

# DeepSeeColor: Realtime Adaptive Color Correction for Autonomous Underwater Vehicles via Deep Learning Methods

Stewart Jamieson<sup>1,2</sup>, Jonathan P. How<sup>2</sup>, and Yogesh Girdhar<sup>3</sup>

**Abstract**—Successful applications of complex vision-based behaviours *underwater* have lagged behind progress in terrestrial and aerial domains. This is largely due to the degraded image quality resulting from the physical phenomena involved in underwater image formation. *Spectrally-selective light attenuation* drains some colors from underwater images while *backscattering* adds others, making it challenging to perform vision-based tasks underwater. State-of-the-art methods for underwater color correction optimize the parameters of image formation models to restore the full spectrum of color to underwater imagery. However, these methods have high computational complexity that is unfavourable for realtime use by autonomous underwater vehicles (AUVs), as a result of having been primarily designed for offline color correction. Here, we present *DeepSeeColor*, a novel algorithm that combines a state-of-the-art underwater image formation model with the computational efficiency of deep learning frameworks. In our experiments, we show that *DeepSeeColor* offers comparable performance to the popular “Sea-Thru” algorithm [1] while being able to rapidly process images at up to 60Hz, thus making it suitable for use onboard AUVs as a preprocessing step to enable more robust vision-based behaviours.

## OPEN-SOURCE SOFTWARE

The datasets collected for and used in this paper are available along with an implementation of *DeepSeeColor* at: <https://warp.who.i.edu/deepseecolor/>

## I. INTRODUCTION

The development of autonomous underwater vehicles (AUVs) has trended towards more complex and adaptive vision-based behaviours, enabled by the increasing performance and energy efficiency of onboard micro-computing resources over time. Many such behaviours, including visual target tracking [2]–[5], adaptive exploration of benthic environments [6]–[12], autonomous docking [13]–[15], and diver following [16] and assistance [17]–[20], strongly depend on an AUV’s capability for visual perception and understanding. Unfortunately, underwater image quality is generally lower and more variable than the quality of similar images taken in air, making vision-based autonomous behaviours far more challenging to perform underwater than in other domains.

The most prominent image quality issues result from spectrally-selective light attenuation and backscattering.

\*This work was supported by NSF-NRI Award Number 1734400, and by grants from NVIDIA. This work utilized an NVIDIA RTX A6000 GPU.

<sup>1</sup>S. Jamieson is with the MIT-WHOI Joint Program in Oceanography/Applied Ocean Science and Engineering [sjamieson@who.i.edu](mailto:sjamieson@who.i.edu)

<sup>2</sup>S. Jamieson and J. P. How are with the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology (MIT) [{sjamieson, jhow}@mit.edu](mailto:{sjamieson, jhow}@mit.edu)

<sup>3</sup>Y. Girdhar is with the Woods Hole Oceanographic Institution (WHOI) Applied Ocean Physics and Engineering Department [yogi@who.i.edu](mailto:yogi@who.i.edu)




Fig. 1: Best viewed on a screen, in color. Images taken in underwater environments (left column) suffer from severe loss of colors due to the combination of spectrally-selective light attenuation and backscattering due to particles in water. This effect is more pronounced as the distance between the camera and imaging target increases. This paper presents *DeepSeeColor*, an efficient technique to reconstruct true-color images (right column) for use in realtime onboard decision making by an autonomous underwater vehicle.

These phenomena cause images taken underwater to be lacking in certain colors, typically red, and to have an excess of others, such as blue and green, as seen in Fig. 1. The magnitude of these effects is correlated with the camera’s water depth and range to the target; objects in the image that are shallow and close to the camera will have more of their natural color than objects that are deeper or further away. However, the specific color distortions in an image depend also on the presence of other light sources (e.g., the sun), and on the optical properties of the camera and the body of water in which the image was captured.

The performance of computer vision algorithms for image classification and 3D reconstruction is severely degraded by color distortions, regardless of whether those algorithms are based on deep neural networks or on more traditional feature detection and matching methods [21], [22]. Even learning-based systems trained to be robust to targets within some range of distances, lighting conditions, and water masses will generally not be robust to real-world targets outside of the training distribution. This makes realtime, adaptive, *in situ* color correction a necessary prerequisite for successful vision-based AUV tasks and behaviours, espe-

cially in potentially safety-critical applications such as diver assistance. This is a gap in the underwater vision field; while previous works have focused on high-quality offline color reconstruction [1], [23]–[27], to our knowledge no methods have been developed specifically for realtime applications.

