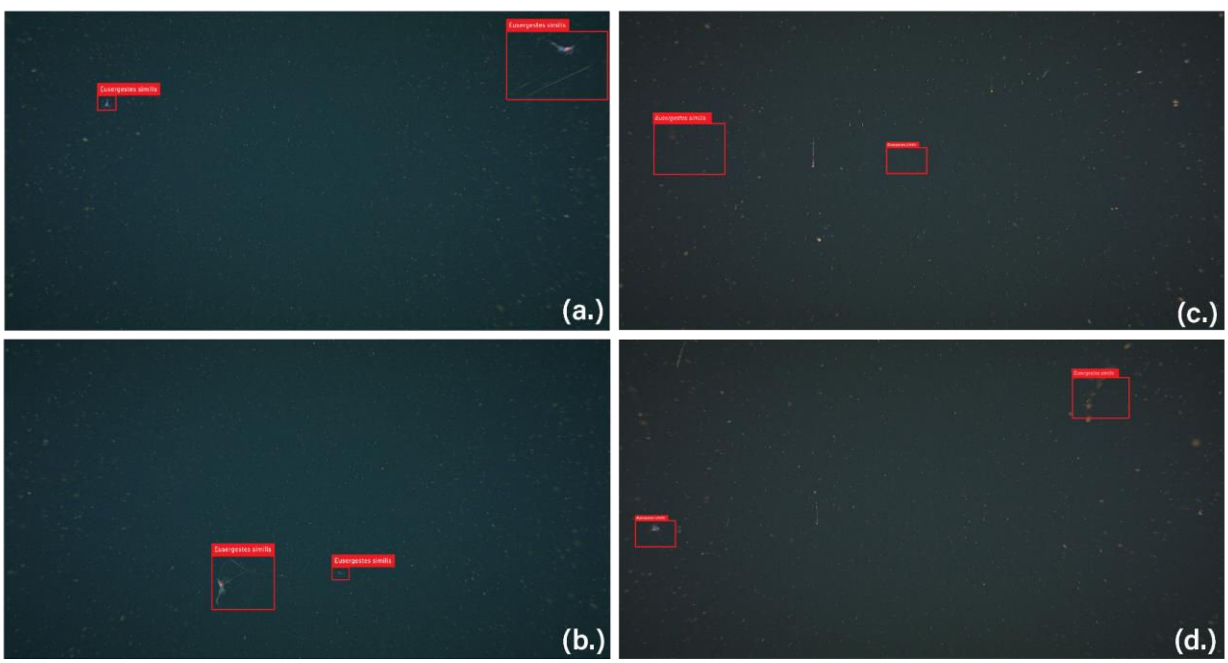

Figure 3: Figures (a) and (b) represent images from the AUV transects at a depth of approximately 200 m, while Figures (c) and (d) present captures from the 400 m depth transects.

These images serve to highlight the distinct lighting conditions encountered at each respective depth.

Table 1: Percentages of *Eusergestes similis* in the ROV data

.

Category	Percentage
Images	4.05%
Bounding boxes	2.90%

Table 2: Sample and instance counts of target species *Eusergestes similis* for AUV datasets at 200 m and 400 m depths.

Depth (m)	Samples	Instances
200	529	1447
400	150	157

2.2. MODEL EXPERIMENTATION

2.2.1. DATA PREPARATION

In this phase, the Regions Of Interest (ROI) were extracted from images captured by both ROV and AUV platforms. For the ROV (source) images, these regions included 21 distinct categories, including the target species of interest (Figure 1). Meanwhile, the AUV (target) images were exclusively centered on the target species, resulting in a single-category set of regions (Table 2). These curated areas served as the foundation for subsequent stages of the project.

2.2.2. DATA PREPROCESSING

A ResNet50 model trained on ImageNet served as the foundational learning model (Figure 4). Rather than fine-tuning the entire ResNet50, only the last block was specifically adapted, a strategy that effectively reduced computation time. During the preprocessing stage, features corresponding to the output of the initial blocks of the ResNet50 were extracted. These features, which were fixed, were subsequently used as input for the model. Consequently, two matrices, denoted as Xs (source) and Xt (target), were obtained for the two distinct domains—ROV and AUV.

Additionally, a custom loading function was designed for managing the last block of the ResNet50 model. In this function, the 'trainable' parameter for the BatchNormalization layer was intentionally set to 'False'. This measure was implemented to avert potential challenges during the fine-tuning stage, such as model instability or overfitting to the training data. By taking this precaution, we ensured that the model remained robust and generalized well, aligning perfectly with the earlier steps of feature extraction.

