HIGH SCHOOL EDITION
Journal of Student Research
Volume 14 Issue 1 (2025)

Impact of Deep Learning Architectures in Coral Bleaching Detection

Mary Robyn Bernabe[1] and Divya Rajagiri[1]

[1]British School Manila, Philppines
[#]Advisor

ABSTRACT

Bleached coral reefs, a result of environmental stress, signal a concerning decline in marine ecosystem health. Coral bleaching is a deadly process which reduces coral populations, causing world-wide environmental issues such as the loss of habitats for wildlife. These detrimental after-effects can be alleviated if the health status of corals are detected early, and when bleaching initially begins. Existing coral bleaching detectors mostly rely on manual imaging and classifications, which are time-consuming and susceptible to human error. Therefore, deep learning techniques were employed to extract patterns and discern the health status of coral reefs from underwater images. We hypothesize that the accuracy and effectiveness of deep learning models in identifying coral bleaching events from underwater imagery are influenced by the underlying architectural design, with models leveraging deeper networks like VGG16 outperforming lighter models such as MobileNetV2 and ResNet50 in terms of recall and overall accuracy. A dataset with diverse underwater images of coral reefs was compiled. This consisted of 923 of total images, with the distribution as follows: 485 (53%) of images were bleached while 438 (47%) were healthy. We further evaluated the efficacy of different convolutional neural network models, including popular architectures like MobileNetV2, ResNet50, and VGG16. Through several experiments, VGG16 was found to be the most effective in accurately classifying coral health status, achieving the accuracy of 89.02, the highest among the tested models.

Introduction

Coral reefs are vital in tropical seas, forming the ecosystems which support significant portions of their oceanic population. It is estimated that coral reefs cover 0.1-0.5% of the oceanic floor (with supported findings of Spalding and Grenfell, 1997: 255,000 km2; Smith, 1978: 617,000 km2; Copper, 1994: 1,500,000 km2). Additionally, one-third of the world's fish species rely and inhabit these reefs, which contributes to 10% of global fish consumption by humans. The global reliance on coral reefs is significant, with over 100 countries having coastlines containing these ecosystems, supporting the livelihoods and protein intake of tens of millions of people (Salvat, 1992). Jennings and Polunin (1996) emphasize the potential of actively growing reef areas, suggesting that 1 km2 could sustain over 300 people if no alternative protein sources were available [1].

Abundant in species and population, corals serve as a primary source of protein for millions in tropical countries while also providing shoreline protection [2]. After coral bleaching, coral reefs face multiple challenges: the reproductive success of dominant reef builders is jeopardized, leading to potential extinction, while habitat degradation favors genetically isolated clusters of smaller, resilient corals. Increased sea temperatures from climate change trigger coral bleaching, and human activities increase competition from seaweeds and predator populations, as a result of the decreased supply of these populations [3].

In many tropical regions, sea temperatures have risen almost 1°C over the past 100 years, whereas it is predicted that it will rise ~1-2°C each century. In the last 20 years, it was found that these rising temperatures expose corals to heat, resulting in a mass loss of zooxanthellae, a type of dinoflagellate living in their tissue which provides their energy as well as their color vibrancy [4]. The escalating number of coral reef bleaching occurrences, illustrated with a surge in major events from 1979 to 1990 and widespread consequences including over 95% mortality in certain areas, present a major concern in marine ecology [5].

This extraction of this dinoflagellate not only lessens corals' visual appeal but also causes them to be vulnerable to disease. This process is primarily resulting from climate change, specifically global warming, highlighting the urgent need for comprehensive measures to mitigate temperature rise and preserve the health of these intricate underwater ecosystems [6].

Preventing coral bleaching is key in order to preserve marine biodiversity and sustain ecosystems supporting various underwater species. Coral reefs are important for the provision for fisheries, shoreline protection, and tourism. Identification involves monitoring for color changes, loss of vibrancy, and visible signs of stress. Early detection enables intervention to address and act upon these underlying conflicts with corals, and alleviate the detrimental effect on these underwater reefs [7].

