

Harmful Cyanobacterial Bloom Detection using Deep Learning and Sentinel-2 Imagery

Conrad Schloer^a and Ronny Hänsch^b

^aTechnische Universität Berlin, Berlin, Germany

^bGerman Aerospace Center (DLR), Oberpfaffenhofen, Germany

Abstract—Harmful cyanobacterial blooms (CyanoHABs) pose a growing threat to human health and aquatic ecosystems, making early detection and monitoring critical. Current satellite-based approaches, such as the Cyanobacteria Assessment Network (CyAN) using the Sentinel-3 OLCI sensor, are optimized for chlorophyll *a* but are limited by a coarse 300 m pixel size. While Sentinel-2 lacks the spectral specialization of Sentinel-3, it provides finer sampling (10–60 m), enabling monitoring of smaller lakes and rivers if suitable models are applied. Here, we present a method to generate a large, geographically diverse CyanoHAB dataset by aligning CyAN data with Sentinel-2 imagery, resulting in 938,607 images and 1.86 billion labeled pixels across the continental United States. Using this dataset, we train a deep learning model that predicts CyanoHAB risk levels with 86.2% accuracy. Bands 4 (665 nm) and 5 (705 nm) are most influential, but randomizing any of 11 bands reduces accuracy by at least 10% in one or more classes, demonstrating the importance of the full spectral range. Both the dataset and code are publicly available (<https://zenodo.org/records/14230064>; https://github.com/cschloer/hab_detection_s2).

Index Terms—Machine Learning, Cyanobacteria, Harmful Algae Blooms, Multispectral Imagery, Satellite Imagery.

I. INTRODUCTION

CYANOBACTERIA are a group of photosynthetic bacteria commonly found in aquatic and terrestrial environments. They are commonly called blue-green algae but are distinct from true algae, which are eukaryotes [10], [30]. A cyanobacterial bloom is a phenomenon in which cyanobacteria multiply rapidly in response to favorable environmental stimulus, such as excess nutrients or warmer temperatures [33]. Blooms are naturally positively buoyant and can grow large enough to be seen from space, as illustrated in Figure 1.

Cyanobacterial blooms threaten human and ecosystem health through cyanotoxin production (a class of toxins produced by 25–75% of cyanobacterial blooms [2]) and by depleting dissolved oxygen, leading to the formation of dead zones [7]. The economic cost of harmful algal blooms (HABs) in the United States is conservatively estimated at an average of \$50 million per year due to impacts on public health, commercial fisheries, and tourism [32]. The term HAB refers more broadly to blooms caused by both cyanobacterial and non-cyanobacterial organisms.

Various methods are employed for detecting and predicting HABs. In situ approaches include measuring water clarity



Fig. 1. A cyanobacterial bloom seen from space. Natural color image from Landsat 8 of Lake Erie [23] (Credit: USGS EROS).

with a Secchi disk [1], quantifying chlorophyll *a* concentrations in water samples [1], and genome sequencing [9]. Although accurate, these approaches are labor-intensive and lack scalability for regional or national monitoring. In contrast, remote sensing using buoys, aircraft, or satellites to measure the spectral signature of cyanobacteria provides an efficient means of obtaining coarse but large-scale estimates of HAB occurrence [1], [11], [12], [17], [34].

A. Cyanobacteria Assessment Network (CyAN)

In 2015, a collaboration between the U.S. Environmental Protection Agency (EPA), the National Aeronautics and Space Administration (NASA), the National Oceanic and Atmospheric Administration (NOAA), and the U.S. Geological Survey (USGS) resulted in the development of the Cyanobacteria Assessment Network (CyAN), a software tool that uses satellite imagery to detect and monitor cyanobacterial blooms [28]. The U.S. Army Corps of Engineers (USACE) joined the collaboration in 2023 [6]. The project includes a web and mobile application for visualizing data and a data repository for downloading CyanoHAB observations.

Historical data are derived from the ESA MERIS sensor (2002–2012). For more recent data, the OLCI sensors onboard Sentinel-3A (2016–present) and Sentinel-3B (2018–present) are used to generate a cyanobacteria concentration digital number estimate (CyanoHAB value) ranging from 0 (no CyanoHAB) to 255 (high CyanoHAB) with a daily temporal resolution.

0000–0000/00\$00.00 © 2025 IEEE

The algorithm used to generate the CyanoHAB values is based on the method proposed by Wynne et al. [35], which utilizes three MERIS bands in the red to near-infrared region of the electromagnetic spectrum, together with meteorological variables such as air temperature and wind patterns, to detect CyanoHABs [25], [29].

