



Article Virtual Underwater Datasets for Autonomous Inspections

Ioannis Polymenis 🔍, Maryam Haroutunian, Rose Norman 🐿 and David Trodden

School of Engineering, Newcastle University, Newcastle upon Tyne NE1 7RU, UK * Correspondence: rose.norman@ncl.ac.uk

Abstract: Underwater Vehicles have become more sophisticated, driven by the off-shore sector and the scientific community's rapid advancements in underwater operations. Notably, many underwater tasks, including the assessment of subsea infrastructure, are performed with the assistance of Autonomous Underwater Vehicles (AUVs). There have been recent breakthroughs in Artificial Intelligence (AI) and, notably, Deep Learning (DL) models and applications, which have widespread usage in a variety of fields, including aerial unmanned vehicles, autonomous car navigation, and other applications. However, they are not as prevalent in underwater applications due to the difficulty of obtaining underwater datasets for a specific application. In this sense, the current study utilises recent advancements in the area of DL to construct a bespoke dataset generated from photographs of items captured in a laboratory environment. Generative Adversarial Networks (GANs) were utilised to translate the laboratory object dataset into the underwater domain by combining the collected images with photographs containing the underwater environment. The findings demonstrated the feasibility of creating such a dataset, since the resulting images closely resembled the real underwater environment when compared with real-world underwater ship hull images. Therefore, the artificial datasets of the underwater environment can overcome the difficulties arising from the limited access to real-world underwater images and are used to enhance underwater operations through underwater object image classification and detection.

Keywords: machine learning; deep learning; computer vision; AUVs; ROVs; autonomous inspection

1. Introduction

Remotely Operated Vehicles (ROVs) and Autonomous Underwater Vehicles (AUVs) are used extensively by the offshore oil and gas industry and the offshore renewables sector, as well as by the marine scientific community, to carry out tasks such as the inspection and maintenance of underwater structures and the survey of the oceans in a manner that is both safe and accurate. To a significant extent, missions carried out by ROVs and AUVs depend on visual inputs to accurately portray and comprehend the surrounding subsea world. As a consequence of this, current underwater intervention requires the use of image classification and object detection techniques. ROVs have been developed and optimised in terms of underwater operations and object manipulations to develop robust vehicles that make physical intervention possible. This was a result of research needs as well as industry requirements for more complex and demanding operations. The same industrial and scientific needs have resulted in the optimisation of AUVs, specifically in terms of power requirements, manoeuvrability, navigation, communication, and autonomy. This establishes AUVs as the predominant means for underwater operations in areas that are hazardous for humans to be in [1].

As a consequence of this, autonomous underwater intervention has to make use of the most promising available technologies of Artificial Intelligence (AI), namely those of Deep Learning (DL) and Computer Vision (CV).

The existing research on autonomous underwater operations takes advantage of advancements in DL and CV and offers practical solutions to the issue of underwater image



Citation: Polymenis, I.; Haroutunian, M.; Norman, R.; Trodden, D. Virtual Underwater Datasets for Autonomous Inspections. *J. Mar. Sci. Eng.* **2022**, *10*, 1289. https://doi.org/10.3390/jmse10091289

Academic Editor: Antoni Burguera

Received: 19 August 2022 Accepted: 7 September 2022 Published: 13 September 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/). enhancement and restoration by increasing the resolution and contrast of underwater images [2]. In addition, the work of [3] improves the image quality by resolving the issue of poor lighting that occurs in an underwater setting. Although these methods provide outcomes that are considered to be state of the art, none of the existing methods deals with the challenge of having access to readily available underwater image datasets.

Currently, the problem of readily available datasets is well known to research areas from Unmanned Aerial Vehicles (UAVs) applications [4] to unmanned underwater operations [5], and it is the fundamental issue that needs to be addressed to improve the performance of the different learning and detection models used by those fields. Particularly, for underwater operations, the issue arises primarily because of the high cost and the difficulties that need to be overcome to collect such a dataset from real-world underwater missions [5]. Because of this, the current work addresses the issues mentioned above and creates a model that enables the generation of custom underwater images by making use of common objects that can be found in underwater structures such as gas pipelines, underwater cables, oil/gas wells, wind turbine piles, etc. In addition, images with characteristics similar to those seen in the natural subsea environment were produced using recent developments in DL models.

Related Work

Unmanned robots need to understand their complex environment to achieve complete autonomous operational capabilities, with object detection being the fundamental low-level task [4]. The demand for underwater vehicles to achieve autonomous capabilities is even more demanding due to the challenging environmental conditions. Underwater Vehicles are equipped with various sensors and instrumentation such as GPS, cameras, LiDAR cameras, and sonars [6,7]. Cameras are essential because they allow for visual interaction between the user/operator and the vehicle [8].

During the past few years, object detection models have become more sophisticated and accurate than ever before, and they are able to take advantage of the modern embedded systems [4]. Some landmark architectures that revolutionise modern CV applications include the R-CNN family of models [9–11], the YOLO architecture and its different versions [12,13], and Feature Pyramid Networks (FPN) [4,14]. Ross Girshick et al. [9] introduced the algorithm designed to overcome the problem of selecting large regions during an object detection task. The model performs a selective search on the image by looking for potential objects (region proposals). Therefore, instead of detecting and classifying a larger region, the model divides the image into smaller regions, which ultimately increases the total training time. The Fast R-CNN model [4,10] was introduced to solve some of the drawbacks of the original model and create a faster model. The two models were approached from the same point of view, with the difference that the input image was fed to the CNN architecture to generate feature maps instead of region proposals, which increased the model's speed and accuracy. The Mask R-CNN model [4,11] was introduced to solve the problem of object detection in images with complex backgrounds, and the model was able to detect objects with high accuracy and segmentation.

All prior object detection methods focus the item inside the picture using areas. The network does not analyse the picture in its entirety. Instead, the probability-rich regions of the picture that contain the item are the focus. You Only Look Once, or YOLO, is an object detection algorithm that differs significantly from the region-based techniques described previously. In YOLO, the bounding boxes and class probabilities for these boxes are predicted by a single neural network. YOLO performs object detection quicker than a conventional object detection algorithms, with speeds ranging from 45 to 145 frames per second. The drawback of the YOLO algorithm is that it struggles to detect small objects in an image, but this changes with more recent versions.

Finally, FPN [14] is not a standalone object detector, and can be classified as a feature extractor that operates in conjunction with object detectors such as R-CNN and Fast R-CNN. The feature extractor accepts a single-scale picture of any arbitrary size as input, and generates correspondingly scaled feature maps at several layers using a fully CNN algorithm. This process is independent of the convolutional architectures' core components. It is a general approach for constructing feature pyramids inside deep convolutional networks for applications such as object identification.