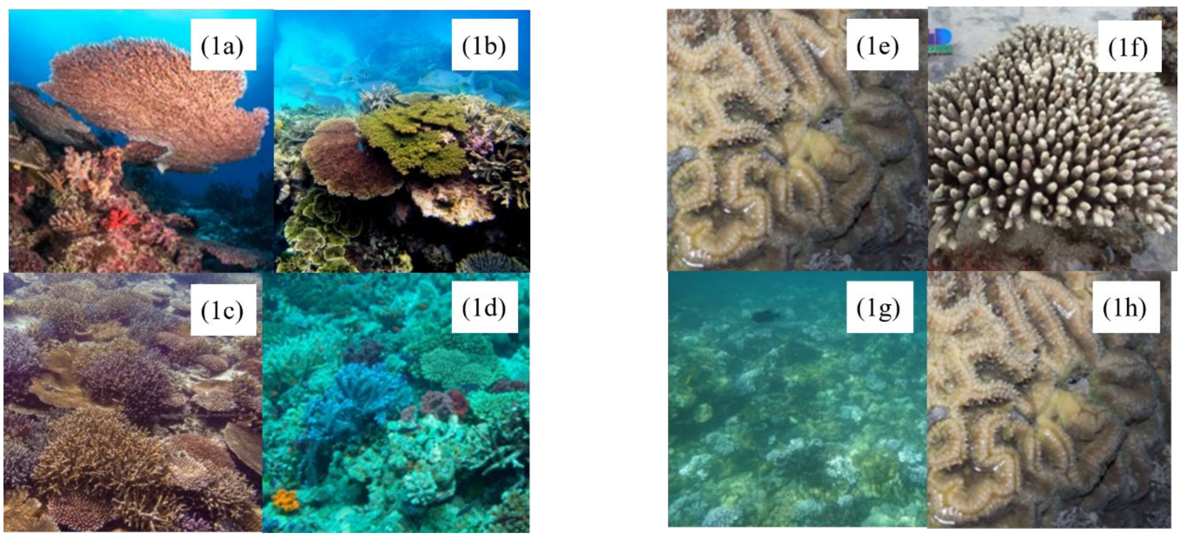

Figure 1. Images of Healthy (1a, 1b, 1c, 1d) and Bleached (1e, 1f, 1g, 1h) Corals. The following presents eight total randomly chosen images from the dataset from the validation classification of images, four images being bleached corals and four being healthy. In this dataset, there were a total of 923 images with the distribution of 485 (53%) bleached coral images and 438 (47%) healthy coral images.

Deep learning allows computational models with multiple processing layers to represent data to mirror how the human brain understands information. By capturing intricate structures within large-scale datasets, deep learning proves to be a versatile family of methods that includes neural networks, unsupervised and supervised feature learning algorithms. There has been a recent surge of interest in these techniques due to its accuracy and efficiency compared to previous techniques proved to perform less [8].

Convolutional neural networks (CNNs) are neural networks which incorporate a convolution operation as a fundamental layer, replacing the traditional fully connected layers. This has been proved to be successful, CNNs excel in scenarios where input data exhibits a grid like topology, such as time series (1-D grid) or images (2-D grid). Serving as a pivotal advancement in computer vision during the digital era, CNNs played a key role in propelling deep learning into contemporary applications [9].

In underwater image analysis, deep learning has emerged as a powerful tool, particularly in marine object detection and recognition [10]. Specifically, there has been similar works of image classification and deep learning used in underwater ecosystem situations, for the analysis of different fish species [11], as well as a similar project of the use of deep learning techniques and convolutional neural networks (CNNs) to understand the health status of coral reefs [12].

Deep learning techniques have demonstrated state-of-the-art results in various computer vision tasks, including image classification, object detection, and scene understanding. The intricacies of marine ecosystems pose unique challenges for computer vision, making automated technologies crucial for effective analysis [13].