This paper presents *DeepSeeColor*, a novel method for realtime and adaptive color correction of images onboard an AUV. This realtime preprocessing enables more robust execution of *in situ* autonomous behaviours such as object recognition and tracking, semantic mapping, and mission summarization. *DeepSeeColor* significantly improves upon the successful “Sea-Thru” algorithm [1] in terms of computational efficiency. It achieves this improvement through using highly efficient gradient-based optimization methods to learn the weights of two simple convolutional neural networks, where these weights correspond to a physics-based image formation model’s unknown parameters. This novel approach replaces far more computationally expensive image processing operations, while leveraging GPU acceleration and other built-in performance optimizations provided by popular deep learning frameworks. We further present validation of *DeepSeeColor* on the dataset provided with [1], as well as on a new dataset of stereo-imagery collected from AUV deployments at coral reef sites in the US Virgin Islands.

## II. BACKGROUND & RELATED WORKS

### A. Underwater Image Formation

A digital camera system creates images by measuring the intensity of incoming light reflected from various targets in its lens’ field of view. Assuming the light originates from a broadband source similar to the sun, the wavelengths of the reflected light accurately represent the colors of the targets. A color filter array on the camera’s image sensor enables measuring the intensity of light within the red, green, and blue wavelength “channels”<sup>1</sup> at each pixel.

Regardless of wavelength, the apparent intensity of light decreases in relation to an observer’s distance from its source according to the inverse-square law. However, the presence of an intervening medium further attenuates the apparent intensity of a target through *absorption* and *scattering*. Furthermore, many media are *spectrally-selective*: they preferentially attenuate some wavelengths of light over others. Water is much more spectrally-selective than air: red light is attenuated over a distance of 1m about as much as green light is over 10m, or blue light is over 100m [28].

*Backscattering* is the process by which particles in the water column between the target and camera reflect light from sources other than the target into the camera. This is particularly problematic at shallow depths (<100m), long target ranges (>1m), and in turbid waters.<sup>2</sup> At shallow depths, light from the surface (e.g., sunlight) is scattered in all directions by tiny particulate in the water; the spectrum of this scattered light consist of wavelengths that were not

absorbed higher up in the water column, and thus tends to be blue or blue-green. The greater the range to the target, the more intervening particulate there is to backscatter this light into the camera; the resulting “haze” in the image makes it more difficult to see the target, especially when it saturates the camera’s dynamic range.

We denote the intensity of light in channel  $c \in \{R, G, B\}$  emitted by the target along the ray to pixel  $(i, j)$  in the camera image sensor as  $J_c(i, j)$ , and the intensity measured by the sensor at that pixel as  $I_c(i, j)$ . Spectrally-selective light attenuation and backscattering can be modelled together using the underwater image formation model,

$$I_c(i, j) = J_c(i, j)A_c(i, j) + B_c(i, j), \quad (1)$$

where  $A_c(i, j)$  and  $B_c(i, j)$  represent wavelength-dependent light attenuation and backscatter, respectively. This reflects the form of models used by many works in underwater color reconstruction [23]–[27], which take  $A$  and  $B$  to be

$$A_c(i, j) = \exp(-a_c \cdot z_{i,j}), \quad (2)$$

$$B_c(i, j) = \gamma_c^\infty (1 - A_c(i, j)), \quad (3)$$

where  $z_{i,j}$  is the range of the target. The values of  $a_c, \gamma_c^\infty \in \mathbb{R}_{\geq 0}$  are determined by the camera system and environmental parameters, including as the water type, target reflectance, illumination sources, image sensor characteristics, and camera depth, which are, for now, all assumed to be fixed.

More recently, [29] found that Eq. (2), derived from an atmospheric dehazing model, neglects the range-dependence of underwater light attenuation coefficient  $a_c$ , and incorrectly assumes that the coefficients governing the range-dependence of attenuation and backscattering are the same. [29] presented a new model capable of capturing these complexities,

$$A_c(i, j) = \exp(-a_c(z_{i,j}) \cdot z_{i,j}), \quad (4)$$

$$B_c(i, j) = \gamma_c^\infty (1 - \exp(-\beta_c z_{i,j})), \quad (5)$$

with scalars  $\gamma_c^\infty, \beta_c \in \mathbb{R}_{\geq 0}$  and a parametric function  $a_c(z)$ .<sup>3,4</sup> This is the model used by Sea-Thru [1] and *DeepSeeColor*, and will be explored further in Section III.