When applying deep learning approaches to problems involving image classification or object detection, one of the most frequent obstacles that arise is a lack of data. Applying data augmentation techniques to a dataset in order to expand its size and variety is a trial-anderror approach to the challenge of dealing with a lack of data [15]. The traditional method of data augmentation involves the use of various libraries, such as those described in [16,17], which provide flexibility and easy-to-use implementation for a variety of augmentations to increase the size and the diversity of the dataset. This approach is known as the "classic" method of data augmentation. These libraries include a variety of enhancement methods, including cropping, blurring, colour saturation, contrast, and greyscale scaling, as well as rotation, changing colour channels, and shifting colour channels.

When it comes to more project specific tasks, the standard data augmentation method cannot generate images that are close to the preferred real-world data, and it requires a significant amount of time and trial and error to produce the desired results. Therefore, DL models such as Generative Adversarial Networks (GANs), CycleGAN, and U-Nets are the current state-of-the-art methods used to augment datasets and increase their size [18–20]. GAN are mainly used to produce synthetic images that follow the same probability distribution as the real images. CycleGAN is a well-known GAN architecture that is typically used to learn image transformations across various patterns, whereas U-Net models focus more on semantic and structural differences between actual and artificial content. These methods is the most advanced currently available.

In addition, the difficulties presented by the underwater environment make the collection of data a laborious job that requires the use of specialised persons and specific equipment. As a consequence of this, it is difficult to build projects that need large underwater datasets. The underwater habitat, the light conditions that are present throughout the picture capturing process, and the task that the image was shot for all play a role in determining the unique problems that come with taking underwater photographs [21]. When it comes to obtaining data for deep learning models, many researchers have focused primarily on underwater image enhancement and restoration to improve the quality of underwater images [2,22,23]. This is to improve the quality of the images obtained from underwater environments.

The technique of enhancing underwater images has the improvement of the image's visual quality as its goal, and it does not often take into account the physical qualities of light in the water, such as the attenuation coefficient or the light scattering [22]. It is generally agreed that picture enhancement may be implemented far more quickly and is simpler to understand than image restoration. Image restoration is a more sophisticated process that has to take into consideration the physical behaviour of light in water. This is because water reflects light differently than air does. Image restoration requires information on the kind of water present, whether it is coastal or ocean water, as well as the quality of the light propagation in the water [2,23]. These methods only produce satisfactory outcomes in a controllable underwater environment, and it is difficult to put them into practice in the real world due to the complexity of their implementation and the large number of parameters that need to be taken into consideration [2].

By including an attenuation coefficient for both the blue–red and blue–green spectrum channels, the technique that was presented by Berman and colleagues [23] was able to consider the various light profiles produced by the various underwater settings. The technique they developed is based on the intensity of the image's colour channel at the pixel level; more specifically, the attenuation coefficient incorporates the two spectrum

characteristics. In addition to this, the topography of the location, the time of year, and the climate were all taken into consideration. Arnold-Bos and colleagues [24] discuss the challenges that the vision of underwater vehicles encounter while operating in underwater conditions and suggest using deconvolution and augmentation approaches. The technique was developed to eliminate light backscattering, the primary feature of noise, and the attenuation inequalities that arise with contrast equalisation. A wavelet filtering approach was then used on the residual picture noise, which may correlate with sensor noise or floating particles. This algorithm helps enhance the edge recognition of underwater images.

To increase the amount and quality of the datasets, more novel techniques as mentioned above, such as data augmentation using GANs, are being employed in various sectors. Some examples of the use of GAN models can be found in the field of neuroscience where, for instance, there is a need to perform segmentation tasks from CT scan images [25]. Additionally, the use of Deep Neural Networks (DNN) and U-Nets to perform segmentation in brain cell representation from Electron Microscopy (EM) images [20,26] is an example of the application of CycleGAN models for the purpose of data augmentation.

Because of the progress that has been made in DL and CV over the past few years in areas such as image classification, image segmentation, and object detection [18,27,28], there is now an opportunity to develop models that are capable of performing image restoration and image enhancement in a manner that is more accurate and precise. These models have the potential to outperform any of the manual approaches that were used in the past [19]. Because of the use of Convolution Neural Networks (CNN) and GANs, it is possible, in certain instances, to identify and detect objects with a higher level of accuracy than is possible for humans to attain [29].

Zhu et al. [28] proposed a CycleGAN model for image-to-image translation in order to learn the mapping functions between two domain images *X* and *Y*, translate the domain of the first image based on the second, $G : X \longrightarrow Y$, and vice versa, to translate the second domain based on the first image, $F : Y \longrightarrow X$. This allowed the model to translate the domain of interest. In addition, the authors incorporated two adversarial discriminators, one for the first domain image and one for the second domain image. The purpose of these discriminators was to assess whether or not the output image had been successfully translated to the target domain. In most instances, the results were adequate, and the translation of one image domain to another image domain delivered acceptable output. Nevertheless, the model might become confused between the domains when there is insufficient feature dispersion in the training set.

Currently, DL models are the standard in underwater applications, and the primary emphasis is on picture restoration, image enhancement, and improvement of underwater settings. Anwar et al. suggested using a CNN model that might improve the quality of photographs taken underwater [22]. The network design is composed of convolutional blocks that are all linked to a dense layer at the end. This provides the whole model with modularity. The model's output is an improved picture of the subsea scene, devoid of the cyan and emerald tones present in the original image.

A similar technique for restoring the colours in underwater photographs was used by Chen et al. [30], where the authors attempted to reduce the effects of the underwater environment, increase the picture details, and fix the colours in the image. The image incorporates several diverse components, each of which is represented by one of the model's three primary elements. The first part is used to estimate the ambient light of the image; the second part is responsible for the direct transmission estimation, which is a function of both the ambient light and the input image, and the third part is responsible for the reconstruction of the enhanced image. Li et al. [31] presented an underwater enhancement method based on GAN models, where they tried to solve underwater degradation effects such as low contrast, colour casts, and haze-like effects using a fusion GAN model on the U45 dataset. The model is utilised by combining the benefits of the inception model architecture [27] with the deep residual learning framework [32]. Another research study by Li et al. [33] approached underwater image enhancement from a different perspective by constructing a large-scale real-world underwater dataset containing 950 images under various light conditions, from natural light to artificial light. The collected data were then tested on the custom Water-Net model to perform image enhancement. Furthermore, Panetta et al. [34] went further in underwater object tracking and image enhancement and introduced a benchmark underwater dataset, UOT100. The dataset comprises 104 underwater videos, from which they generated a complete set of 74 K annotated image frames. Additionally, they introduced the CRN-UIE GAN model to perform image enhancement. The model tries to improve underwater object detection performance by correcting the underwater environment's effects.