The labels were one-hot encoded as an additional preprocessing step, which was done to convert the categorical labels into a format that could be more easily processed by the machine learning algorithms. Specifically, one-hot encoding transforms each categorical label into a binary vector, allowing for a more straightforward and efficient representation of the classes. By incorporating this encoding method, we eliminated potential numerical hierarchies or biases that standard integer labeling might introduce, thereby facilitating a more accurate model training process.

Building upon the foundational ResNet50 model, a specialized task network was designed to classify the data into 21 distinct classes. This network was positioned subsequent to the ResNet50's last block and incorporated additional regularization features to enhance performance. Notably, beyond implementing dropout techniques for model generalization, we also imposed a constraint on the norm of the network's weights. This constraint proved essential for the successful training of the domain adaptation methods that would be employed later in the project. For optimization, a Stochastic Gradient Descent (SGD) algorithm was utilized. Additionally, a custom learning rate decay strategy, termed 'MyDeepOceanDecay,' was implemented to adjust the learning rate during training dynamically. These augmentations to the modeling architecture were meticulously crafted to optimize both accuracy and computational efficiency, offering a robust and streamlined approach for the classification tasks across the ROV and AUV domains.

2.2.3. MODEL FITTING WITHOUT DOMAIN ADAPTATION

Fine-tuning of the ResNet50 model was conducted solely on the source (ROV) data, without incorporating the target (AUV) data. The fine-tuning function was parameterized with several components including, loss function selected as categorical cross-entropy, while the evaluation metric utilized was accuracy. Upon configuration, the model was trained on the source data Xs and labels ys,lab for 100 epochs with a batch size of 32. Concurrently, validation was conducted using the target data Xt and its corresponding labels yt,lab. This focused approach on the source data ensured that the model was specifically attuned to the characteristics of the domain for which we had labeled information, laying a strong foundation for subsequent stages involving domain adaptation techniques. Model inference was conducted on distinct subsets of the datasets, specifically: a reserved split of the ROV dataset that was not employed during training, as well as the AUV datasets at both 200 m and 400 m depths. This approach provided a comprehensive evaluation across varying environmental conditions (varying domains) and ensured that the model's generalization capabilities were rigorously tested.

2.2.4. MODEL FITTING WITH DOMAIN ADAPTATION

To incorporate domain adaptation into our computational model, we utilized the ADAPT toolkit. Within the scope of the project, we experimented with three feature-based unsupervised methods—Maximum Classifier Discrepancy (MCD) (Saito et al., 2018), Margin Disparity Discrepancy (MDD) (Zhang et al., 2019), and Deep CORrelation ALignment (DeepCORAL) (Sun & Saenko, 2016) (Figure 4).

MCD, aims to minimize the feature discrepancy between source and target domains through adversarial training involving three networks: an encoder and two classifiers. A reversal layer is inserted between the encoder and classifiers to facilitate this adversarial process, ultimately reducing the domain gap. The method was parameterized with loss function selected as categorical cross-entropy, while the evaluation metric utilized was accuracy.

Similar to MCD, MDDs objective is to minimize the feature discrepancy between the source and target domains. However, the key difference lies in the network architecture and the metric for estimating this discrepancy. MDD employs an encoder, a task network, and a discriminator for adversarial training. While MCD focuses on minimizing classifier discrepancy through dual classifiers, MDD aims to minimize the disparity discrepancy using a task network and a discriminator. This subtle yet crucial difference allows MDD to approach domain adaptation from a slightly different angle, offering an alternative method for reducing the domain gap. In our implementation of the MDD (Margin Disparity Discrepancy) technique, we utilized specific hyperparameters to fine-tune the model's performance. Specifically, we set the lambda parameter as a TensorFlow variable initialized at 0, denoted as lambda_ = tf.Variable(0.). This parameter serves as a weighting factor in the loss function to balance the contributions from the source and target domains. Additionally, we employed a gamma value of 2 (gamma = 2.), which is used to modulate the impact of the discrepancy term in the model's objective function.

DeepCORAL is the third feature-based domain adaptation method we employed, and it extends the original CORAL method. It uses an encoder and a task network to learn a nonlinear transformation that aligns the correlation matrices of layer activations in deep neural networks. The encoder maps input features into a new space where the task network is trained. The objective is to minimize a loss function that considers both task loss on the source data and the Frobenius norm between the correlation matrices of the source and target data. In our DeepCORAL implementation, the lambda parameter was set to a default of 1, which serves as a trade-off factor in the loss function. This default value implies equal importance between task-specific loss and domain alignment loss, providing a balanced starting point for initial experiments. However, the optimal value may vary and could require fine-tuning for specific applications in underwater image analysis.

2.2.5. MODEL VALIDATION

The baseline ROV model, which did not incorporate any domain adaptation techniques, was initially validated on both ROV and AUV validation sets (Table 3, Figure 4). Subsequently, models fine-tuned using domain adaptation methods were specifically evaluated against AUV validation sets (Table 3, Figure 4). To assess the performance of the models, we utilized TensorFlow's “predict” function and visualized the results using t-SNE representations.