### B. Methods for Color Reconstruction

Spectrally-selective light attenuation and backscattering have been considered by underwater color reconstruction algorithms since the seminal works of [23]–[25]. These works presented each presented a method to estimate their respective image formation model parameters from a single image. Given these parameter estimates, the true color of the target at pixel location  $(i, j)$  can be reconstructed as

$$J_c(i, j) = D_c(i, j)A_c(i, j)^{-1}, \quad (6)$$

where  $D_c(i, j) := I_c(i, j) - B_c(i, j)$  denotes the “direct” signal. The components of this approach to color reconstruction are depicted in Fig. 2. More recent works have explored techniques for better estimating the parameters

<sup>1</sup>These channels are roughly defined as the wavelength ranges 575-725nm (red), 475-625nm (green), and 400-550nm (blue), respectively.

<sup>2</sup>Note “depth” refers to the vertical distance of the target from the water surface, while “range” refers to the distance from the camera to the target.

<sup>3</sup> $\beta_c$  is most accurately modelled as a function of the range  $z$ , like  $a_c(z)$ , but this effect is negligible within a fixed water type [29].

<sup>4</sup>Details on the interpretation of each parameter are given in [1].




Fig. 2: Examples of an input image  $I$  and corresponding range map  $Z$  (color-coded such that blue is “close” and red is “far”), best viewed on a screen in color. The DeepSeeColor method efficiently learns to approximate the backscatter image  $B$  and attenuation factor image  $A$  (both shown with contrast enhanced), enabling recovery of the true-color image  $J$ .

$a_c, \gamma_c^\infty$  through leveraging additional information, such as the optical properties of known water profiles [30].

Among the most relevant of prior works is [27], which recognized that leveraging an accurate range map, which describes the distance of each pixel in an image, could better constrain estimates of the image formation model parameters. The method was designed for use on AUV collected data, and applied structure-from-motion estimation to generate range maps based on the overlapping field-of-view and real-world camera displacement between pairs of captured images [27]. This displacement was estimated using navigational information collected by other sensors on the AUV, but estimation could be avoided by using a synchronized stereo-camera system to simultaneously capture pairs of images with a known, fixed camera displacement.

The previously discussed methods were developed for estimating the parameters of the traditional underwater image formation model described by Eqs. (2) and (3). Estimating the parameters of the improved underwater image formation model presented in Eqs. (4) and (5) originally required the use of multiple color chart calibration targets *in situ* [29]. The idea of using range maps, instead of color charts, to guide parameter estimation was adapted by the *Sea-Thru* algorithm [1], which enabled learning the improved model parameters using only the captured image and range map.

There has also been progress in learning-based methods for color correction [31], including for underwater imagery [32], [33]. However, these methods do not guarantee consistency or accuracy in the color reconstruction of an image stream, or images captured in different environmental conditions. This makes them less suitable than physics-based color correction methods for use in supporting realtime, possibly safety-critical, autonomous underwater behaviours.

### C. AUV Vision System Considerations

Though [27] specifically explored color reconstruction for images *collected* by AUVs, to our knowledge no prior works have developed robust, physics-based underwater color reconstruction methods intended to run *onboard* an AUV’s highly constrained computational resources in realtime. In fact, each of the aforementioned color reconstruction methods relies on solving optimization problems with computational complexity that is at least linear in the number of pixels in the input image [1], [23]–[27]. The significant computational complexity of Sea-Thru [1], the only prior work to solve the parameters of the improved image formation model in Eqs. (5) and (12), will be explored further in Section III.

Computationally expensive methods could be acceptable if the image formation parameters only needed to be computed once. However, prior works have found that these image formation parameters can change significantly in time and space due to factors including variation in depth, lighting conditions, exposure time, turbidity, and imaging angle [1], [29]. This demands the usage of an *adaptive* color correction method that is robust to changes in lighting and other environmental parameters, and can be run in realtime on each image as it is collected.

## III. THE DEEPSEECOLOR METHOD

The DeepSeeColor method estimates the backscatter and attenuation parameters of the underwater image formation model from [29] using a sequence of two convolutional neural networks [34] depicted in Figures 3 and 4, respectively. The networks are trained under self supervision using the captured image  $I$  and range map  $Z$ ; this training process can take advantage of highly energy-efficient deep learning hardware accelerators increasingly found onboard autonomous platforms [35].

### A. Backscatter Estimation

DeepSeeColor makes use of the same backscatter model as Sea-Thru,

$$\hat{B}_c(i, j) = \gamma_c^\infty (1 - \exp(-\beta_c \cdot z_{i,j})) + \eta_c \exp(-\alpha_c \cdot z_{i,j}), \quad (7)$$

which corresponds to the backscatter model presented in Eq. (5) augmented with a residual term, described in [1], characterized by two new scalar parameters  $\eta_c, \alpha_c \in \mathbb{R}_{\geq 0}$ .