Underwater image restoration using real-world images from coral reefs (HICRD) was proposed by Han et al. [35]. They created the custom HICRD dataset to overcome the limitation of previous datasets to capture a more diverse underwater environment. The dataset consists of 9676 images and is used on the Constructive UndeWater Restoration (CWR) model for image restoration. The CWR model at its core utilises GAN models and Representation Learning [36], essentially an unsupervised method used to perform image restoration. The CWR model performed satisfactorily, and the end result was close to the reference images without content or structural losses on the generated images.

In addition, during the last few years, there has been a shift from the conventional techniques toward the substantial use of CNN and GAN models for underwater image repair and enhancement [19,30,37,38]. Because of such networks' characteristics, DL models represented a significant advancement in the analysis of underwater photographs. The task of processing and interacting with the underwater world poses a number of difficulties for any autonomous vehicle. DL makes it possible to create more accurate data-driven models of the environment, improving one's capacity to analyse and comprehend that environment. The most notable benefit of DL models is that they can be put into action without the necessity of explicitly describing every facet of the environment and manually coding everything that is required for the operation line by line. This is the distinct advantage that sets them apart from other types of models. DL models can be trained and can learn the most valuable features on their own, provided the necessary data are fed through the network during the training process. As a result, the model would be able to learn the necessary characteristics and parameters for every given job, notwithstanding the complexity of the underwater scenes.

Consequently, as a result of the development of advanced deep learning algorithms, it is now more conceivable than it has ever been to generate underwater photographs that may be as similar as possible to the real world, despite the complexity of such an environment. Because of this, the data collection for marine imagery, which is an essential component of any project relating to the underwater environment, can be made more accessible and will not require direct underwater data, at least in the initial stages of the model development. This will result in savings of time, resources, and funding.

This paper is organised as follows. Section 2 describes the methodology used to generate the underwater photographs. Section 3 describes the results of the model development. Section 4 presents the conclusions of the paper.

2. Methods for Data Capture and Processing

In the technique that is proposed in this paper, a dataset was compiled based on items that have the potential to be discovered in sub-sea structures, and this was carried out in a laboratory environment. The next step is to compile a dataset that is as accurate a representation as possible of the real-life underwater environment. This is accomplished by utilising common photographs captured under typical atmospheric conditions and image-to-image translation powered by CycleGAN models. Therefore, the production of artificial or "fake" underwater datasets will make it possible to circumvent the challenges associated with the acquisition of real-world underwater photographs.

A primary dataset was compiled by taking pictures of components that are typical of underwater structures. These components include bolts, hex nuts, flanges, anodes, and pipelines. The next step was constructing an expanded data set using traditional Data Augmentation methods such as rotation, cropping, blurring, and changing the colour channel. The CycleGANs model was then trained on the images that were obtained together with additional photographs of the underwater environment using publicly available datasets [39,40]. Both datasets, one of which was produced by using data augmentation and the other of which was produced by using CycleGANs (a learning-based method), were compared using the Frechet Inception Distance (FID) metric [41] in order to determine which method is more appropriate.

2.1. Formulation of the Proposed Method

When they were initially presented in the area of DL by Goodfelow et al. in 2014 [42], the Generative Adversarial models and the adversarial loss made a significant contribution in the area of DL. The network consists of two parts: the Generator, which produces pictures that include characteristics of a particular domain, and the Discriminator, which attempts to accurately categorise actual images based on the created ones. Zhu et al. [28] proposed an enhanced adversarial loss in which the loss function employs the least-squared loss of the original loss. This is due to the fact that it demonstrates more stable behaviour during the training process.

Therefore, the adversarial loss can be expressed using Equations (1a) and (1b) as follows:

$$\mathcal{L}_{GAN_{uw \to lab}} = \mathbb{E}_{X_{uw}^{real}} [(D_{lab}(G_{lab}(X_{uw}^{real})) - 1)^2] \\ + \mathbb{E}_{X_{lab}^{real}} [(D_{lab}(X_{lab}^{real}))^2]$$
(1a)

$$\mathcal{L}_{GAN_{lab\to uw}} = \mathbb{E}_{X_{lab}^{real}}[(D_{uw}(G_{uw}(X_{lab}^{real})) - 1)^2]$$
(1b)

$$+\mathbb{E}_{X_{uw}^{real}}[(D_{uw}(X_{uw}^{real}))^2]$$

$$\mathcal{L}_{GAN} total = \mathcal{L}_{GAN_{lab \to uw}} + \mathcal{L}_{GAN_{uw \to lab}}$$
(1c)

where \mathbb{E} is the expected value for the underwater domain X_{uw} and the lab domain X_{lab} . The discriminator D_{uw} is responsible for the mapping of the lab images with underwater features and can be expressed as $G_{uw} : X_{uw} \to X_{lab}$, and discriminator D_{lab} is responsible for the inverse operation $G_{lab} : X_{lab} \to X_{uw}$. The total adversarial loss \mathcal{L}_{GAN} total is the summation of the lab loss $\mathcal{L}_{GAN_{lab} \to uw}$ and the underwater loss $\mathcal{L}_{GAN_{uw} \to lab}$.

2.2. Cycle Consistency Loss

The adversarial model has been trained to learn the properties of both the X_{uw} and X_{lab} domains in the cycleGAN models that were presented by Zhu et al. [28]. This indicates that a lab picture may be converted to another image that has characteristics from the area of interest, such as an image of an underwater environment. Because of this, the model has to be able to meet the cycle consistency between the two domains, which can be seen in Equations (2a) and (2b), respectively.

$$X_{uw}^{real} \to G_{lab}(X_{uw}^{real}) \to G_{uw}(G_{lab}(X_{uw}^{real})) = X_{uw}^{rebuild}$$
(2a)

$$X_{lab}^{real} \to G_{uw}(X_{lab}^{real}) \to G_{lab}(G_{uw}(X_{lab}^{real})) = X_{lab}^{rebuild}$$
(2b)

Then, the total cycle consistency loss $\mathcal{L}_{cycle}^{total}$ is $\mathcal{L}_{cycle}uw + \mathcal{L}_{cycle}lab$ where:

$$\mathcal{L}_{cycle}uw = \mathbb{E}_{X_{uvn}^{real}}[\|G_{uw}(G_{lab}(X_{uw}^{real})) - X_{uw}^{real}\|_{1}]$$
(3a)

$$\mathcal{L}_{cycle}lab = \mathbb{E}_{X_{lab}^{real}}[\|G_{lab}(G_{uw}(X_{lab}^{real})) - X_{lab}^{real}\|_{1}]$$
(3b)

Combining the adversarial and the cycle consistency loss, the total loss of the model will be

$$\mathcal{L}_{total} = \mathcal{L}_{GAN} + \lambda \mathcal{L}_{cycle} \tag{4}$$

where λ is the regularisation hyperparameter factor. The value of λ is utilised to regulate and optimise the performance of the network loss [42].