The 300 m pixel size of Sentinel-3 limits the effective resolution of CyAN predictions. Small water bodies are often represented by mixed pixels containing both water and land, which reduces accuracy.

Sentinel-2 provides finer spatial sampling (10–60 m depending on band), enabling more detailed mapping of inland water bodies. However, the spectral configuration of Sentinel-2 differs from that of Sentinel-3. In particular, Sentinel-2 has fewer bands in the red and red-edge regions, reducing its direct suitability for use with the Wynne et al. [35] algorithm employed in CyAN.

B. Improving the Spatial Resolution of CyAN

This study investigates the use of machine learning, specifically convolutional neural networks (CNNs), to predict cyanobacterial concentration from Sentinel-2 satellite imagery. Because Sentinel-2 offers finer sampling than Sentinel-3, the task can be formulated as super-resolving low-resolution reference data. Specifically, the reference data (derived from Sentinel-3) have a pixel size of 300 m, whereas the Sentinel-2 input data range from 10 m to 60 m.

To improve the spatial resolution of CyAN predictions, the model is trained using Sentinel-2 spectral data. All Sentinel-2 bands are resampled to a common resolution of 20 m. This resolution is selected because Bands 5 and 6 (705 nm and 740 nm), which have been shown to be particularly relevant for chlorophyll detection, are natively sampled at 20 m. Furthermore, downsampling reduces noise, whereas upsampling may amplify noise; therefore, minimizing the number of upsampled bands is advantageous. Unifying the data at 20 m requires upsampling only three bands (1, 9, and 10), downsampling five bands (2, 3, 4, 8), and retaining six bands (5, 6, 7, 8a, 11, 12) at their native resolution.

The CyAN reference data are interpolated to 20 m to match the Sentinel-2 input resolution. Discrete CyanoHAB values (0–253) are aggregated into categorical classes for training. The model is trained using the paired input and reference data. Incorporating spatial context, which is not considered in the original CyAN algorithm, should further improve predictive performance.

This large-scale, automatically generated dataset for CyanoHAB detection substantially increases both geographic and temporal diversity relative to existing datasets. By aligning CyAN risk assessments with temporally and spatially corresponding Sentinel-2 imagery, a comprehensive training resource is created for high-resolution modeling. The dataset comprises 938,607 images with 1.86 billion labeled pixels across the continental United States. Both the dataset and the code used to generate it are publicly available, supporting reproducibility and enabling further research in this domain. Together, the automated data generation approach and the

dataset itself represent a second key contribution of this work, complementing the development of the deep learning model.

II. RELATED WORK

Numerous recent studies have explored the use of satellite imagery for harmful algal bloom (HAB) detection, with a growing emphasis on deep learning methods.

Traditional machine learning approaches include Joshi et al. [15], who applied a Random Forest classifier to predict four HAB indicators, including chlorophyll *a*, from Sentinel-3 imagery and ranked spectral bands by importance. Caballero et al. [3] developed reflectance algorithms using in situ measurements from the Iberian Peninsula and compared the spectral capabilities of Sentinel-2, Sentinel-3, and Landsat-8, concluding that the finer sampling of Sentinel-2 makes it more suitable for inland HAB monitoring.

Deep learning methods have been investigated extensively. Hill et al. [14] combined CNNs and long short-term memory (LSTM) networks to detect HABs in Florida using ground-truth data from the Florida Fish and Wildlife Conservation Commission (2003–2018). The data were binarized into HAB/no-HAB classes, converted into spatiotemporal “data cubes,” aligned with MODIS-Aqua and Terra imagery, and used to train the model, which achieved 91% accuracy. Yan et al. [36] developed a U-Net model for CyanoHAB detection in Chaohu Lake, China, using training labels derived from a floating algae index and expert visual interpretation, outperforming three alternative extraction methods. Zhao et al. [38] introduced a U-Net with improved channel attention and a modified loss function for detecting HABs in the East China Sea using HY-1D CZI imagery (50 m resolution), outperforming five other approaches despite the sensor’s limited spectral coverage. Cui et al. [5] proposed a deep learning-based super-resolution method for detecting ocean green tides in the Yellow Sea, using Gaofen-1 (16 m) and MODIS (250 m) imagery and achieving an average accuracy of 93.13% across 10 test images.

Alternative labeling and clustering strategies have also been explored. Medina-López et al. [22] used unsupervised clustering of satellite imagery combined with human expert labeling to monitor macroalgal blooms in Spain’s Mar Menor lagoon, achieving >98% accuracy with a Classification and Regression Tree (CART) model.