Fig. 3: The backscatter net inputs a range map  $Z$  (visualized in color) and estimates the corresponding backscatter image  $\hat{B}$ . The kernel parameters in each convolutional layer map to the parameters of the backscatter estimation model in Eq. (7).

1) *Inference*: DeepSeeColor performs backscatter estimation by using the neural network presented in Fig. 3. We define a novel nonlinear activation function, Gated Exponential Decay (GED), as

$$\text{GED}(x) := \begin{cases} 1 & x \leq 0, \\ \exp(-x) & x > 0. \end{cases} \quad (8)$$

We also introduce its complement,  $\text{CGED}(x) := 1 - \text{GED}(x)$ . Observe that Eq. 7 can now be rewritten as

$$\hat{B}_c(Z) = \gamma_c^\infty \cdot \text{CGED}(\beta_c Z) + \eta_c \cdot \text{GED}(\alpha_c Z). \quad (9)$$

With its final sigmoid activation layer removed, the network depicted in Fig. 3 computes  $\hat{B}_c$  for the backscatter parameters encoded in its convolutional layers' kernels. The operations  $\beta_c Z$  and  $\alpha_c Z$  for  $c \in \{r, g, b\}$  are modelled as convolutions between the tensor  $Z = [z_{i,j}]$  with dimensions  $(W, H, 1)$  and a kernel of shape  $(1, 1, 1, 3)$  to produce an output tensor with dimensions  $(W, H, 3)$ , where the values of the kernel elements correspond to  $\beta$  and  $\alpha$ , respectively.<sup>5</sup> Similarly, multiplication by the coefficients in  $\gamma_c^\infty$  and  $\eta$  can be modelled as convolutions with kernels of dimension  $(1, 1, 3, 3)$ , where each kernel is constrained to be diagonal and thus has only 3 free parameters, resulting in the scaling of each channel by the corresponding  $\gamma_c^\infty$  or  $\eta_c$ . These parallel sequences of operations correspond to the top and bottom branches of the network in Fig. 3, respectively. The final sigmoid layer applies the activation function  $\sigma(x) :=$

<sup>5</sup>Note tensor shapes are given in order of (width, height, channels), while kernel shapes are given as (width, height, input channels, output channels).

$1/(1 + e^{-x})$ , ensuring that the backscatter estimates remain bounded in  $[0, 1]$  regardless of any outliers in the range map.

2) *Training*: DeepSeeColor leverages the assumption from [1], [36] that in any given range interval there are some pixels in the image which should have zero intensity. It trains the backscatter network using a novel loss function,

$$\mathcal{L}_{\text{bs}}(\hat{D}) = \sum_{(i,j)} \sum_c (\max\{\hat{D}_c(i,j), 0\} + k \min\{\hat{D}_c(i,j), 0\}), \quad (10)$$

with hyperparameter  $k > 1$  and  $\hat{D}_c := I_c - \sigma(\hat{B}_c)$ .

As  $k \rightarrow \infty$ , optimizing Eq. (10) becomes equivalent to finding backscatter parameters that minimize the minimum intensity of pixels at every range, while ensuring the network *never* predicts that a pixel in the direct signal would have negative intensity. As  $k \rightarrow 1$ , the loss function becomes tolerant of some pixels in the direct signal being predicted to have negative intensity, if that enables bringing more pixels closer to zero intensity. Intuitively, larger values of  $k$  would produce more accurate backscatter estimates if the range map was noiseless and the assumption of zero-intensity pixels at every range was satisfied, but are also less robust to noise and outliers in the image and range map. We find empirically that a value of  $k = 1000$  produces accurate backscatter estimates.

For comparison, the backscatter parameters are estimated in Sea-Thru by binning the image pixels into 10 evenly spaced range windows, and assigning the darkest 1% of pixels in each bin to the set  $\Omega$  [1]. Assuming these pixels should have zero-intensity, the backscatter parameters are then found by solving the nonlinear optimization problem,

$$\min_{\gamma_c^\infty, \beta_c, \eta_c, \alpha_c} \sum_{(i,j) \in \Omega} \sum_c |I_c(i,j) - \hat{B}_c(i,j)|^2. \quad (11)$$

Constructing  $\Omega$  using a bitonic merge sort has time complexity  $O(NP^{-1} \log^2 N)$  for  $N$  pixels in the input image, when run in parallel on a processor with  $P \leq N$  cores. In contrast, computing the loss function in Eq. (10) and backpropagating its gradients has  $O(NP^{-1} + \log P)$  time complexity.