Table 3: Count of target species *Eusergestes similis* specimens in the validation set

Validation set	Sample count	Instances count
ROV	4542	184
AUV – 200 m	106	151
AUV – 400 m	150	157

The use of the predict function allows for a straightforward and efficient way to generate model predictions on new, unseen data. It is particularly useful for obtaining probability scores or class labels, which can then be compared to the actual labels for performance assessment. This method provides a consistent and scalable approach to model evaluation, making it a suitable choice for our analysis. Furthermore, t-SNE diagrams were employed because they provide an effective way to visualize high-dimensional data by reducing it to a two-dimensional space. This allows for a clear depiction of data clusters and separations, offering insights into the model's ability to distinguish between different classes or categories.

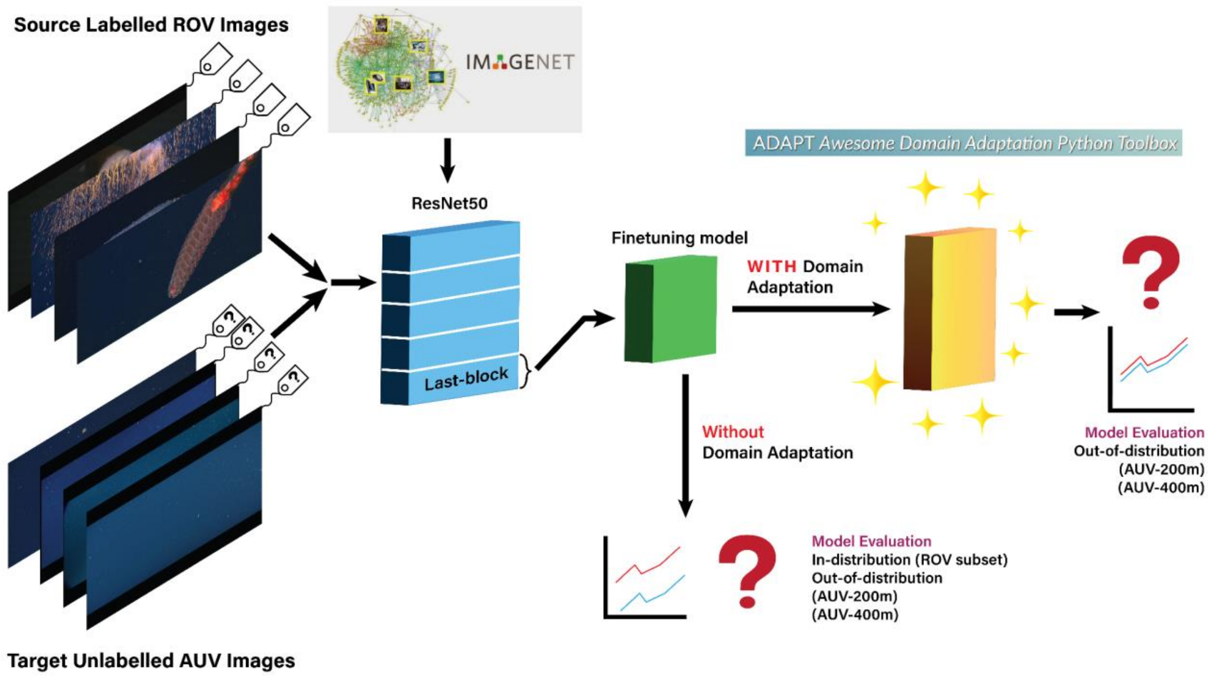

Figure 4: Schematic representation of the implemented model architecture.

3. RESULTS

In this study, we aimed to evaluate the efficacy of various domain adaptation techniques in managing out-of-distribution data in classifying small target organism, *Eusergestes similis*. Our comparison encompassed three distinct methods: MCD, MDD, and DeepCORAL. As a benchmark, we also examined the results when no domain adaptation was applied.

3.1. MODEL PERFORMANCE SCORES

3.1.1. BASELINE (NO DOMAIN ADAPTATION):

For the in-distribution data, which is a subset of ROV, we achieved an accuracy of 0.9845. This high performance is expected as models tend to perform well on data they are familiar with (Table 4, Figure 5).

When transitioning to out-of-distribution data, the accuracy significantly dropped. For the AUV-200m dataset, the accuracy was 0.7222, and for the AUV-400m dataset, it further declined to

0.5432 (Table 4). This decline underscores the challenges of managing out-of-distribution data without any domain adaptation.