Finally, Niroumand-Jadidi and Bovolo [26] proposed a neural network to generate a virtual orange band (~620 nm) for Sentinel-2 and Landsat imagery. Because phycocyanin exhibits a strong absorption near 620 nm, this virtual band can improve the differentiation of cyanobacteria from other phytoplankton.

Building on these efforts, the present study employs Sentinel-2 imagery and deep learning for CyanoHAB detection. Unlike prior work that relies on limited in situ measurements or manual labeling, we automatically generate a large and geographically diverse dataset using the CyAN algorithm, resulting in 938,607 images with 1.86 billion labeled pixels across the continental United States. Its size and diversity enable spatially independent training and test sets, mitigating




Fig. 2. Workflow for dataset creation and model training. Relevant CyAN and Sentinel-2 images are downloaded, divided into matching subtiles, filtered based on HAB occurrence, and masked for clouds and land. The resulting dataset is used to train and evaluate a semantic segmentation model that produces high-resolution HAB estimates.

data leakage, and allow all Sentinel-2 bands to be used as input, facilitating a detailed analysis of spectral contributions to CyanoHAB detection.

III. METHODS

This section details the methodology employed in this study. Each component of the workflow, illustrated in Figure 2, is described in the following subsections, including dataset construction, model architecture and training configuration, and the loss function design.

A. Dataset

The dataset is created using Sentinel-2 multispectral images (12 bands) as features and v4 CyAN images as reference. The CyAN HAB detector has covered the entire continental USA and Alaska daily since 2016 [25]. Sentinel-2 revisits every location on the planet every 3–5 days [8]. Given the vast temporal and spatial scale of available data, it is necessary to limit the number of scenes in the dataset. Furthermore, the majority of lakes forecasted by CyAN are relatively free of CyanoHAB. A naïve dataset would thus have significant class imbalance, potentially leading to problems during training. Therefore, the dataset is limited to scenes that contain interesting features, namely scenes that contain CyanoHAB above a given threshold.

CyAN divides the USA into tiles, as illustrated in Figure 3. An automated process queries every CyAN image available for a given time period and CyAN tile. Each CyAN tile is downloaded separately and divided into 1600 subtiles, each 50×50 pixels large, corresponding to a 15×15 km area. Each subtile is evaluated for inclusion in the dataset separately. First, it is confirmed that the subtile contains enough water pixels to be of interest. CyAN images have a built-in land mask, i.e., land pixels are labeled with 254. If more than 90% of a subtile's pixels are land, the subtile is excluded from the dataset.

Next, the number of CyanoHAB occurrences in the subtile is counted. A CyanoHAB occurrence is defined as any




Fig. 3. The CyAN tile designation in the continental USA [24].




Fig. 4. Example of subtile selection in CyAN region 6_2. Black areas are cloud-covered and not visualized. Areas without any squares do not contain enough water to be considered. Areas with black squares do not contain enough CyanoHAB to be included. Red squares represent tiles selected for training, blue squares are tiles selected for testing. The reference image is the CyAN image from the sampling period with the least cloud cover and does not necessarily represent the mean CyanoHAB values for each subtile.

CyanoHAB level between 1 and 253. Values of 0 (no CyanoHAB) and 255 (clouds) are both considered as no CyanoHAB. The percentage of CyanoHAB in all water pixels (ignoring land pixels) is then calculated across the entire time period. If more than 10% of the water pixels in a subtile contain CyanoHAB, the subtile is added to the dataset. The time period used is from July 1, 2019, to August 29, 2019, and the chosen CyAN tiles are 6_2, 6_5, 7_2, 7_5, 8_3 (see Figure 3). The time period is chosen because CyanoHAB levels are highest in the summer season. The CyAN tiles are chosen because they contain lakes identified as CyanoHAB areas of interest by NOAA domain experts during the project planning phase. See Figure 4 for a visualization of the subtile selection process.

This method results in 183 subtiles selected from the 8000

available in the 5 CyAN tiles. These chosen subtiles are then divided into train and test sets. To prevent data leakage between sets, where adjacent areas could appear in both train and test sets, an iterative region-growing algorithm is used. First, a subtile is randomly assigned to test or train based on the desired ratio. Then, all adjacent chosen subtiles are recursively assigned to the same set, ending once there are no more adjacent subtiles available. This process is repeated until all selected subtiles are assigned. The ratio used is 70% train and 30% test, though the stochastic nature of the method leads to a slightly different final ratio.

The next step finds the corresponding Sentinel-2 products. Data from January 1, 2019, to December 31, 2020, is used. Data after 2022 is excluded due to changes in European Space Agency (ESA) processing. The dataset is limited to two years because this yields sufficient data for the scope of this paper. The cloud cover percentage is restricted to 0-10% to minimize downloading data with too much occlusion. Applying these filters results in a total of 1,789 Sentinel-2 products across all subtiles and dates, with some products repeated because they completely cover multiple subtiles.