The architecture of the UnderWaterCycleGAN (UWCycleGAN) model is shown in Figure 1. In this model, the image characteristics of the underwater domain X_{uw} are transferred to the required lab image domain X_{lab} using the generator G_{uw} . Then, the generator G_{lab} applies the newly acquired features to the original lab pictures X_{lab}^{real} , which ultimately leads to the production of the artificial images X_{lab}^{fake} . The discriminator D_{lab} is responsible for monitoring and comparing the false pictures with the genuine ones at the very last phase of the network to guarantee that the results are satisfactory. In addition, throughout each step of the process, the model will reconstruct the pictures by making use of the cycle consistency losses. In particular, during the initial step of the process, the model recreates the photographs the underwater environment. $X_{uw}^{rebuild}$ is calculated using the cycle loss $\mathcal{L}_{cycle}uw$. The adversarial loss \mathcal{L}_{GAN} is used to optimise the final image output of the discriminator D_{lab} . This is carried out for the fake image X_{lab}^{fake} , first by using the \mathcal{L}_{GAN} loss to optimise the output of G_{lab} , and then using it to optimise the output of D_{lab} .



Figure 1. Cyclegan Model for Underwater Data Augmentation.

2.3. Dataset

Newcastle University's towing tank was used for the data collection needed for the image classification and object identification tasks. In particular, a dataset consisting of five different classes of objects was compiled, including bolts, hex nuts, flanges, pipes, and lead blocks (representing anodes). These objects were chosen because they are common in underwater structures and were readily available during the image data acquisition step. The above objects were initially placed inside the towing tank. Then, videos were captured using an underwater camera. During the acquisition process, different lighting conditions were used to record the underwater objects, including high and low illumination levels (which simulated the underwater environment as much as possible). In addition, to broaden the variety of the dataset, the above objects were used outside the towing tank for image collection. Subsequently, images were extracted from the video files, which resulted in the creation of around 2670 photographs in total.

In addition to the original laboratory dataset, a supplementary 300-image dataset was prepared to serve as the foundation for image classification algorithms. The items in these 300 photographs were manually labelled, and these labelled objects were subsequently removed to generate a dataset including images of five-class object classes. The image classification challenge required only particular items and not the complete scene to be present in the photographs. In order to achieve this, a script was created to crop the photographs in the dataset. As inputs, the custom software received the image to be cropped and the file containing the data of the item labels, and it cropped the objects on the provided image. Finally, the script was applied on the prior dataset, resulting in the generation of 2372 object pictures.

In order to avoid overfitting and to develop a more robust and accurate image classification model, the original dataset of 300 photographs was increased to 550 images. The same procedure for labelling and extracting the objects from photographs was used, and in the end, the total object dataset had 4700 pictures. The images and objects that were extracted to construct the first and second datasets are shown in Table 1, and as was indicated previously, the combination of the two datasets includes 550 individual photographs and 4700 unique class objects. The distribution of the five objects that were used to build the object dataset for the image classification task can be seen in Figure 2. The steps followed from collecting the video data in the towing tank to the generation of the object images dataset are presented in Figure 3. The custom laboratory datasets are available on the (figshare repository https://doi.org/10.6084/m9.figshare.20944354.v1, accessed on 15 August 2022).

Table 1. Towing Tank Image Datasets.

	Number of Images	Number of Objects
First set	300	2372
Second set	250	2328
Total set	550	4700



Figure 2. Object Class Distribution.

Last but not least, open source underwater datasets were used to create the dataset that contains the underwater environment. To be more specific, 1500 photos were utilised from the UFO-120 dataset [40], while 1170 images were collected from films of the Deep Sea Debris Dataset [39] from the Japan Agency for Marine-Earth Science and Technology.



Figure 3. Data Processing Cycle.

2.4. Data Augmentation

Traditional data augmentation refers to the process of improving visual data using various machine learning frameworks and libraries. In deep learning, and notably in computer vision applications, data augmentation is utilised extensively to enhance the number of datasets as well as the diversity of pictures. For the purpose of this research, the Albumentations [16] library was used to add additional features to the photographs included in the dataset. Albumentations is a well-established library that enables a range of image transformations and augmentation methods. It is specially tailored to work with any of the current Machine Learning frameworks. Therefore, it can be used with almost any of these frameworks.

As a result, the goal was to apply data augmentation in the original dataset, which consisted of 550 images featuring underwater items, to expand the dataset's variety and richness. The enhancement consisted of methods such as rotating and flipping the picture horizontally, increasing the image's saturation, increasing the exposure, and adding noise, as well as transforming it to greyscale. A total of 1980 pictures were produced as a consequence of applying the enhancement techniques to the dataset in a random fashion.

2.5. Deep Learning Augmentation

The second approach of data augmentation makes use of DL models, especially the use of GAN models. GAN models are essentially two neural networks competing to make more accurate predictions by generating their own training data and automatically detecting and learning patterns to create new samples that plausibly may have been selected from the original dataset. This makes them perfect for data augmentation. A CycleGAN model can take the images that were collected from the towing tank as its first input and an image of an underwater environment as its second input in order to generate outputs that will contain the images that were collected from the towing tank, but with characteristics that are unique to the underwater environment. The fact that this technique is an unsupervised learning process [21] means that it may utilise the original dataset without the need to perform any augmentation; nonetheless, it will need photographs that depict the underwater environment in order to function properly. The specifics of the CycleGAN implementation that were utilised to produce the augmented pictures are shown in Figure 1.

Figure 4 illustrates some of the image results that were achieved by employing DL data augmentation. It is evident that these images are far more accurate representations of the subsea environment than the initial data augmentation described in the results section when performing model evaluation with the FID technique. The model requires more optimisation, since in some of the photographs, it might be challenging to identify the individual objects that are there.



Figure 4. CycleGAN model images output.

2.6. Image Classification

The primary aim of the DL data augmentation process is to produce pictures of the type X_{lab}^{fake} that incorporate features from the domain X_{uw}^{real} . The last phase involves training an object identification model to determine how well it can identify the items of interest in the data. Before applying the object detection model to the simulated underwater photographs, an image classification model was developed to determine whether or not the five-class objects had been correctly categorised.

The most essential parameters in the CNN model are listed in Table 2. The first step is to crop the input image dimensions to 150 pixels on the horizontal (width, w) and vertical (height, h) planes, and it contains three colour channels (RGB colour depth). The next step is feature extraction, which happens as the image is processed through the convolutional layers (Conv layer). After the first Conv layer, the channels are increased to 16. After the second Conv layer, they are increased to 32, and after the third Conv layer, they are increased to 64. Additionally, the Max Pooling technique [43] is used after each convolution. This leads to a decrease in the size of the original input, which decreases from $w \times h = 150 \times 50$ pixels to $w \times h = 18 \times 18$ pixels. Next, the output of the Max Pooling layer has to be flattened so that it can appropriately link to the Dense layers, and then

lastly, to the classification layer, where the prediction will be made. The details of the image classification CNN model are shown in Figure 5.