In conclusion, this study aimed to investigate the application of deep neural network based models in the early detection of coral bleaching in underwater imagery. The primary hypothesis of our work was to understand if depth and complexity of a convolutional neural network (CNN) significantly impacts its ability to accurately detect coral bleaching in underwater imagery, specifically in terms of recall, precision, and overall accuracy. Our findings show insights into the potential of deep learning techniques in the field of marine ecology. To achieve these findings, we used three convolutional neural network architectures with varying complexity: VGG16, ResNet50, and MobileNetV2. We achieved an accuracy rate of 89.02 percent using the VGG16 architecture, validating the hypothesis that a more complex model architecture with a larger number of parameters results in better performance. VGG16, the most complex model, has 138 million parameters, followed by ResNet50 with 25.6 million parameters, and MobileNetV2, the lightest model, with only 3.4 million parameters. Our study highlights the importance of leveraging cutting-edge technology to address the pressing challenge of coral bleaching, while underscoring the potential of deep learning in safeguarding the future of coral reefs and marine ecosystems.

Results

The objective of the experiments was to assess the impact of model architecture complexity on the ability of deep learning models to detect coral bleaching events from underwater imagery. Specifically, we aimed to determine whether deeper networks like VGG16 outperform lighter models such as MobileNetV2 and ResNet50 in terms of recall and overall accuracy.

Three pre-trained models—MobileNetV2, ResNet50, and VGG16—were utilized. VGG16, the largest model with 138 million parameters, was hypothesized to outperform ResNet50 (25.6 million parameters) and MobileNetV2 (3.4 million parameters). A dataset of 923 images (53% bleached coral, 47% healthy coral) was used, and hyperparameters such as learning rates (ranging from 0.000001 to 0.05) and epochs (spanning from 10 to 50) were adjusted to optimize performance. Recall and overall accuracy were used as key metrics, with recall being particularly important in this context due to its relevance in identifying bleached coral (i.e., minimizing false negatives).

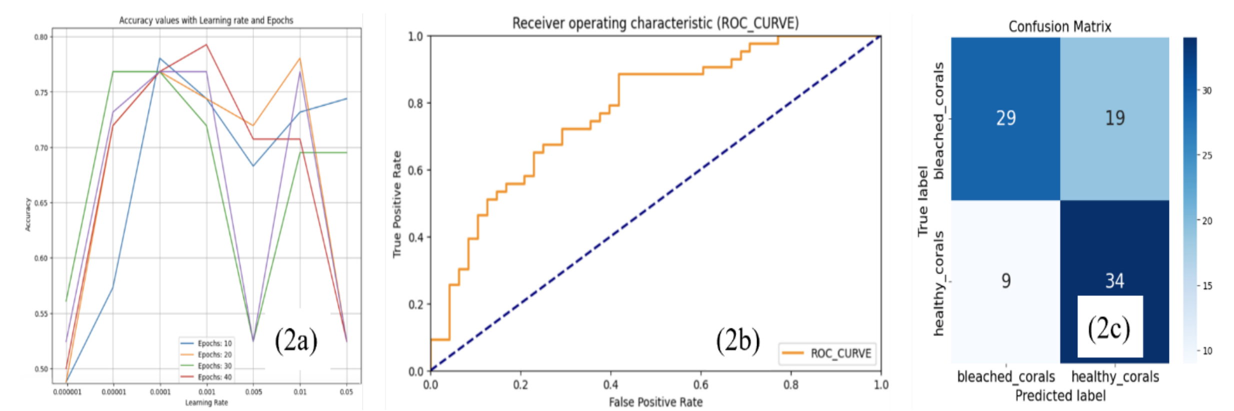

Figure 2. Results from experiments with the MobileNetV2 model are presented in this figure. The Multi-Line plot (2a) shows the results of hyper-parameter tuning on the validation dataset. X-axis represents the learning rate and y-axis represents the accuracy value. Results from using different epoch values are represented by different lines in the graph where the legend describes the value of the epoch used. ROC Curve (2b) showing the performance of the model on the test subset. Confusion Matrix (2c) for MobileNetV2 calculated using the test dataset.