The Sentinel-2 Level-2A products are downloaded and converted into single GeoTIFF files with all bands unified at 20 m resolution. Bands not already at 20 m resolution (e.g., Band 9 at 60 m or Band 2 at 10 m) are resampled to 20 m using simple nearest-neighbor interpolation. Next, the corresponding CyAN tile is downloaded if not already available from the previous subtile selection step. Since CyAN HAB images are 300 m resolution, they are upsampled to 20 m using bilinear interpolation. The bilinear method was chosen by NOAA domain experts after visually comparing interpolation results from multiple methods.

Once both files are at 20 m resolution, they are cropped to fit the bounding box of the relevant subtile. The CyAN image already contains a mask for cloud and land cover. However, the Sentinel-2 image is taken at a different time of day, meaning cloud cover may differ. Furthermore, the land mask is designed for 300 m resolution. Upsampling to 20 m may cause some pixels to be misclassified as land or water. The former is not an issue, as land pixels are ignored. The latter is problematic, as pixels incorrectly labeled as water may contribute noise during training and testing.

To address the cloud mask problem, a simple cloud detection algorithm proposed by Zhai et al. [37] is used. It first computes two cloud indices CI_1 and CI_2 :

$$CI_1 = \frac{B_2 + 2 \cdot B_{11}}{B_2 + B_3 + B_4} \quad (1)$$

$$CI_2 = \frac{B_2 + B_3 + B_4 + B_8 + B_{11} + B_{12}}{6}, \quad (2)$$

where B_i refers to band index i . CI_1 and CI_2 are then used to determine if a pixel is a cloud based on the thresholds T_1 and T_2 :

$$(|CI_1 - 1| < T_1) \text{ or } (CI_2 > T_2), \quad (3)$$

where values $T_1 = 1$ and $T_2 = 0.1$ are used as recommended by Zhai et al. Equation 3 is applied to all pixels in the Sentinel-2 image, and the resulting cloud-labeled pixels are




Fig. 5. Demonstration of the cloud and land filter. Top left: original Sentinel-2 image; top right: corresponding CyAN HAB product. Bottom left: cloud and land filter, with clouds and land in black. Bottom right: CyAN HAB product with cloud and land filter overlay, with black pixels corresponding to “no data” (clouds or land). The clouds are clearly shown as black spots on the lake, and the land pixels are clear when comparing the shoreline with that of the original CyAN HAB product.

added to the CyAN mask. Notably, cloud shadows are not detected and therefore remain unmasked.

More sophisticated cloud detection algorithms exist, such as Sen2Cor [21], Fmask [27], and MAJA [13], which use multi-temporal analysis or advanced machine learning methods to improve cloud and shadow detection. However, these approaches are significantly more computationally intensive and can introduce delays when processing large-scale datasets such as ours. Given that the focus of our work is on spectral and spatial feature contributions to CyanoHAB detection rather than optimizing pre-processing pipelines, the method of Zhai et al. is chosen for its computational efficiency, reproducibility, and adequate performance for reducing noise in this context.

To mitigate the land mask problem, a simple land detection algorithm is used:

$$(B_{NIR} > B_G) \text{ and } (B_{SWIR} > B_G) \quad (4)$$

The spectral signature of water absorbs NIR light and reflects green light; therefore, Band 8A (B_{NIR}) is compared with Band 3 (B_G , green). To avoid false positives caused by dense vegetation in water, Band 11 (B_{SWIR}) is also compared with Band 3, since middle-infrared light is absorbed strongly by water but not by vegetation [16]. This equation is applied to all pixels in the Sentinel-2 image, and the resulting land-labeled pixels are incorporated into the CyAN mask.

This spectral thresholding method is chosen for its simplicity, efficiency, and robustness across diverse regions. More advanced approaches (e.g., NDWI-based segmentation or machine learning classifiers) require regional tuning or extra training data, reducing scalability.

Since misclassified cloud and land pixels are rare, these lightweight methods sufficiently minimize noise during training. See Figure 5 for a demonstration of the cloud and land filter.

The last step divides each subtitle into 64×64 patches. This is done by iterating through the image from the top left to the bottom right with a stride of 64 (no overlap), cropping the Sentinel-2 and CyAN images into geographically identical 64×64 images. Patches containing more than 5% “no data” pixels are discarded.

B. Model and Training Setup

The semantic segmentation model is based on a DeepLabv3 architecture with a ResNet-18 backbone. To accommodate the multispectral input data, the first convolutional layer of the encoder is modified to accept 12 input channels instead of the standard three. Input data are normalized per channel and model weights are initialized randomly.