## B. Attenuation Coefficient Estimation

The attenuation coefficient function  $a_c(z)$  is well approximated as a double exponential function [1], characterized using the parameters  $v_c, w_c, x_c, y_c \in \mathbb{R}_{\geq 0}$  as

$$a_c(z) = w_c \exp(-v_c \cdot z) + y_c \exp(-x_c \cdot z). \quad (12)$$

1) *Inference*: From Eqs. (4) and (12), it follows that

$$A_c(Z) = \text{GED}\left(Z \cdot (w_c \cdot \text{GED}(v_c Z) + y_c \cdot \text{GED}(x_c Z))\right). \quad (13)$$

Accordingly, the attenuation coefficient network depicted in Fig. 4 takes the same structure as the backscatter net, except the CGED and sigmoid activation functions are replaced with GEDs. The true-color image  $J$  is then estimated as,

$$\hat{J}_c(i,j) = \hat{D}_c(i,j) \cdot \exp(a_c(z_{i,j}) \cdot z_{i,j}). \quad (14)$$




Fig. 4: The attenuation net produces the attenuation map  $A$  from a range map  $Z$ . The kernel parameters in each convolutional layer map to the parameters of the attenuation coefficient function  $a(z)$  in Eq. (12).

2) *Training*: DeepSeeColor trains the attenuation network using a novel composite loss function designed to learn attenuation coefficient parameters similar to those of the Sea-Thru method,

$$\mathcal{L}_{ac}(\hat{J}; \hat{D}) = \mathcal{L}_{\text{saturation}}(\hat{J}) + \mathcal{L}_{\text{intensity}}(\hat{J}) + \mathcal{L}_{\text{var}}(\hat{J}; \hat{D}). \quad (15)$$

The saturation loss penalizes over-saturating pixels in the output image, and is defined as

$$\mathcal{L}_{\text{saturation}}(\hat{J}) = \frac{1}{3N} \sum_c \|\max\{\hat{J}_c - 1, 0\}\|^2,$$

where  $N$  is the number of pixels in  $\hat{J}$ . The intensity loss penalizes image channels which have a very low or very high average intensity, and is defined as

$$\mathcal{L}_{\text{intensity}}(\hat{J}) = \frac{1}{3} \sum_c \left( \frac{1}{N} \sum_{i,j} (\hat{J}_c(i,j) - 0.5) \right)^2.$$

Finally, the variation loss penalizes changing the variation in each color channel across the image, compared to the variation measured in the direct signal,

$$\mathcal{L}_{\text{var}}(\hat{J}; \hat{D}) = \frac{1}{3} \sum_c (s(\hat{J}_c) - s(\hat{D}_c))^2,$$

where  $s(M)$  is the standard deviation of the values in  $M$ .

Minimizing the intensity loss function is similar to following the popular “gray-world” approach to white balancing [37], except the network is constrained to modifying the attenuation coefficients. This drives it to modify each channel’s coefficients in a way that stretches contrast but

TABLE I: Grayscale Patch Mean Angular Error, in degrees.

Image	Raw	Sea-Thru	DeepSeeColor (ours)
D1_3272	26	8	14
D2_3647	26	8	10
D3_4910	22	8	5
D4_0209	23	4	4
D5_3374	17/16/15/17	4/3/5/3	9/10/10/11

biases the change towards targets far from the camera, and were thus most attenuated. The saturation and variation losses help to avoid very large attenuation coefficients; this is most relevant for images where attenuation has nearly eliminated a color channel such that very large coefficients would be required to recover it. Large coefficients cause over-saturation, and amplify noise in the input image; including the saturation and variation losses mitigates this issue.

This training process has  $O(NP^{-1} + \log P)$  time complexity per training iteration when parallelized on a processor with  $P \leq N$  cores. Furthermore, each iteration produces a color-corrected output image with increased quality, so the method is anytime optimal. For comparison, the Sea-Thru algorithm finds the attenuation model parameters using the LSAC algorithm [38] followed sequentially by a nonlinear optimization algorithm, which are each iterative methods with  $O(N)$  time complexity per training iteration.

## IV. EXPERIMENTAL RESULTS

The DeepSeeColor method has been implemented in PyTorch [39] and demonstrates strong performance on the Sea-Thru dataset [1] and on stereo-camera imagery collected with AUV deployments in the US Virgin Islands.