Figure 5. Image Classification model.

Table 2. CNN model implementation.

Layer	Filters	Output
Input image	3	150 imes 150
Convolutional	16	150×150
Max Pooling	16	75 imes 75
Convolutional	32	75 imes 75
Max Pooling	32	37 imes 37
Convolutional	64	37 imes 37
Max Pooling	64	18 imes18
Flatten	20,736	18 imes18
Dense	128	18 imes18
Dense	5	18 imes18
Softmax	Prec	liction

It is vital to change the original basic CNN model in a manner that will not be prone to overfitting to increase the model's overall accuracy and decrease the amount of loss that occurs during training and validation. This objective may be accomplished by performing several fundamental picture modifications to random images included within the dataset. These transformations include image rotation, image flip, and zoom. After the final Conv layer, an additional layer called a dropout layer is also added. This layer will retain just a subset of the filters or "neurons" and remove the others from the network. The Dropout parameter is now set to 0.2 in this arrangement, which indicates that twenty per cent of the filters will not be used.

Two different models were used during the preliminary training. The first model implemented a basic CNN architecture, consisting of three convolutional layers, and it was used to classify images. The first model was then used in the development of the second model, with the exception of the application of certain fundamental transformations prior to the input to the network and the employment of the dropout layer after the convolution block (Convolutional-Max Pooling). In addition, the dataset used for the CNN model was segmented into training and validation sets, each of which included eighty and twenty percent of the total 4700 object picture dataset, respectively. After this, each model was trained and validated using these two sets.

The results of running the first CNN model on the initial object dataset consisting of 4700 pictures are shown in Figure 6. The model structure may be seen in Table 2. The model's accuracy improves at a comparable pace over the whole training period of 100 epochs in both the training and validation sets; however, the validation accuracy is somewhat lower than the training accuracy. Training loss and validation loss follow the same diminishing trend, with the validation loss starting a little higher than the training loss.



Figure 6. CNN model results trained on 4700 images.

2.7. Object Detection

Object detection models were used in the last phase to identify the five different class items. Initially, the YOLO v1 algorithm [12] was selected, but this was later upgraded to the YOLO v4 algorithm [44]. Based on the original paper [12], an implementation of the YOLO v1 model was selected primarily due to its speed during the training of the model, the simplicity of the model architecture, which involved a single forward CNN architecture, and lastly, its efficiency, because it involves fewer detection bounding boxes while training in comparison to other more demanding models such as R-CNN models [9]. The YOLO v4 model was then used to produce more accurate detection results on the dataset at shorter training times and with greater hardware utilisation. The utilised code is publicly accessible from the authors [13].

The training and assessment of all models were carried out on a double NVIDIA Tesla V100 GPU using the Rocket High-Performance Computing (HPC) Service offered by Newcastle University.

3. Results

The approach used in this study to construct a bespoke dataset comprised of lab photographs and including underwater features has been outlined in Section 2. Essentially, the method takes an input picture called X_{uw}^{real} and then passes this image on to the generator G_{uw} in order to extract the features, and finally rebuilds the image called $X_{uw}^{rebuild}$. The extracted features, together with the input image X_{lab}^{real} are then used in the generator G_{lab} to generate X_{uw}^{fake} , which is then assessed by the discriminator D_{lab} . Lastly, the extracted features are evaluated by the discriminator D_{lab} . In addition, the outputs of the models are applied to the image classification model, as discussed.

The UWCycleGAN model was trained on both the full-scale and object image datasets. The training of the model was performed on the Rocket HPC. Some key experimental details include the size of each dataset, the image size, the training time required to complete the task, the batch size, and the learning rate of the model. Table 3 gives a summary of those experimental details. For the full-scale image dataset, the dimensions of the input images were 256×256 pixels with batch size 4 images per batch and the learning rate was 2×10^{-4} . The model training was performed on an NVIDIA Tesla V100 GPU, consuming 14.2 GB of RAM for 18 h. Similarly, the object image dataset was trained on the same hardware, but the input image was set to 128×128 pixels with a batch size of 8 images per batch and a learning rate of 2×10^{-4} , consuming 15 GB RAM for 22 h of training. Additionally,

different image transformations were used in the model during the preprocessing phase to allow the model to generalise better in an unseen dataset. Such transformations include Gaussian Blur, Horizontal Flip, and Random Rotation.

Table 3. UWCycleGAN model experimental details.

CycleGAN	Dataset Size	Image Size	Hardware	RAM Usage	Training Time	Batch Size	Learning Rate
Full Image	2670	256 × 256	Tesla V100	14.2 GB	18 h	4	$2 imes 10^{-4}$
Object Image	4700	128×128	Tesla V100	15.0 GB	22 h	8	$2 imes 10^{-4}$

3.1. Model Evaluation

The end output of an adversarial model is an image and, more specifically for the purpose of this study, the result of the UWCycleGAN model is the false image X_{lab}^{fake} . Because of this, the generated output image X_{lab}^{fake} needs to be compared to the real input image X_{real} in order to determine whether the artificial image is an accurate representation of the underwater domain. The FID technique, which is the metric that produces the distance between two feature vectors of the X_{lab}^{fake} and X_{real} , is the most effective approach to achieve such an analysis [45].

The FID metric is described by the following equations:

$$\delta^2(\mu_{1,2}, C_{1,2}) = M_{1,2} + \tau_r(C_{1,2}) \tag{5}$$

$$M_{1,2} = \|\mu_1 - \mu_2\|_2^2 \tag{6}$$

$$C_{1,2} = (C_1 + C_2 - 2\sqrt{(C1 * C2)})$$
(7)

where δ is the Frechet distance, which is also known as the Wasserstein-2 distance [41]; μ_1 and μ_2 are the feature-wise mean of the real and the fake output images. C_1 and C_2 are the covariance matrices of the real and generated images, and τ_r is the trace linear operation of the square matrices. The FID criterion should be as low as practicable and, therefore, should be zero in the case of identical photos. In addition, the assessment of the Frechet distance is predicated on the first implementation in [41], which was carried out using the PyTorch ML framework [46].

The outcomes of the UWCycleGAN model are shown in Table 4. The UWCycleGAN model was tested with two distinct varieties of images: the first test was performed on the dataset consisting of the original 2670 images that were gathered in the laboratory, and the second test was carried out on the dataset consisting of the 4700 objects that were extracted from those original images. Finally, the FID evaluation was carried out on the images produced by the classical data augmentation. Given that the augmentation is carried out directly on the dataset, the only metrics that can be compared are those that compare X_{uw}^{real} vs. X_{uw}^{fake} and X_{lab}^{real} vs. X_{lab}^{fake} , respectively.

Table 4. FID scores.