Given the highly imbalanced distribution of CyanoHAB concentration values, the continuous range $[0, 253]$ is discretized into three risk levels:

- Class 1 (low risk): values 0-99,
- Class 2 (moderate risk): values 100-199,
- Class 3 (high risk): values 200-253.

This binning both reflects domain-relevant risk categories and mitigates the problem of underrepresented high-concentration values in the dataset. The network’s final classification layer is therefore configured to output three channels corresponding to the three classes.

During each training epoch, images are randomly flipped horizontally and vertically with a probability of 50%. Random cropping and blurring were omitted, as preliminary experiments indicated that they had minimal impact on overfitting for this model and dataset.

The model is trained using the Adam optimizer [18] with cross-entropy loss (see Section III-C). Hyperparameters, including the model architecture, learning rate, batch size, and weight decay, are selected via three-fold cross-validation on a subset of 10,000 samples. The final selected values are a learning rate of 0.0001, a batch size of 128, and a weight decay of 0.01.

C. Loss Function

The loss function used is the Cross Entropy Loss with additional custom pixel distribution weighting.

Pixel distribution weighting is necessary because of the CyanoHAB value imbalance. For example, CyanoHAB level 225 is about 100,000 times less likely to occur in the training set than CyanoHAB level 0 (see Figure 6). Therefore, the contribution of each pixel to the loss must be weighted according to the distribution of its label value. Furthermore, this weighting cannot be applied only at the class level, as CyanoHAB values near class boundaries have similar distributions but might otherwise be assigned very different weights based on the overall class frequency. This could bias the model toward misclassifying edge pixels into the neighboring class with a lower distribution, since the penalty for misclassifying the rare class would be disproportionately higher than that for misclassifying the common class, even though their edge distributions are similar. Pixel distribution weighting is performed via:




Fig. 6. Pixel count fraction by CyanoHAB level for the full dataset as well as the train and test subsets. Fractions are relative to the size of each set. Note that the y-axis uses a logarithmic scale.

$$l_i^{new} = l_i^{old} \cdot \frac{D(0)}{D(v_i)}, \quad (5)$$

where l_i is the loss for a specific pixel i , $D(\cdot)$ is a function that returns the distribution frequency for a specific CyanoHAB value (see Figure 6), and v_i is the actual CyanoHAB value for pixel i . The numerator uses the distribution of CyanoHAB value 0 (the most common value) to normalize all other weights.

IV. RESULTS

A. Dataset

The dataset generation process yields a total of 938,607 image pairs. Across these images, approximately 1.86 billion usable pixels (excluding cloud- and land-masked areas) are available. These pixels are distributed across the three risk classes at ratios of 89.51%, 10.34%, and 0.15% for low (0-99), medium (100-199), and high (200-253) risk levels, respectively. The proposed train-test split procedure, which enforces geographic disjointness, produces an approximately 80/20 split. The training set contains 776,864 image pairs, with class distributions of 89.90%, 10.04%, and 0.06% for the low, medium, and high risk classes, respectively. The test set contains 161,743 image pairs with corresponding distributions of 87.34%, 12.00%, and 0.67%. See Figure 6 for a visualization of the pixel count fraction by CyanoHAB level for each subset and for the entire dataset.

Figure 7 presents the mean spectral signature of each risk class across all pixels in the test set.

B. Model Performance

Figure 8 shows a visualization of the loss throughout training.

The model achieved 88.29% accuracy on low-risk, 82.76% accuracy on moderate-risk, and 87.86% accuracy on high-risk CyanoHAB. The model very rarely misclassifies by more than one class: when the true label is low-risk, it predicts high-risk




Fig. 7. Mean spectral signature of each risk class across all pixels in the test set. Vertical lines indicate the Sentinel-2 spectral bands. The plotted reflectance values correspond to the post-processed inputs used for model training.




Fig. 8. Training and test loss over 50 epochs. Semi-transparent lines represent raw loss values, while opaque lines show a smoothed version.

TABLE I
METRICS FOR EACH CYANOHAB RANGE ON THE TEST DATASET.

CyanoHAB Range	0-99	100-199	200-253	Average
Accuracy	0.8829	0.8276	0.8776	0.8627
Precision	0.9811	0.4737	0.4613	0.6387
Recall	0.8829	0.8276	0.8776	0.8627
Specificity	0.8721	0.8830	0.9954	0.9168
F1-Score	0.9294	0.6026	0.6047	0.7122

only 0.02% of the time, and when the true label is high-risk, it predicts low-risk only 0.09% of the time. See Figure 9 for the confusion matrix. For a detailed view of performance at each individual CyanoHAB level, see Figure 10.