3.1.2. MCD:

With the MCD method applied to the AUV-200m dataset, there was an accuracy of 0.759, representing an increase of approximately 3.7% compared to the no adaptation scenario. For the AUV-400m dataset, the accuracy was 0.5921, showcasing a roughly 4.9% enhancement from the baseline (Table 4, Figure 6).

3.1.3. MDD:

The MDD technique, when applied to the AUV-200m data, resulted in an accuracy of 0.8267. This is an improvement of approximately 10.5% from the no adaptation approach. The AUV-400m dataset yielded an accuracy of 0.6738, reflecting a notable increase of around 13% compared to the baseline (Table 4, Figure 7).

3.1.4. DeepCORAL :

DeepCORAL, when tested on the AUV-200m dataset, achieved an impressive accuracy of 0.8945, marking a substantial enhancement of about 17.2% over the no adaptation method. On the AUV-400m dataset, the accuracy was 0.7847, indicating a significant gain of approximately 24.1% compared to the baseline (Table 4, Figure 8).

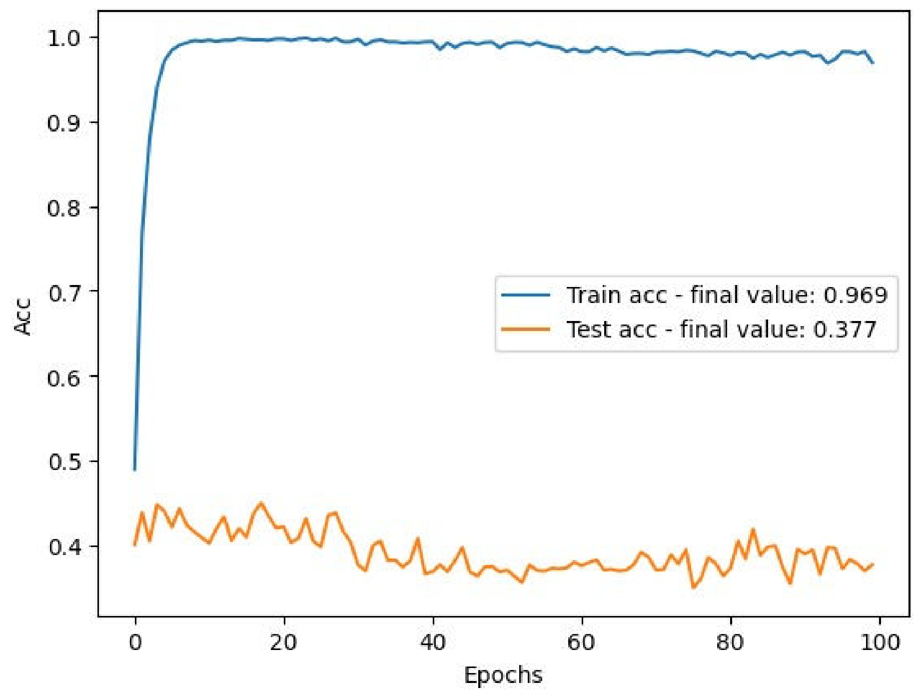

Figure 5: Evolution of Model Accuracy Across Epochs Without Domain Adaptation.