FID Score		X_{uw}^{real} vs. X_{lab}^{real}	X_{uw}^{real} vs. X_{uw}^{fake}	X_{lab}^{real} vs. X_{lab}^{fake}	X_{uw}^{real} vs. $X_{uw}^{rebuild}$	X_{lab}^{real} vs. $X_{lab}^{rebuild}$	X _{real} vs. X ^{fake} _{lab}
UWCycleGAN	Full Image Object Image	301.05 326.10	21.92 29.19	20.31 36.79	11.93 20.58	7.38 15.90	42.12
Classical Augnentation		-	174.26	212.91	-	-	-

As stated previously, a lower value of the FID metric implies that the images are more comparable to one another. The FID value is high for images that are entirely distinct from one another, denoted by the notation X_{uw}^{real} and X_{lab}^{real} , respectively, as shown in Table 3, but it is much lower for images that are comparable to one another. The score may drop as low as 7.38 when comparing X_{lab}^{real} vs. $X_{lab}^{rebuild}$ and X_{uw}^{real} vs. $X_{uw}^{rebuild}$, respectively. Each picture that was processed using the cycleGAN model has a counterpart that was stored in the

object dataset. However, the score for each comparison is higher because the object dataset contains images that include only the area of interest (pipes, flanges, etc.). This makes it significantly more difficult for the model to perform as satisfactorily as it did in the first scenario, which used the entire image.

The most important comparison is X_{lab}^{real} and X_{lab}^{fake} , where the FID score is 20.31. This is because the goal is to make images that include underwater features. Given this score, it can be deduced that the fake lab underwater picture is quite similar to the real underwater image. Equivalent results were obtained with the object dataset; however, with the classical data augmentation, the values are significantly different due to the absence of underwater features in the images, which instead make use of various transformations such as hue transformation, image blur, noise, and saturation.

Additionally, to test the artificial underwater pictures that the UWCycleGAN model produced, actual underwater photographs were employed. These pictures were taken from reports of ship hull cleaning processes that were found on MAST Maritime Services' website [47]. The photos utilised for the FID assessment are shown in Figure 7. The photographs of the ship's hull before the cleaning procedure are shown in the top row, while those produced by the UWCycleGAN model are displayed in the bottom row. The analysis of these pictures can be seen in Table 3, which presents a side-by-side comparison of the underwater picture taken in the actual world and the one created in a lab. The FID score of 20.31 for X_{lab}^{real} vs. X_{lab}^{fake} is lower than the score of 42.12 when comparing X_{lab}^{fake} and X_{real} . This discrepancy is to be expected given that the underwater photographs were taken in the actual subsea environment and only show the unclean portion of the ship hull, which does not include any items such as those found in the laboratory data. As a result, these pictures do not have the same level of detail as the lab items, resulting in lower ratings.



Figure 7. Real and artificial underwater images.

3.2. Object Detection

After the image classification model had produced the anticipated results, it was necessary to continue to the next step, which was the training of the object detection model. As discussed in Section 2, YOLO v1 and YOLO v4 were the models that were selected.

3.2.1. YOLO v1

In order for the YOLO v1 algorithm to function appropriately on the towing tank dataset, the first version of the method had to be updated. It was necessary to change the parameters for the input classes because the initial model was trained on the Pascal VOC dataset, which contains 20 classes [12], but the towing tank dataset only has five, and the dataset that was used for the object detection model was the initial towing tank dataset, which contains 550 images. The object identification process requires a significant amount of computing resources due to the fact that the algorithm must pinpoint several items inside a picture. Hence, the Rocket HPC was used for this phase of the project. The

YOLO v1 was trained on a single NVIDIA Tesla V100 GPU, and it took approximately four hours to complete the task.

3.2.2. YOLO v4

The aim is to develop computer vision models that apply to as many real-world scenarios as possible, and this can be achieved by vigorously training the model and allowing it to generalise, which will allow it to use unseen data. Building deep learning models that predict an environment that may not have prior knowledge, particularly the underwater environment, is quite challenging. Training a DL model with a limited dataset may lead to overfitting, and hence poor results compared to a larger image dataset. Given the uneven or sparse sampling of points in the high-dimensional input data space, small datasets may also provide a more difficult mapping challenge for neural networks to learn. Adding noise to inputs during training is one way to make the input space smoother and easier [48,49]. Figure 8 shows the effects of the salt-and-pepper noise on the training dataset. During the training of the YOLO v4 object detection model, the following augmentation techniques were applied:

- 50% probability of horizontal flip.
- Random Gaussian blur between 0 and 1.25 pixels.
- Salt-and-pepper noise was applied to 8 per cent of pixels.



Figure 8. YOLO v4 Dataset preprocessing.

The amount of data input into an object detection model is significant, and this results in the need for computing capacity, such as graphics processing units (GPUs), to reduce the training time. Even if YOLO v1 achieved good results and can detect the objects it was trained on, it still needs to be utilised for the object detection task to be as efficient and effective as possible and reduce the training time as much as possible. To achieve better model performance, the algorithm needs to function properly on the available hardware. The most reliable option to solve this problem is to use an already optimised detection algorithm such as the YOLO v4 model. This model delivers more accurate results for the object identification tasks, and shorter training periods. In comparison with the initial YOLO v1 training time, the YOLO v4 was able to achieve the same results in one hour. The YOLO v4 model that was used for the training was based on the official model [44]. Experimental training details of YOLO v1 and YOLO v4 are shown in Table 5.

Table 5. YOLO object detection model experimental details.

Object Detection	Dataset Size	Image Size	Hardware	RAM Usage	Training Time	Batch Size	Learning Rate
YOLO v1	1128	416×416	Tesla V100	4.2 GB	3.8 h	16	0.01
YOLO v4	1128	416×416	Tesla V100	2.1 GB	1 h	16	0.01

The model was altered to function based on the five classes; also, parameters such as the batch number for the GPU that was utilised and the picture input size were updated to conform to the model's needs. Figure 9 shows the model's outcomes in terms of training and validation Loss. As can be seen from the plot, both losses follow the same decreasing trend, and the model converges with just a minor difference between them. In a perfect world, the training Loss and validation Loss should be equal; if they are not, this is an indicator of some overfitting. The gap between them remains the same after approximately epoch 50, and there is no indication that the validation loss will deviate further compared to the training Loss. Nevertheless, it is impossible to achieve ideal results, and some overfitting is always acceptable.

Finally, the inference YOLO v4 model, which was trained on the laboratory dataset, was used to validate the UWCycleGAN model's artificial images. Figure 10 presents the results of the detection model. The model is able to identify the various objects in the majority of the photographs; nevertheless, there are some instances in which it misclassifies the objects in the image; for example, in the first image, the model identifies the pipe as a lead block. The model needs to be trained in a larger and more diverse dataset to achieve better results and increased accuracy. The model's performance shows that it is able to detect the objects in the images, and it can be used to identify the objects in the artificial images generated by the UWCycleGAN model. However, future work focused on collecting relevant real-world underwater data to use for the training of the object detection model should be conducted.