Across all classes on the test set, the average precision is 64%, the average recall is 86%, the average specificity is 92%, and the average F1-score is 71.22% (see Table I).

Figure 11 shows a qualitative example of the model's performance on a selected image.




Fig. 9. Confusion matrix showing the fraction of predictions for each class. The y-axis represents the predicted label, and the x-axis represents the actual label. Percentages are normalized along the x-axis.




Fig. 10. Per-class performance across the CyanoHAB index on the test set. Thick opaque lines represent correct classifications, while transparent lines represent misclassifications. The sharp drops near the right end of the x-axis correspond to CyanoHAB levels not represented in the training set.




Fig. 11. Lake Winnebago results. Top left: actual CyanoHAB values. Top right: actual values converted into three classes. Bottom right: model predictions. Bottom left: difference map between ground truth and predictions.

V. DISCUSSION

A. Model Performance

The model's performance meets our expectations. We initially set a benchmark goal of 80% accuracy, which the model surpasses with an average accuracy of 86% (see Table I). Recall and specificity are also strong, with averages of 86% and 92%, respectively. Notably, the specificity for CyanoHAB levels 200-253 (high risk) is 99.54%, which is particularly important for CyanoHAB monitoring. This means that when the model predicts a pixel is not high risk, we can be highly confident that this classification is correct.

The weakest metric is precision. While CyanoHAB 0-99 (low risk) achieves a precision of 98.1%, the precision for moderate-risk (100-199) and high-risk (200-253) classes is much lower, at 47.4% and 46.1%, respectively. This result can be attributed to the extreme class imbalance in the dataset. As shown in the middle-right of Figure 9, pixels that are actually moderate risk are predicted as high risk only 3.9% of the time. However, because there are so many moderate-risk pixels, this small percentage corresponds to 900,177 pixels - making up 52.16% of all pixels predicted as high risk. A similar pattern occurs with low-risk pixels misclassified as moderate risk: only 12% of low-risk pixels are predicted as moderate risk, but this represents 20,863,560 pixels, which exceeds the number of correctly predicted moderate-risk pixels (18,879,988). In fact, they account for 52.4% of all pixels predicted as moderate risk.

Thus, precision is the poorest-performing metric for moderate- and high-risk CyanoHAB. However, this outcome is acceptable for the intended application of this model. Ideally, all metrics would perform well, but if one metric must underperform, low precision at higher risk levels is the most tolerable. This model is designed to serve as a coarse, first-layer "alarm" for CyanoHAB outbreaks. If the model flags a potential outbreak, further analysis - such as in situ testing - would typically follow before any management action is taken.

B. Band Contribution

Figure 7 shows that there is a clear difference in spectral signature between the three classes. To measure the impact of each spectral band, we tested model performance across the entire test set while a given band was randomized. Because the data are normalized to have a mean of 0 and a standard deviation of 1, random values could be drawn from a standard normal distribution without affecting the overall variance of the input to the model. We also repeated the test by setting each band to zero rather than randomizing it, and observed similar performance, so we report only the randomization results. Figure 12 shows the result of this experiment.

Previous studies on remote sensing of CyanoHAB, including research that informed the development of the CyAN algorithm, have shown that the red portion of the electromagnetic spectrum is most relevant for detecting cyanobacteria. Specifically, a spectral shape algorithm centered on the 681 nm band from the European Space Agency's Medium Resolution Imaging Spectrometer (MERIS) and a similar band from the Sentinel-3 satellite was used to create the CyanoHAB values




Fig. 12. Model performance (from top to bottom: overall, as well as for low-, moderate-, and high-risk class, respectively) after randomizing a single band. Random values are sampled from a normal distribution with mean 0 and standard deviation 1. Accuracy is reported as the mean per-class accuracy.

TABLE II

CHANGE IN ACCURACY FOR EACH CLASS WHEN RANDOMIZING EACH BAND. VALUES OVER ABSOLUTE VALUE 10% ARE BOLDED.