MobileNetV2 achieved a maximum accuracy of 79.27%, with a recall of 72.35%. While lightweight and computationally efficient, MobileNetV2 demonstrated lower overall performance across all metrics. ResNet50 showed improved performance, achieving a maximum accuracy of 85.37% and a recall of 80.12%. This suggests that an increase in model depth and complexity led to enhanced detection capabilities. VGG16, as hypothesized, performed the best, achieving the highest accuracy of 89.02% and the highest recall of 84.65%.

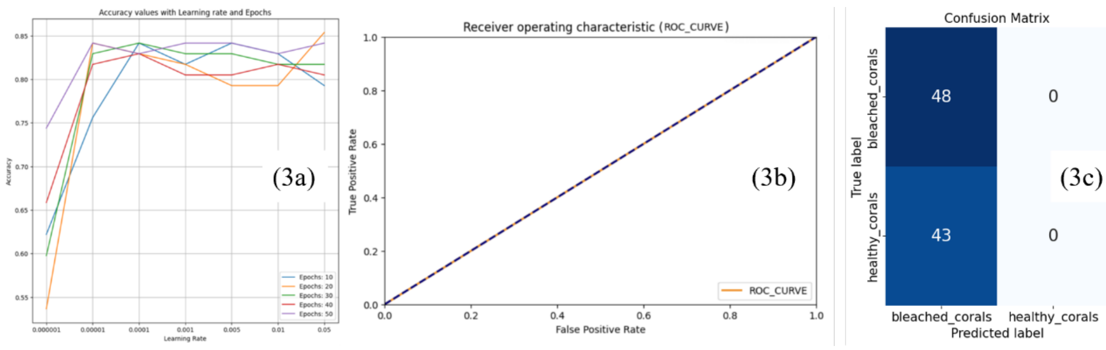

Figure 3. Results from experiments with the ResNet50 model are presented in this figure. The Multi-Line plot (3a) shows the results of hyper-parameter tuning on the validation dataset. X-axis represents the learning rate and y-axis represents the accuracy value. Results from using different epoch values are represented by different lines in the graph where the legend describes the value of the epoch used. ROC Curve (3b) showing the performance of the model on the test subset. Confusion Matrix (3c) for ResNet50 calculated using the test dataset.

The best model from each pre-trained architecture is used to evaluate its performance on the test data (Table). The confusion matrix along with the ROC curves are reported for each model on the test data. Based on all the best data collected from each pre-trained model, it was found that the VGG16 model had the highest overall accuracy of 89.02 supporting our hypothesis of a model with a larger number of parameters performing the best.

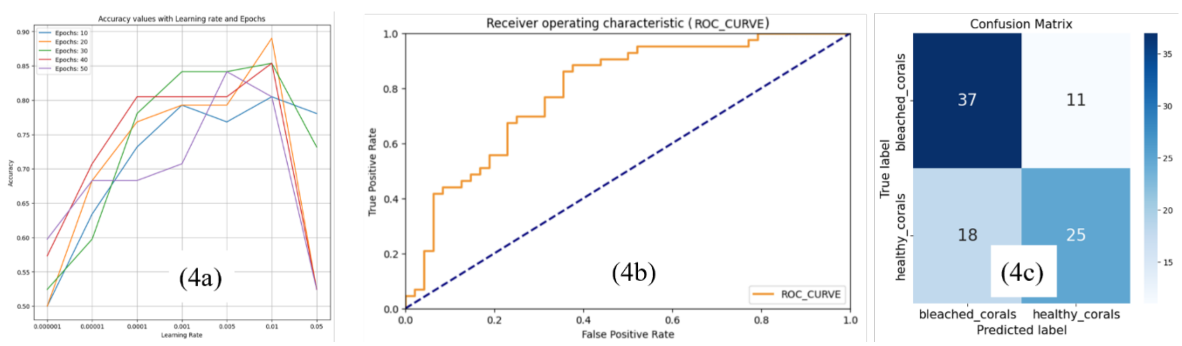

Figure 4. Results from experiments with the VGG16 model are presented in this figure. The Multi-Line plot (4a) shows the results of hyper-parameter tuning on the validation dataset. X-axis represents the learning rate and y-axis represents the accuracy value. Results from using different epoch values are represented by different lines in the graph where the legend describes the value of the epoch used. ROC Curve (4b) showing the performance of the model on the test subset. Confusion Matrix (4c) for VGG16 calculated using the test dataset.