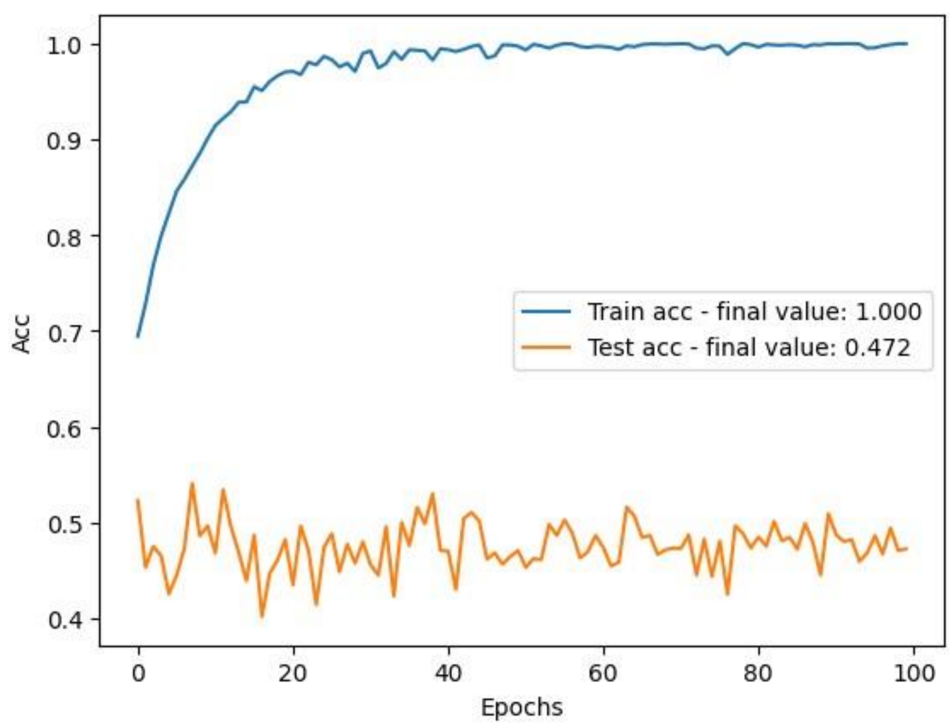

Figure 6: Evolution of Model Accuracy Across Epochs with MCD (Maximum Classifier Discrepancy) domain adaptation technique.

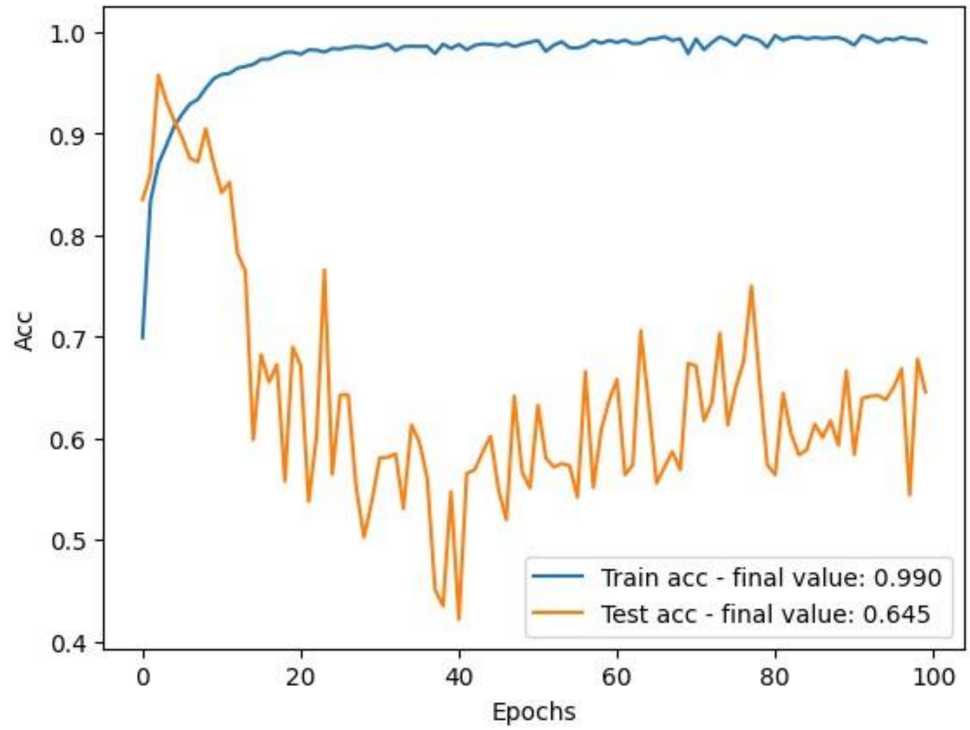

Figure 7: Evolution of Model Accuracy Across Epochs with Margin Disparity Discrepancy (MDD) domain adaptation technique.