CyanoHAB Range	0-99	100-199	200-253	Average
Band 1	+0.0454	-0.1253	-0.045	-0.0416
Band 2	+0.0186	-0.1535	+0.0365	-0.0328
Band 3	-0.0748	-0.0410	-0.0856	-0.0671
Band 4	-0.1105	-0.5654	-0.1251	-0.2670
Band 5	-0.5728	-0.0921	-0.8616	-0.5088
Band 6	+0.0027	-0.3176	-0.0604	-0.1251
Band 7	+0.0505	-0.4435	+0.0176	-0.1251
Band 8	-0.0372	-0.1455	-0.1840	-0.1222
Band 8a	-0.3469	-0.064	-0.0108	-0.1406
Band 10	+0.0399	-0.1132	+0.0079	-0.0218
Band 11	+0.0706	-0.2833	-0.0015	-0.0714
Band 12	+0.0493	-0.1961	+0.0045	-0.0474

that serve as reference data for our model. That spectral shape algorithm uses the 665 nm, 681 nm, and 709 nm bands [4], [31], [34]. Although Sentinel-2 data do not have as high a spectral resolution within the red region, Band 4 (665 nm) and Band 5 (705 nm) are closest to the bands used in the CyAN algorithm. Therefore, the poor model performance observed when these two bands are removed, as shown in Figure 12, is expected.

Figure 12 and Table II show how the results vary by class. Interestingly, the contribution of individual bands differs significantly between classes. Both low-risk and high-risk classes suffer significant performance drops when Band 5 is removed, whereas moderate risk is comparatively unaffected. Conversely, moderate-risk performance depends strongly on Band 4. High risk also relies on Band 8 (842 nm), while low risk relies more on Band 8A (865 nm). Moderate risk, which has the lowest accuracy among the three classes, appears to depend on a wider range of bands, as its performance decreases by more than 0.1 when Bands 1 (443 nm), 2 (490 nm), 6 (740 nm), 7 (783 nm), 8, 9 (940 nm), 11 (1610 nm), and 12 (2190 nm) are removed. Interestingly, removing almost the same set of bands (1, 2, 6, 7, 9, 11, 12) slightly improves the performance of the low-risk class, suggesting that these bands provide a weak but informative signal for distinguishing the boundary between the two classes.

Overall, these results indicate that the spectral information contained in Sentinel-2 data is sufficient for detecting cyanobacteria. Although MERIS and Sentinel-3 data have higher spectral resolution in the red region, the model appears to compensate by leveraging a broader portion of the spectrum, particularly for detecting moderate-risk CyanoHAB. These findings suggest that hyperspectral satellite imagery with additional spectral bands - such as that provided by NASA's PACE mission - may further improve the accuracy of CyanoHAB detection when combined with machine learning approaches [16].

C. Spatial Contribution

To determine the impact of surrounding pixels on a given pixel's classification, we conducted two experiments in which we replaced all other pixels in the image with synthetic data.

TABLE III

ACCURACY RESULTS FOR THE EXPERIMENT IN WHICH ALL PIXELS IN A 64×64 IMAGE WERE SET TO THE VALUE OF THE TARGET PIXEL.

CyanoHAB Range	0-99	100-199	200-253	Average
Original	0.8829	0.8276	0.8776	0.8627
Experiment	0.8820	0.7697	0.7363	0.7960
Δ	-0.0009	-0.0579	-0.1413	-0.0667

TABLE IV

ACCURACY RESULTS FOR THE EXPERIMENT IN WHICH ALL PIXELS OTHER THAN THE TARGET PIXEL WERE REPLACED WITH RANDOM VALUES SAMPLED FROM A NORMAL DISTRIBUTION WITH MEAN 0 AND STANDARD DEVIATION 1.

CyanoHAB Range	0-99	100-199	200-253	Average
Original	0.8829	0.8276	0.8776	0.8627
Experiment	0.9985	0.0015	0.0000	0.3333
Δ	0.1156	-0.8261	-0.8776	-0.5294

In the first experiment, for each pixel in the test dataset, we set all other pixels in the 64×64 image to the same value as the target pixel. The performance was then measured on a per-pixel basis, and the results are shown in Table III.

Model performance is notably reduced for both the moderate- and high-risk classes. In contrast, the low-risk class shows minimal change, indicating that its predictions are less dependent on spatial context.

In the second experiment, we replaced all surrounding pixels with random values sampled from a normal distribution with mean 0 and standard deviation 1. As described in Section V-B, the normalized preprocessing step ensures that this substitution does not substantially alter the overall input distribution. The results are shown in Table IV and indicate that the model's performance is highly dependent on spatial context for the moderate- and high-risk classes. When surrounding pixels are replaced with random noise, accuracy for these classes drops dramatically, approaching zero. In contrast, the accuracy for the low-risk class increases, suggesting that the model may occasionally misclassify low-risk pixels as higher risk due to contextual information from neighboring pixels. Overall, this experiment highlights that spatial context plays a critical role in the detection of CyanoHAB hotspots, particularly for the more severe classes.