Figure 9. YOLO V4 loss.



Figure 10. YOLO object detection.

4. Conclusions

Inspection and maintenance of underwater structures are among the numerous underwater-related applications that encounter difficulties due to a lack of data. Researchers are tackling the issue using new technologies made available by the rapid developments in the field of artificial intelligence in recent years. However, they have primarily focused on aspects of image enhancement or restoration. This is because of the difficulties that have arisen as a result of the lack of publicly available datasets from underwater environments. This paper attempts to solve the problem of easily accessible underwater datasets in light of the growing complexity of the underwater world and the absence of a tool to produce artificial underwater photos.

The main contribution of the present work is the development of a DL model used to generate images with underwater features. The model uses images of objects that can be found in underwater structures, such as pipes, anodes, flanges, bolts and nuts, taken under lab conditions, and images containing underwater scenes taken from public datasets.

As the research has demonstrated, it is clear that the UWCycleGAN model can be used to generate images with underwater features. Furthermore, the underwater domain of the artificially generated images was validated against real-world underwater images of the underwater ship hull section. The results between the generated and real-world images shown that the model can generate realistic underwater features. Finally, to investigate if the artificial underwater images can be used for object detection, they were tested on the initial YOLO v4 object detection model and the results show that the model can generalise satisfactorily and detect the objects. Since object detection is crucial for modern underwater operations such as underwater structural inspection and maintenance, the proposed method can be used to create rapidly datasets containing the desired objects and features to test the initial performance of such applications. Although it is possible to produce synthetic images of the underwater environment only using custom images and perform object detection with good results, to complete validation of the method proposed in this paper, new real-world underwater images from offshore structures need to be collected. The two real-world and artificial datasets will be compared for similarities using the FID metric introduced in Section 3.1, and then the YOLO v4 model will be trained on the artificial dataset and tested on the real-world dataset.

Author Contributions: Conceptualisation, I.P., M.H., R.N. and D.T.; methodology, I.P., M.H., R.N. and D.T.; software, I.P.; validation, I.P.; formal analysis, I.P.; investigation, I.P.; resources, I.P.; data curation, I.P.; writing—original draft preparation, I.P.; writing—review and editing, I.P., M.H., R.N. and D.T.; visualization, I.P.; supervision, M.H., R.N. and D.T.; project administration, I.P., M.H., R.N. and D.T.; funding acquisition, I.P., M.H., R.N. and D.T. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by EPSRC Doctoral Training Programme grant number EPN5095281.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets for the UWCycleGAN, and YOLO v4 models are available at: https://doi.org/10.6084/m9.figshare.20944354.v1, accessed on 15 August 2022.

Acknowledgments: The authors would like to express our gratitude to the staff of Hydrodynamics Lab at Newcastle University for their support and help during the image acquisition in the towing tank.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AIArtificial Intelligence.AUVsAutonomous Underwater Vehicles.

CNN	Convolutional Neural Network.
CV	Computer Vision.
DL	Deep Learning.
DNN	Deep Neural Networks.
EM	Electron Microscopy.
GANs	Generative Adversarial Networks.
HPC	High Performance Computing.
ROVs	Remote Operated Vehicles.
UWCycleGAN	UnderWaterCycleGAN.
X ^{real}	Real-world underwater environment image.
$X_{uw}^{rebuild}$	Reconstructed real world underwater environment image.
X ^{real}	Laboratory image.
X ^{fake}	Artificial underwater image.
D_{uw}	Underwater environment generator.
D_{lab}	Laboratory environment generator.

References

- 1. Ridolfi, A.; Conti, R.; Costanzi, R.; Fanelli, F.; Meli, E. A Dynamic Manipulation Strategy for an Intervention Autonomous Underwater Vehicle. *Adv. Robot. Autom.* **2015**, *4*. [CrossRef]
- Corchs, S.; Schettini, R. Underwater Image Processing: State of the Art of Restoration and Image Enhancement Methods; Hindawi Publishing Corporation: London, UK, 2010; Volume 14. [CrossRef]
- Han, F.; Yao, J.; Zhu, H.; Wang, C. Underwater Image Processing and Object Detection Based on Deep CNN Method. J. Sens. 2020, 2020, 6707328. [CrossRef]
- 4. Cazzato, D.; Cimarelli, C.; Sanchez-Lopez, J.L.; Voos, H.; Leo, M. A Survey of Computer Vision Methods for 2D Object Detection from Unmanned Aerial Vehicles. *J. Imaging* **2020**, *6*, 78. [CrossRef]
- 5. Chen, L.; Tong, L.; Zhou, F.; Jiang, Z.; Li, Z.; Lv, J.; Dong, J.; Zhou, H. A Benchmark dataset for both underwater image enhancement and underwater object detection. *arXiv* 2020, arXiv:2006.15789.
- 6. Xu, Z.; Haroutunian, M.; Murphy, A.J.; Neasham, J.; Norman, R. An Underwater Visual Navigation Method Based on Multiple ArUco Markers. *J. Mar. Sci. Eng.* 2021, *9*, 1432. [CrossRef]
- Yang, H.; Xu, Z.; Jia, B. An Underwater Positioning System for UUVs Based on LiDAR Camera and Inertial Measurement Unit. Sensors 2022, 22, 5418. [CrossRef]
- 8. Ludvigsen, M.; Sørensen, A.J. Towards Integrated Autonomous Underwater Operations for Ocean Mapping and Monitoring. *Annu. Rev. Control.* **2016**, *42*, 145–157. [CrossRef]
- 9. Girshick, R.; Donahue, J.; Darrell, T.; Malik, J. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. *arXiv* 2014, arXiv:1311.2524.
- 10. Girshick, R. Fast R-CNN. *arXiv* 2015, arXiv:1504.08083.
- 11. He, K.; Gkioxari, G.; Dollár, P.; Girshick, R. Mask R-CNN. arXiv 2017, arXiv:1703.06870.
- Redmon, J.; Divvala, S.; Girshick, R.; Farhadi, A. You only look once: Unified, real-time object detection. In Proceedings of the Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; IEEE Computer Society: Washington, DC, USA, 2016; Volume 2016, pp. 779–788. [CrossRef]
- 13. Bochkovskiy, A.; Wang, C.-Y.; Mark Liao, H.-Y. Yolo v4, v3 and v2 for Windows and Linux. Versions 2, 3 and 4. 2021. Available online: https://github.com/AlexeyAB/darknet (accessed on 19 August 2022).
- 14. Lin, T.Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; Belongie, S. Feature Pyramid Networks for Object Detection. *arXiv* 2017, arXiv:1612.03144.
- 15. Perez, L.; Wang, J. The Effectiveness of Data Augmentation in Image Classification Using Deep Learning. *arXiv* 2017, arXiv:1712.04621.
- 16. Buslaev, A.; Parinov, A.; Khvedchenya, E.; Iglovikov, V.I.; Kalinin, A.A. Albumentations: Fast and flexible image augmentations. *Information* **2020**, *11*, 125. [CrossRef]
- Riba, E.; Mishkin, D.; Ponsa, D.; Rublee, E.; Bradski, G. Kornia: An open source differentiable computer vision library for pytorch. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Waikoloa, HI, USA, 3–8 January 2020; pp. 3674–3683.
- 18. Goodfellow, I.J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative Adversarial Networks. *arXiv* 2014, arXiv:1406.2661.
- 19. Park, J.; Han, D.K.; Ko, H. Adaptive Weighted Multi-Discriminator CycleGAN for Underwater Image Enhancement. *J. Mar. Sci. Eng.* **2019**, *7*, 200. [CrossRef]
- Ronneberger, O.; Fischer, P.; Brox, T. U-net: Convolutional networks for biomedical image segmentation. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Munich, Germany, 5–9 October 2015; Springer: Berlin/Heidelberg, Germany, 2015; pp. 234–241.