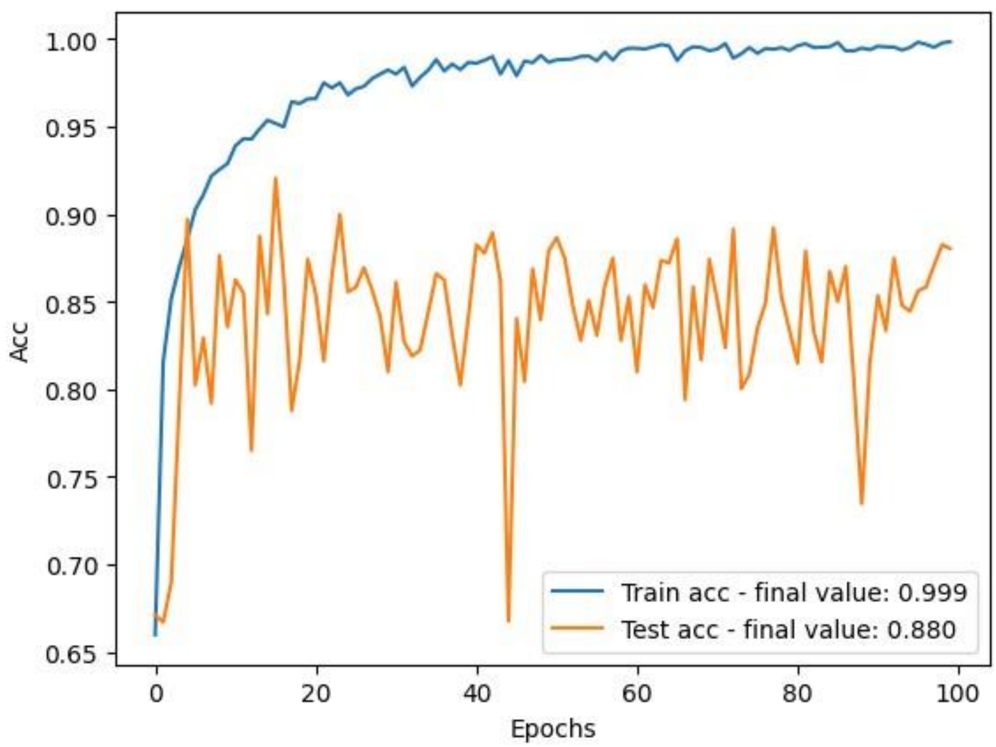

Figure 8: Evolution of Model Accuracy Across Epochs with Deep CORrelation ALignment (DeepCORAL) domain adaptation technique.

Table 4: Comparative inference scores as a percentage for models, with and without domain adaptation techniques.

	No domain adaptation	MCD	MDD	DeepCORAL
In-distribution (ROV subset)	98.45%	-	-	-
Out-of-distribution (AUV-200m)	72.22%	75.90%	82.67%	89.45%
Out-of-distribution (AUV-400m)	54.32%	59.21%	67.38%	78.47%

3.2. ENCODED FEATURE SPACE

For the baseline model, the t-SNE visualization of the encoded feature space for the target AUV data of *Eusergestes similis* displayed a random scattering (Figure 9). This randomness indicates that, without domain adaptation, the model struggled to differentiate or group the features of the target species effectively.

In stark contrast, the DeepCORAL model's t-SNE visualization revealed a much more organized feature space. The encoded features of the *Eusergestes similis* from the AUV data were not just randomly scattered but instead formed discrete clusters (Figure 10). This clustering suggests that the DeepCORAL adaptation technique significantly improved the model's ability to recognize and group similar features, enhancing its classification capabilities.

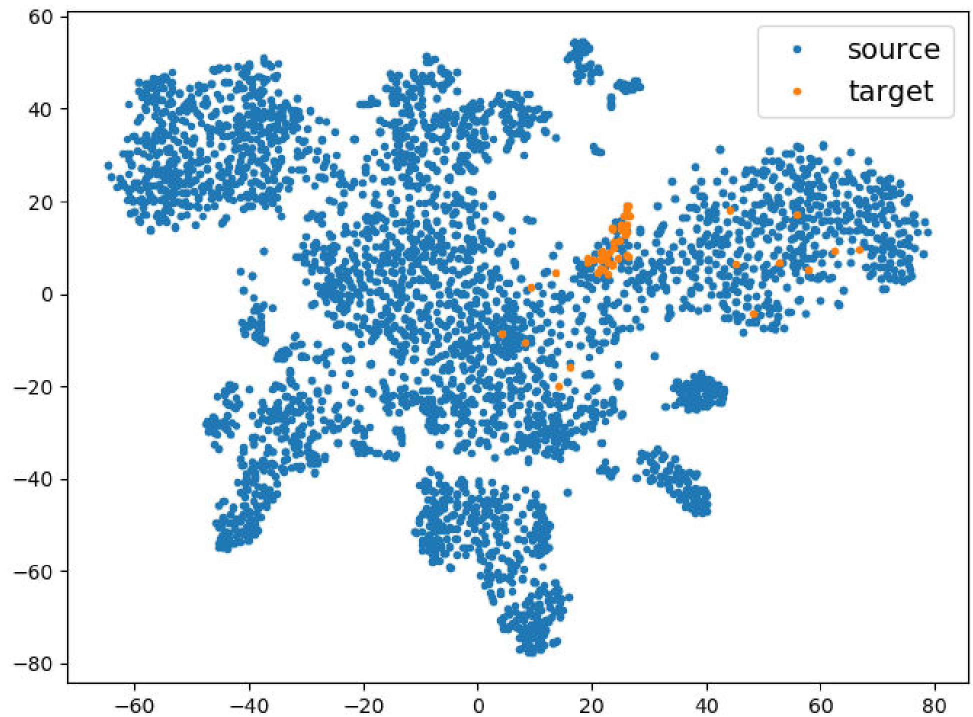

Figure 9: t-SNE Visualization of the Encoded Feature Space for the model without domain adaptation.