Discussion

Experiments for each pre-trained deep learning model (MobileNetV2, ResNet50, VGG16) were conducted which obtained accuracies ranging from 48.78 to 89.02 percent in detecting the health status of coral reefs, demonstrating varying levels of effectiveness in classifying bleached and healthy corals. These experiments were held through different learning rates and epochs, which found that the VGG16 model outperformed both MobileNetV2 and ResNet50 achieving the optimal accuracy of 89.02 percent. Generally, all models found low accuracies with instability at lower learning rates, whereas higher and stable accuracies were observed at varied, higher learning rates.

In the case of MobileNetV2, the optimal learning rate for this task appears to be around 0.001, where the model achieves consistent and high performance across different epochs. Moderate learning rates allow for steady and stable learning without the risk of overfitting and underfitting too quickly. The performance at very low and very high learning rates illustrates the importance of choosing appropriate learning rate for the task and dataset at hand. It's also clear that monitoring model performance across epochs is crucial for identifying the best point to stop training before overfitting occurs.

The optimal learning rates for training the ResNet50 model on binary classification appear to be in the range of 0.0001 to 0.005, with 0.001 showing particularly stable and high performance across all epochs. These rates balance the speed of convergence with the risk of instability or overfitting. Very high learning rates (0.05) demonstrate the potential for rapid learning but at the risk of initial volatility. Conversely, very low rates (0.000001) are too conservative, requiring many epochs to reach comparable accuracy levels. The data underscores the importance of selecting a learning rate that aligns with the model's complexity and the dataset's characteristics to ensure efficient and effective learning.

Table 1. Test results from the models: MobileNetV2, ResNet50, VGG16.

Results	MobileNetV2	ResNet50	VGG16
Accuracy	79.27	85.37	89.02
Precision	76.32	52.75	69.44
Recall	79.07	100.00	77.08

The performance of the VGG16 model for binary classification across various learning rates and epochs reveals critical insights into its learning dynamics. Initially, at very low learning rates (0.000001, 0.00001), the model shows minimal improvement, struggling to significantly learn from the data, which is evident from the stagnant performance at the lowest rate and modest gains at the slightly higher rate. However, as the learning rate increases to moderate levels (0.0001, 0.001), there's a notable improvement in accuracy, demonstrating the model's capacity to effectively learn and adapt to the dataset with an optimal learning rate around 0.001 and 0.005, where it achieves its peak performance (89.02 at 20 epochs for 0.005). Surprisingly, at high learning rates (0.01, 0.05), the model's performance becomes inconsistent, peaking at mid-epochs before deteriorating, which could indicate issues related to overfitting or the inability to stabilize due to the aggressive learning rates.

Analyzing the validation results across MobileNetV2, ResNet50, and VGG16 for binary classification reveals how each architecture responds to different learning rates. Common across all models is the identification of an optimal learning rate range that maximizes performance while minimizing risks such as overfitting and underfitting. For MobileNetV2, a learning rate around 0.001 is identified as optimal, emphasizing the model to achieve high and stable performance without rapid overfitting. ResNet50 displaying a slightly broader optimal range from 0.0001 to 0.005, suggests flexibility in adapting to various rates efficiently. The optimal learning rate for VGG16 is typically identified as around 0.01. This value was chosen to promote the model's ability to achieve consistent and high performance while mitigating the risk of overfitting. All models demonstrate diminished performance at very low learning rates, highlighting the slow convergence and inefficiency of overly conservative learning strategies. At high learning rates, the models exhibit either volatility or a sharp decline in performance beyond mid-epochs, underscoring the delicate balance required to avoid overfitting while still encouraging robust learning. These findings show the critical role of learning rate optimization in model training, underscoring a shared principal across different architectures. Moderate learning rates tend to offer the best balance between convergence, speed and accuracy, although the specific optimal point varies slightly with each model's unique characteristics.