- 21. Li, J.; Skinner, K.A.; Eustice, R.M.; Johnson-Roberson, M. WaterGAN: Unsupervised Generative Network to Enable Real-time Color Correction of Monocular Underwater Images. *IEEE Robot. Autom. Lett.* **2017**, *3*, 387–394. [CrossRef]
- 22. Anwar, S.; Li, C.; Porikli, F. Deep Underwater Image Enhancement. arXiv 2018, arXiv:1807.03528.
- 23. Berman, D.; Avidan, S. Diving into Haze-Lines: Color Restoration of Underwater Images. In Proceedings of the British Machine Vision Conference (BMVC 2017), London UK, 4–7 September 2017.
- 24. Arnold-Bos, A.; Malkasse, J.P.; Kervern, G. A preprocessing framework for automatic underwater images denoising. In Proceedings of the European Conference on Propagation and Systems, Copenhagen, Denmark, 15–20 March 2020.
- 25. Sandfort, V.; Yan, K.; Pickhardt, P.J.; Summers, R.M. Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in CT segmentation tasks. *Sci. Rep.* **2019**, *9*, 16884. [CrossRef]
- 26. Fakhry, A.; Peng, H.; Ji, S. Deep models for brain EM image segmentation: Novel insights and improved performance. *Bioinformatics* **2016**, *32*, 2352–2358. [CrossRef]
- 27. Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; Rabinovich, A. Going Deeper with Convolutions. *arXiv* 2015, arXiv:1409.4842.
- Zhu, J.Y.; Park, T.; Isola, P.; Efros, A.A. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 22–29 October 2017.
- Geirhos, R.; Janssen, D.H.J.; Schütt, H.H.; Rauber, J.; Bethge, M.; Wichmann, F.A. Comparing deep neural networks against humans: Object recognition when the signal gets weaker. *arXiv* 2018, arXiv:1706.06969.
- Chen, X.; Zhang, P.; Quan, L.; Yi, C.; Lu, C. Underwater Image Enhancement based on Deep Learning and Image Formation Model. arXiv 2021, arXiv:2101.00991.
- 31. Li, H.; Li, J.; Wang, W. A Fusion Adversarial Underwater Image Enhancement Network with a Public Test Dataset. *arXiv* 2019, arXiv:1906.06819.
- 32. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. arXiv 2016, arXiv:1512.03385.
- Li, C.; Guo, C.; Ren, W.; Cong, R.; Hou, J.; Kwong, S.; Tao, D. An Underwater Image Enhancement Benchmark Dataset and Beyond. *IEEE Trans. Image Process.* 2019, 29, 4376–4389. [CrossRef]
- 34. Panetta, K.; Kezebou, L.; Oludare, V.; Agaian, S. Comprehensive Underwater Object Tracking Benchmark Dataset and Underwater Image Enhancement With GAN. *IEEE J. Ocean Eng.* **2021**, *47*, 59–75. [CrossRef]
- Han, J.; Shoeiby, M.; Malthus, T.; Botha, E.; Anstee, J.; Anwar, S.; Wei, R.; Armin, M.A.; Li, H.; Petersson, L. Underwater Image Restoration via Contrastive Learning and a Real-world Dataset. *Remote Sens.* 2022, 14, 4297. [CrossRef]
- 36. He, K.; Fan, H.; Wu, Y.; Xie, S.; Girshick, R. Momentum Contrast for Unsupervised Visual Representation Learning. *arXiv* 2020, arXiv:1911.05722.
- Li, C.; Anwar, S.; Porikli, F. Underwater Scene Prior Inspired Deep Underwater Image and Video Enhancement. *Pattern Recognit.* 2020, 98, 107038. [CrossRef]
- Islam, M.J.; Enan, S.S.; Luo, P.; Sattar, J. Underwater image super-resolution using deep residual multipliers. In Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA), Paris, France, 31 May–31 August 2020; pp. 900–906.
- 39. Japan Agency for Marine-Earth Science and Technology. Deep-Sea Debris Database JAMSTEC; JAMSTEC: Kanagawa, Japan, 2021.
- 40. Md Jahidul, I.; Peigen, L.; Junaed, S. *The UFO-120 Dataset*; Interactive Robotics and Vision Lab.: Minneapolis, MN, USA, 2022.
- 41. Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; Hochreiter, S. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. *arXiv* **2018**, arXiv:1706.08500.
- 42. Goodfellow, I. NIPS 2016 Tutorial: Generative Adversarial Networks. *arXiv* 2016, arXiv:1701.00160.
- Wu, H.; Gu, X. Max-pooling dropout for regularization of convolutional neural networks. In Proceedings of the International Conference on Neural Information Processing, Istanbul, Turkey, 9–12 November 2015; Springer: Berlin/Heidelberg, Germany, 2015; pp. 46–54.
- 44. Bochkovskiy, A.; Wang, C.Y.; Liao, H.Y.M. YOLOv4: Optimal Speed and Accuracy of Object Detection. *arXiv* 2020, arXiv:2004.10934.
- 45. Brownlee, J. How to Implement the Frechet Inception Distance (FID) for Evaluating GANs. Retrieved Dec. 2019, 5, 2019.
- Seitzer, M. Pytorch-Fid: FID Score for PyTorch. Version 0.2.1. 2020 Available online: https://github.com/mseitzer/pytorch-fid (accessed on 19 August 2022).
- 47. MAST Maritime Services S.A. Underwater Services. 2021. Available online: https://mastms.gr/ (accessed on 19 August 2022).
- 48. Vincent, P.; Larochelle, H.; Lajoie, I.; Bengio, Y.; Manzagol, P.A. Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion. *J. Mach. Learn. Res.* **2010**, *11*, 3371–3408.
- 49. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT Press: Cambridge, MA, USA, 2016. Available online: http://www.deeplearningbook.org (accessed on 19 August 2022).