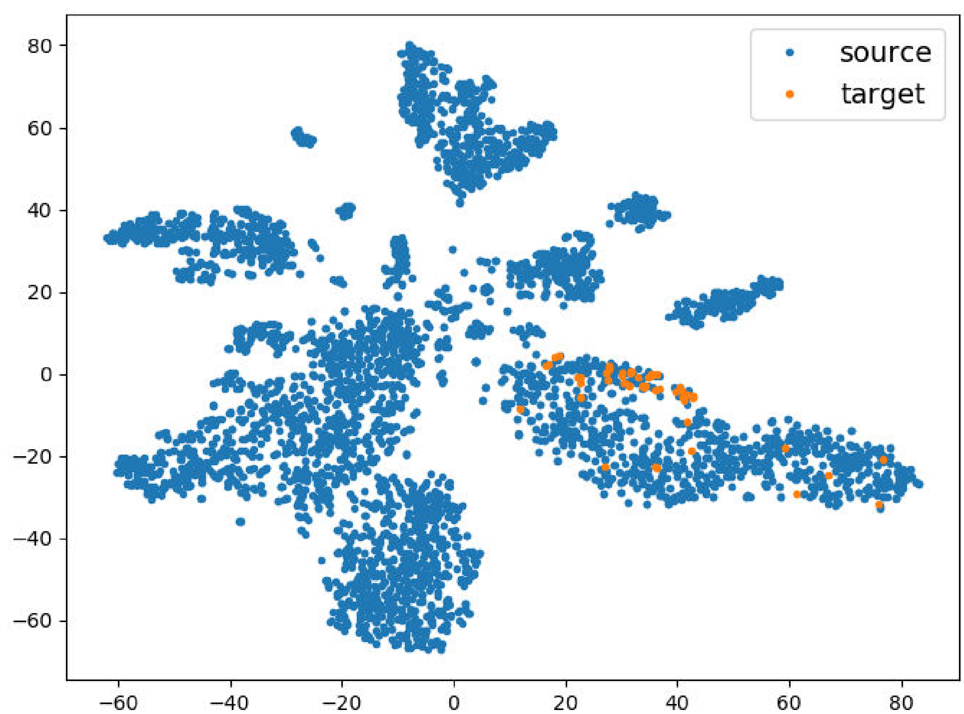

Figure 10: t-SNE Visualization of the Encoded Feature Space for the DeepCORAL Model

4. DISCUSSION

The results clearly demonstrate the benefits of employing domain adaptation techniques when dealing with out-of-distribution data (Table 4). All three methods (MCD, MDD, and DeepCORAL) outperformed the baseline, with DeepCORAL showing the most significant gains across both out-of-distribution datasets (Table 4). This suggests that for tasks where out-of-distribution data is a concern, investing in domain adaptation strategies can significantly boost performance.

One pivotal observation to discuss is the distinct performance variations between ROV images and AUV images taken at depths of 200m and 400m. Several factors can contribute to these differences:

- Environmental Factors: The deeper the ocean, the more complex and less understood the environment becomes. Images from 200m might still retain some semblance of surface light, leading to better clarity compared to the more light-deprived 400m depth. This light variation can significantly affect the quality and characteristics of captured images.
- Technological Variations: ROVs and AUVs might be equipped with different imaging technologies. The variation in camera quality, resolution, and light sensitivity can result in images with distinct characteristics, affecting model performance (Fujimura et al., 2020).
- Image Noise and Artifacts: At greater depths, images might also be more susceptible to noise due to factors like water turbidity, camera motion, or interference. Such noise can further exacerbate the challenges of out-of-distribution data handling (Er et al., 2023).

5. CONCLUSION

In this study, the challenges of underwater image analysis, especially in the realm of deep-water image classification, were explored. The unique environmental intricacies of underwater settings, coupled with the technological variations in imaging methods, introduce significant domain shifts in data, which traditional machine learning models might struggle to handle effectively.

The primary focus was to assess the efficacy of domain adaptation techniques in managing out-of-distribution data for classifying the marine organism, *Eusergestes similis*. The comparison spanned several methods: Baseline (without domain adaptation), MCD, MDD, and DeepCORAL. Among these, the DeepCORAL method stood out, showcasing substantial enhancements in accuracy across both out-of-distribution datasets and forming discrete clusters in the encoded feature space.

The t-SNE visualizations of the encoded feature spaces further emphasized the importance of domain adaptation. While the baseline model struggled with a scattered representation, the DeepCORAL model adeptly formed organized clusters, indicating its robustness in recognizing and categorizing the target species.

In light of the results and discussions, it is evident that domain adaptation techniques, particularly DeepCORAL, hold significant promise in underwater image analysis. They not only bridge the gap between diverse data distributions but also ensure that machine learning models are both theoretically sound and practically applicable for real-world deep-water research and conservation efforts. As marine research continues to delve deeper into the oceans, harnessing the power of domain adaptation will be paramount in ensuring the accuracy and reliability of automated image analysis.