It is important to consider several factors and limitations that could influence outcomes. Firstly, the diversity and the size of the dataset play crucial roles in the training process. A limited dataset may not capture the full variability of the classes involved, potentially skewing model performance and generalization capability. Future experiments could benefit from larger and more diverse datasets to ensure models are well-generalized and robust against overfitting. Additionally, investigating how these models perform across a wider range of tasks and in more complex, real-world scenarios could further validate the generalizability of these findings.non proident, sunt in culpa qui.

Materials and Methods

The dataset used for the experiments was found at a public domain source providing images of healthy and bleached corals collected from underwater sites. This consisted of 923 of total images, in which the images were distributed as follows: 485 (53%) of images were bleached while 438 (47%) were healthy. After the images were correctly classified into their respective categories, the data was split into three groups when uploaded into Google Drive: 10% of images for the validation dataset, 10% of images sorted for the testing dataset, and 80% of the images were used for the training dataset.

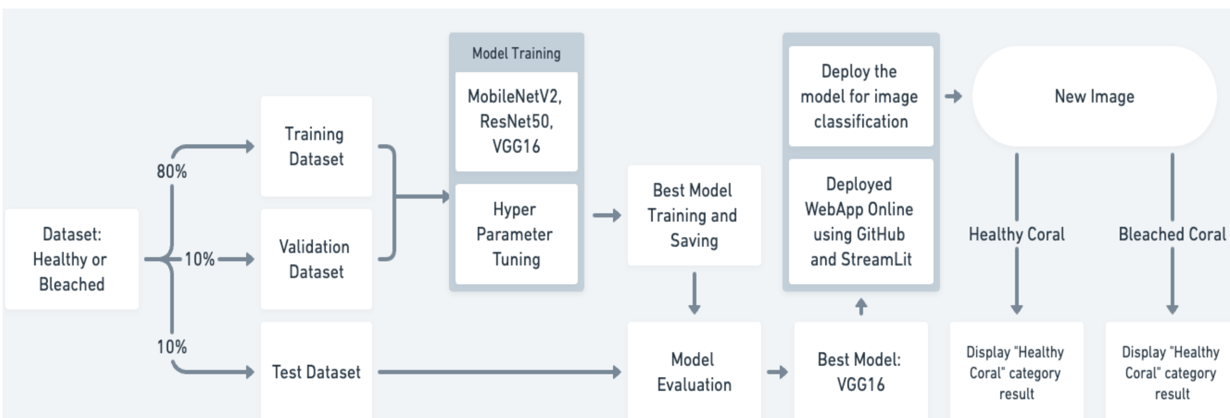

Figure 5. Flowchart Process. The following diagram exhibits the creation process of the final product of this image classification.

Firstly, three pre-trained deep learning models were chosen (MobileNetV2, ResNet50, VGG16) for model evaluation. These Keras convolutional neural networks (CNNs) were selected for evaluation due to their popular use for image classification.

MobileNetV2 is a neural network designed to be lightweight and efficient, making it ideal for applications on mobile devices with limited resources. This efficiency is achieved through unique convolutional operations, reducing computational demands while maintaining good performance.

ResNet50 is a deep neural network used for its depth and the introduction of residual connections. These connections address challenges associated with training very deep networks, enabling the model to effectively capture complex features. Particularly, this is powerful for image classification tasks on large and diverse datasets.

VGG16, a straightforward and effective neural network architecture, has a strength which lies in its uniform structure with stacked convolutional layers and max-pooling layers, allowing it to capture intricate patterns in images. VGG16 is often used for transfer learning, leveraging pre-trained weights on large datasets for tasks with smaller datasets.