### A. Sea-Thru Dataset

The Sea-Thru dataset consists of 1157 raw images, each with a corresponding range map generated using structure-from-motion techniques [1]. DeepSeeColor was run on all of the images available on the hosting site,<sup>6</sup> and a sampling of the outputs are presented in Fig. 5.

As in [1], the accuracy of the color reconstruction is estimated using the angular error between true gray and the grayscale patches on the color charts present in many of the dataset images. This angular error is computed at each grayscale patch in the color chart as

$$\psi(i,j) = \cos^{-1} \left( \frac{\sum_c J_c(i,j)}{\left(3 \cdot \sum_c [J_c(i,j)]^2\right)^{\frac{1}{2}}} \right). \quad (16)$$

The average angular error averaged over the six grayscale patches on each color chart for each of Sea-Thru and DeepSeeColor is presented in Table I. As there is no public release of Sea-Thru, its errors are taken directly from [1].

<sup>6</sup>[http://csms.haifa.ac.il/profiles/tTreibitz/datasets/sea\\_thru/index.html](http://csms.haifa.ac.il/profiles/tTreibitz/datasets/sea_thru/index.html)




Fig. 5: Best viewed on a screen, in color. The left column shows 10 raw images from the Sea-Thru dataset [1]. The center column shows the corresponding images after a contrast stretch, while the right column shows the outputs of DeepSeeColor. DeepSeeColor significantly outperforms contrast stretching, especially in scenes with distant targets (e.g. 3647, 3681, and 3385), and achieves comparable performance to Sea-Thru (refer to Fig. 6 in [1] for comparison).

TABLE II: DeepSeeColor Runtime.

# Pixels	Backscatter (per iter)	Attenuation (per iter)	Total (per iter)	Max Frequency
0.7M	6.2 ms	9.9 ms	16.1 ms	62.1 Hz
2.4M	6.4 ms	10.5 ms	16.9 ms	59.2 Hz

### B. US Virgin Islands Dataset

The DeepSeeColor method was also evaluated to process imagery collected in the US Virgin Islands by the CUREE AUV [40]. The AUV was equipped with a downwards facing color stereo-camera, which enabled the collection of color imagery and range maps like those seen in Figs. 1 and 2. The post-processed imagery has comparable quality to that generated by Sea-Thru, but is generated at a very high rate.

As seen in Table II, DeepSeeColor can perform 60 training iterations of each network per second. This corresponds to performing up to 60 training iterations on a single image, or correcting up to 60 images per second once the image formation model parameters have stabilized. This offers significant flexibility for an AUV to dedicate more processing power to training the DeepSeeColor model when imaging conditions change and the image formation parameters need updating, and less when imaging conditions are stable. Interestingly, the training time is largely insensitive to the input image resolution up to 2.4 megapixels; this is a result of the fixed number of unknown model parameters and of the GPU processing capabilities of modern deep learning frameworks like PyTorch [39], which process many sections of an image in parallel.

### V. CONCLUSIONS & FUTURE WORK

In this paper we have presented DeepSeeColor, a novel color correction method that uses convolutional neural network training operations to learn the parameters of a physics-based underwater image formation model [29]. This approach to color correction is more robust than heuristic methods, and thus more appropriate for usage on AUVs performing safety-critical tasks. We demonstrated our proposed method on the Sea-Thru dataset [1], as well as on images collected during field experiments in the US Virgin Islands. DeepSeeColor was able to achieve comparable performance in color reconstruction to the Sea-Thru algorithm [1], with significant improvements made in terms of computational complexity. The results support that DeepSeeColor is well-suited for use in realtime preprocessing of imagery collected by an AUV to enable more robust execution of more sophisticated autonomous tasks. In the future, we plan to deploy DeepSeeColor on an AUV to evaluate how access to robust and realtime color correction improves mission performance on various tasks such as place recognition and target tracking.

### ACKNOWLEDGEMENTS

The authors thank Nathan McGuire, Seth McCammon, and Levi Cai for their assistance in the development and deployment of the CUREE platform [40], which generated much of the data used to develop and test DeepSeeColor.