ACKNOWLEDGEMENTS

I extend my sincere gratitude towards all my mentors from the Bio Inspiration Lab for your guidance and motivation throughout the internship. Thank you to the Video Lab and Robison Lab for annotating and labelling ROV and AUV datasets needed to make this project a success. A special thanks to George Matsumoto for your compassion and support provided during the internship. I especially acknowledge NSF-IOS#2137977(to Kakani Katija), the David and Lucile Packard Foundation, and the Monterey Bay Aquarium Research Institute for providing the resources to support this project. A special thanks to Gayani Thilakarathne, my university mentor for recommending me for the internship program. Finally, a heartfelt gratitude towards my fellow interns for making the internship a memorable experience.

BIBLIOGRAPHY

Beijbom, O., Edmunds, P. J., Kline, D. I., Mitchell, B. G., & Kriegman, D. (2012, 16-21 June 2012). Automated annotation of coral reef survey images. 2012 IEEE Conference on Computer Vision and Pattern Recognition,

Bellingham, J., & Rajan, K. (2007). Robotics in Remote and Hostile Environments. *Science (New York, N.Y.)*, *318*, 1098-1102. https://doi.org/10.1126/science.1146230

Christ, R. D., & Wernli, R. L. (2014). Chapter 1 - The ROV Business. In R. D. Christ & R. L. Wernli (Eds.), *The ROV Manual (Second Edition)* (pp. 3-20). Butterworth-Heinemann. https://doi.org/https://doi.org/10.1016/B978-0-08-098288-5.00001-4

de Mathelin, A., Deheeger, F., Richard, G., Mougeot, M., & Vayatis, N. (2021). Adapt: Awesome domain adaptation python toolbox. *arXiv preprint arXiv:2107.03049*.

Er, M. J., Chen, J., Zhang, Y., & Gao, W. (2023). Research Challenges, Recent Advances, and Popular Datasets in Deep Learning-Based Underwater Marine Object Detection: A Review. *Sensors*, *23*(4), 1990. https://doi.org/10.3390/s23041990

Fujimura, Y., Iiyama, M., Hashimoto, A., & Minoh, M. (2020). Photometric Stereo in Participating Media Using an Analytical Solution for Shape-Dependent Forward Scatter. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, *42*(03), 708-719. https://doi.org/10.1109/tpami.2018.2889088

González-Rivero, M., Beijbom, O., Rodriguez-Ramirez, A., Bryant, D. E. P., Ganase, A., Gonzalez-Marrero, Y., Herrera-Reveles, A., Kennedy, E. V., Kim, C. J. S., Lopez-Marcano, S., Markey, K., Neal, B. P., Osborne, K., Reyes-Nivia, C., Sampayo, E. M., Stolberg, K., Taylor, A., Vercelloni, J., Wyatt, M., & Hoegh-Guldberg, O. (2020). Monitoring of Coral Reefs Using Artificial Intelligence: A Feasible and Cost-Effective Approach. *Remote Sensing*, *12*(3), 489. https://www.mdpi.com/2072-4292/12/3/489

Han, M., Lyu, Z., Qiu, T., & Xu, M. (2020). A Review on Intelligence Dehazing and Color Restoration for Underwater Images. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, *50*(5), 1820-1832. https://doi.org/10.1109/TSMC.2017.2788902

Jian, M., Liu, X., Luo, H., Lu, X., Yu, H., & Dong, J. (2021). Underwater image processing and analysis: A review. *Signal Processing: Image Communication*, *91*, 116088. https://doi.org/https://doi.org/10.1016/j.image.2020.116088

Quionero-Candela, J., Sugiyama, M., Schwaighofer, A., & Lawrence, N. D. (2009). Dataset Shift in Machine Learning.

Saito, K., Watanabe, K., Ushiku, Y., & Harada, T. (2018). Maximum classifier discrepancy for unsupervised domain adaptation. Proceedings of the IEEE conference on computer vision and pattern recognition,

Salas, E., & Argialas, D. (2022). *Automatic Identification of Marine Geomorphologic Features Using Convolutional Neural Networks in seafloor Digital Elevation Models: Segmentation of DEM for Marine Geomorphologic Feature Mapping with Deep Learning Algorithms* Proceedings of the 12th Hellenic Conference on Artificial Intelligence, Corfu, Greece. https://doi.org/10.1145/3549737.3549766

Sun, B., & Saenko, K. (2016). Deep coral: Correlation alignment for deep domain adaptation. Computer Vision–ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8-10 and 15-16, 2016, Proceedings, Part III 14, Zhang, Y., Liu, T., Long, M., & Jordan, M. (2019). Bridging theory and algorithm for domain adaptation. International conference on machine learning,