These models were tuned by varying the hyper-parameters learning rate and epochs, where the following combinations were used: learning rates of 0.000001, 0.00001, 0.0001, 0.001, 0.005, 0.01, 0.05 and epochs of 10, 20, 30, 40, 50, with a batch size of 32. The model with the highest validation accuracy was saved and used for evaluating the test subset of the dataset. Through the comparisons of these accuracies (Table 1) from different deep learning models, it was found that VGG16 had the highest accuracy of 89.02 with the combination of an epoch of 20 and a learning rate of 0.005.

In the process of the creation of this app, Google Colaboratory was used for all experiments. TensorFlow and Keras were employed as the primary deep learning frameworks [14], with the latter benefiting from pre-trained models available for various applications. Python served as the programming language, supported by essential libraries such as NumPy [15] for numerical operations. Additionally, Matplotlib [16] played a crucial role in visualizing and interpreting the results, enhancing the overall analytical capabilities of the development process.

Acknowledgments

I would like to thank my advisor for the valuable insight provided to me on this topic.

References

Moberg, F., & Folke, C. (1999). Ecological goods and services of coral reef ecosystems. *Ecological economics*, *29*(2), 215-233. https://doi.org/10.1016/S0921-8009(99)00009-9

Sheppard, C., Davy, S., Pilling, G., & Graham, N. (2017). *The biology of coral reefs*. Oxford University Press.

Knowlton, N. (2001). The future of coral reefs. *Proceedings of the National Academy of Sciences*, *98*(10), 5419-5425. https://doi.org/10.1073/pnas.091092998

Hoegh-Guldberg, O. (1999). Climate change, coral bleaching and the future of the world's coral reefs. *Marine and freshwater research*, *50*(8), 839-866. https://doi.org/10.1071/MF99078

Glynn, P. W. (1993). Coral reef bleaching: ecological perspectives. *Coral reefs*, *12*, 1-17. https://doi.org/10.1007/BF00303779

Douglas, A. E. (2003). Coral bleaching—how and why?. *Marine pollution bulletin*, *46*(4), 385-392. https://doi.org/10.1016/S0025-326X(03)00037-7

Solayan, A. (2016). *Biomonitoring of coral bleaching—A glimpse on biomarkers for the early detection of oxidative damages in corals* (pp. 101-117). InTech. http://dx.doi.org/10.5772/61831

Chai, J., Zeng, H., Li, A., & Ngai, E. W. (2021). Deep learning in computer vision: A critical review of emerging techniques and application scenarios. *Machine Learning with Applications*, *6*, 100134. https://doi.org/10.1016/j.mlwa.2021.100134

Ketkar, N., Moolayil, J., Ketkar, N., & Moolayil, J. (2020). Deep Learning with Python: Learn Best Practices of Deep Learning Models with PyTorch. *Apress LP*. https://doi.org/10.1007/978-1-4842-5364-9

Lumini, A., Nanni, L., & Maguolo, G. (2020). Deep learning for plankton and coral classification. *Applied Computing and Informatics*, *19*(3/4), 265-283. https://doi.org/10.1016/j.aci.2019.11.004

Rathi, D., Jain, S., & Indu, S. (2017, December). Underwater fish species classification using convolutional neural network and deep learning. In *2017 Ninth international conference on advances in pattern recognition (ICAPR)* (pp. 1-6). Ieee. DOI: 10.1109/ICAPR.2017.8593044

Narayan, R. Using Deep Learning to Monitor Coral Reef Health.

Mahmood, A., Bennamoun, M., An, S., Sohel, F., Boussaid, F., Hovey, R., ... & Fisher, R. B. (2017). Deep learning for coral classification. In *Handbook of neural computation* (pp. 383-401). Academic Press. https://doi.org/10.1016/B978-0-12-811318-9.00021-1

Manaswi, N. K., & Manaswi, N. K. (2018). Understanding and working with Keras. *Deep learning with applications using Python: Chatbots and face, object, and speech recognition with TensorFlow and Keras*, 31-43. https://doi.org/10.1007/978-1-4842-3516-4_2

Bressert, E. (2012). SciPy and NumPy: an overview for developers.

Tosi, S. (2009). *Matplotlib for Python developers*. Packt Publishing Ltd.