## REFERENCES

- [1] D. Akkaynak and T. Treibitz, "Sea-Thru: A Method for Removing Water From Underwater Images," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Long Beach, CA, USA: IEEE, June 2019, pp. 1682–1691.
- [2] L. Cai, N. E. McGuire, R. Hanlon, T. A. Mooney, and Y. Girdhar, "Semi-supervised Visual Tracking of Marine Animals Using Autonomous Underwater Vehicles," *International Journal of Computer Vision*, Mar. 2023.
- [3] D. R. Yoerger, A. F. Govindarajan, J. C. Howland, J. K. Llopiz, P. H. Wiebe, M. Curran, J. Fujii, D. Gomez-Ibanez, K. Katija, B. H. Robison, B. W. Hobson, M. Risi, and S. M. Rock, "A hybrid underwater robot for multidisciplinary investigation of the ocean twilight zone," *Science Robotics*, vol. 6, no. 55, June 2021.
- [4] K. Katija, P. L. D. Roberts, J. Daniels, A. Lapides, K. Barnard, M. Risi, B. Y. Ranaan, B. G. Woodward, and J. Takahashi, "Visual tracking of deepwater animals using machine learning-controlled robotic underwater vehicles," in *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*. Waikoloa, HI, USA: IEEE, Jan. 2021, pp. 859–868.
- [5] Z. Hao, J. Qiu, H. Zhang, G. Ren, and C. Liu, "UMOTMA: Underwater multiple object tracking with memory aggregation," *Frontiers in Marine Science*, vol. 9, 2022.
- [6] S. Jamieson, J. P. How, and Y. Girdhar, "Active Reward Learning for Co-Robotic Vision Based Exploration in Bandwidth Limited Environments," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. Paris, France: IEEE, May 2020, pp. 1806–1812.
- [7] E. Guerrero, F. Bonin-Font, and G. Oliver, "Adaptive Visual Information Gathering for Autonomous Exploration of Underwater Environments," *IEEE Access*, vol. 9, pp. 136 487–136 506, 2021.
- [8] Y. Girdhar and G. Dudek, "Modeling curiosity in a mobile robot for long-term autonomous exploration and monitoring," *Autonomous Robots*, vol. 40, no. 7, pp. 1267–1278, 2016.
- [9] F. Bonin-Font, G. Oliver, S. Wirth, M. Massot, P. Lluís Negre, and J.-P. Beltran, "Visual sensing for autonomous underwater exploration and intervention tasks," *Ocean Engineering*, vol. 93, pp. 25–44, Jan. 2015.
- [10] B. Joshi, M. Xanthidis, M. Roznere, N. J. Burgdorfer, P. Mordohai, A. Q. Li, and I. Rekleitis, "Underwater Exploration and Mapping," in *2022 IEEE/OES Autonomous Underwater Vehicles Symposium (AUV)*. Singapore: IEEE, Sept. 2022, pp. 1–7.
- [11] S. Manjanna, N. Kakodkar, M. Meghjani, and G. Dudek, "Efficient Terrain Driven Coral Coverage Using Gaussian Processes for Mosaic Synthesis," in *2016 13th Conference on Computer and Robot Vision (CRV)*, June 2016, pp. 448–455.
- [12] D. Rao, A. Bender, S. B. Williams, and O. Pizarro, "Multimodal information-theoretic measures for autonomous exploration," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 4230–4237.
- [13] P. Singh, E. Gregson, J. Ross, M. Seto, C. Kaminski, and D. Hopkin, "Vision-Based AUV Docking to an Underway Dock using Convolutional Neural Networks," in *2020 IEEE/OES Autonomous Underwater Vehicles Symposium (AUV)*, Sept. 2020, pp. 1–6.
- [14] T. Matsuda, T. Maki, K. Masuda, and T. Sakamaki, "Resident autonomous underwater vehicle: Underwater system for prolonged and continuous monitoring based at a seafloor station," *Robotics and Autonomous Systems*, vol. 120, p. 103231, Oct. 2019.
- [15] A. M. Yazdani, K. Sammut, O. Yakimenko, and A. Lammas, "A survey of underwater docking guidance systems," *Robotics and Autonomous Systems*, vol. 124, p. 103382, Feb. 2020.
- [16] H.-M. Chou, Y.-C. Chou, and H.-H. Chen, "Development of a Monocular Vision Deep Learning-based AUV Diver-Following Control System," in *Global Oceans 2020: Singapore – U.S. Gulf Coast*, Oct. 2020.
- [17] N. Stilinović, D. Nad, and N. Mišković, "AUV for diver assistance and safety — Design and implementation," in *OCEANS 2015 - Genova*, May 2015, pp. 1–4.
- [18] A. Gomez Chavez, A. Ranieri, D. Chiarella, E. Zereik, A. Babić, and A. Birk, "CADDY Underwater Stereo-Vision Dataset for Human-Robot Interaction (HRI) in the Context of Diver Activities," *Journal of Marine Science and Engineering*, vol. 7, no. 1, p. 16, Jan. 2019.
- [19] J. R. Pelletier, B. W. O'Neill, J. J. Leonard, L. Freitag, and E. Galimore, "AUV-Assisted Diver Navigation," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 208–10 215, Oct. 2022.
- [20] A. Birk, "A Survey of Underwater Human-Robot Interaction (U-HRI)," *Current Robotics Reports*, vol. 3, no. 4, pp. 199–211, Dec. 2022.
- [21] Q.-Y. Zhou and V. Koltun, "Color map optimization for 3D reconstruction with consumer depth cameras," *ACM Transactions on Graphics*, vol. 33, no. 4, pp. 155:1–155:10, July 2014.
- [22] K. De and M. Pedersen, "Impact of Colour on Robustness of Deep Neural Networks," in *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*. Montreal, BC, Canada: IEEE, Oct. 2021, pp. 21–30.
- [23] N. Carlevaris-Bianco, A. Mohan, and R. M. Eustice, "Initial results in underwater single image dehazing," in *OCEANS 2010 MTS/IEEE SEATTLE*, Sept. 2010, pp. 1–8.
- [24] H.-Y. Yang, P.-Y. Chen, C.-C. Huang, Y.-Z. Zhuang, and Y.-H. Shiau, "Low Complexity Underwater Image Enhancement Based on Dark Channel Prior," in *2011 Second International Conference on Innovations in Bio-Inspired Computing and Applications*, Dec. 2011, pp. 17–20.
- [25] J. Y. Chiang and Ying-Ching Chen, "Underwater Image Enhancement by Wavelength Compensation and Dehazing," *IEEE Transactions on Image Processing*, vol. 21, no. 4, pp. 1756–1769, Apr. 2012.
- [26] D. Berman, T. Treibitz, and S. Avidan, "Diving into haze-lines: Color restoration of underwater images," in *Proceedings of the British Machine Vision Conference*. BMVA Press, 2017.
- [27] M. Bryson, M. Johnson-Roberson, O. Pizarro, and S. B. Williams, "True Color Correction of Autonomous Underwater Vehicle Imagery," *Journal of Field Robotics*, vol. 33, no. 6, pp. 853–874, 2016.
- [28] W. S. Pegau, D. Gray, and J. R. V. Zaneveld, "Absorption and attenuation of visible and near-infrared light in water: Dependence on temperature and salinity," *Applied Optics*, vol. 36, no. 24, Aug. 1997.
- [29] D. Akkaynak and T. Treibitz, "A Revised Underwater Image Formation Model," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Salt Lake City, UT: IEEE, June 2018, pp. 6723–6732.
- [30] D. Berman, D. Levy, S. Avidan, and T. Treibitz, "Underwater Single Image Color Restoration Using Haze-Lines and a New Quantitative Dataset," *arXiv:1811.01343 [cs]*, Mar. 2019.
- [31] H. Porav, W. Maddern, and P. Newman, "Adversarial Training for Adverse Conditions: Robust Metric Localisation Using Appearance Transfer," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 1011–1018.
- [32] J. Li, K. A. Skinner, R. M. Eustice, and M. Johnson-Roberson, "WaterGAN: Unsupervised Generative Network to Enable Real-time Color Correction of Monocular Underwater Images," *IEEE Robotics and Automation Letters*, 2017.
- [33] C. Fabbri, M. J. Islam, and J. Sattar, "Enhancing Underwater Imagery Using Generative Adversarial Networks," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 7159–7165.
- [34] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, Dec. 2015.
- [35] A. Reuther, P. Michaleas, M. Jones, V. Gadepally, S. Samsi, and J. Kepner, "Survey of Machine Learning Accelerators," in *2020 IEEE High Performance Extreme Computing Conference (HPEC)*, Sept. 2020, pp. 1–12.
- [36] K. He, J. Sun, and X. Tang, "Single Image Haze Removal Using Dark Channel Prior," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 12, pp. 2341–2353, Dec. 2011.
- [37] G. Buchsbaum, "A spatial processor model for object colour perception," *Journal of the Franklin Institute*, vol. 310, no. 1, July 1980.
- [38] M. Ebner and J. Hansen, "Depth map color constancy," *Bio-Algorithms and Med-Systems*, vol. 9, no. 4, Jan. 2013.
- [39] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "PyTorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035.
- [40] Y. Girdhar, N. McGuire, L. Cai, S. Jamieson, S. McCammon, B. Claus, J. E. San Soucie, J. E. Todd, and T. A. Mooney, "CUREE: A Curious Underwater Robot for Ecosystem Exploration," in *2023 International Conference on Robotics and Automation (ICRA)*, London, UK, May 2023.