D. Comparison with normalized difference indices

To better understand the benefit of using a deep learning model that incorporates all available spectral bands as input, we compare its performance to that of five different normalized difference indices (NDIs) commonly used to detect CyanoHAB. The Floating Algae Index (FAI), described by Liu et al. [20], is a spectral formula that can be applied to Sentinel-2 or LANDSAT products. The Normalized Difference Vegetation Index (NDVI), Normalized Difference Chlorophyll Index (NDCI), the Band 8A–Band 4 (B8AB4) index, and the Band 3–Band 2 (B3B2) index are existing NDIs that Kislik et al. [19] applied to Sentinel-2:

TABLE V
AVERAGE METRICS ACROSS 3 GROUPS FOR THE NDIs AND THE PAPER'S CNN MODEL ON THE TEST DATASET.

Model	Accuracy	Precision	Recall	Specificity	F1-Score
FAI	0.5759	0.3706	0.5759	0.7265	0.2787
NDVI	0.2845	0.3030	0.2845	0.6239	0.2934
NDCI	0.2772	0.2879	0.2773	0.6143	0.2825
B8AB4	0.2899	0.3053	0.2899	0.6306	0.2970
B3B2	0.3181	0.3011	0.3181	0.6547	0.3057
CNN	0.8627	0.6387	0.8627	0.9168	0.7122

$$FAI = B_{8a} - \left(B_{04} + (B_{11} - B_{04}) \cdot \frac{\lambda_{8a} - \lambda_4}{\lambda_{11} - \lambda_4} \right) \quad (6)$$

$$NDVI = \frac{B_{08} - B_{04}}{B_{08} + B_{04}} \quad (7)$$

$$NDCI = \frac{B_{05} - B_{04}}{B_{05} + B_{04}} \quad (8)$$

$$B8AB4 = \frac{B_{8a} - B_{04}}{B_{8a} + B_{04}} \quad (9)$$

$$B3B2 = \frac{B_{03} - B_{02}}{B_{03} + B_{02}}, \quad (10)$$

where B_{03} , B_{04} , B_{05} , B_{08} , B_{8a} , and B_{11} are the reflectance values of bands 3, 4, 5, 8, 8a, and 11 respectively. The λ values are constants that refer to the central wavelength of the respective bands.

In order to compare the NDI output to the output of our model, we converted each NDI's continuous values into a discrete, three-class representation consistent with our model. We used a Bayesian optimization algorithm to optimize the thresholds for each NDI on the training dataset before applying them to the test dataset. We chose to optimize the macro F1-score, defined as the unweighted mean of the F1-scores across all classes, because it is well suited to imbalanced datasets and jointly accounts for precision and recall. The resulting performance metrics on the test dataset using the optimized thresholds are shown in Table V. The results clearly demonstrate that the CNN model, which utilizes all spectral bands, substantially outperforms all NDIs. This finding supports the conclusion in Section V-B that spectral information across the entire wavelength range contributes meaningfully to CyanoHAB detection performance.

E. Performance on Smaller Bodies of Water

One of the primary motivations of this paper is to improve the spatial resolution of CyanoHAB detection in order to enable more effective monitoring of smaller bodies of water. Quantitatively evaluating this improvement is challenging due to the scarcity of high-resolution reference data. As a result, direct comparisons must rely on synthetically interpolated data from CyAN, and evaluation is largely qualitative.

Figure 13 shows a section of the Allegheny Reservoir, a smaller body of water that was not used for training or testing. The 300 m resolution of the CyAN prediction is clearly limited in this narrow reservoir, as the coarse pixelation makes it difficult to determine which specific areas are experiencing CyanoHAB events. Our model's prediction (bottom right) is




Fig. 13. Images of the Allegheny Reservoir in Pennsylvania from August 20, 2019 comparing the reference data (CyAN) and the prediction. The bottom right image, despite a limited cloud and land filter, clearly demonstrates the advantage of using a higher spatial resolution model to monitor CyanoHAB in smaller bodies of water. Circled in blue is an area that our model predicts as high-risk CyanoHAB but which could not be identified by CyAN due to its coarse predictions. Note that this region was included in neither the training nor test datasets. The prediction was generated without tiling.

imperfect, primarily due to a relatively simple cloud and cloud-shadow filter: cloud edges are often misclassified as high-risk CyanoHAB, and cloud shadows are frequently classified as low or moderate risk instead of being masked as no data. Nonetheless, the advantage of our model's prediction is clear. The reservoir's shoreline is much more clearly delineated, and our model is able to produce a low-risk prediction for Quaker Lake (right side of the image), which is too small to be resolved in the CyAN prediction. Even more importantly, the model predicts high-risk CyanoHAB in the narrow channel connecting the reservoir and Quaker Lake (circled in blue) - a region that CyAN is unable to resolve. This represents actionable information for monitoring efforts that would be missed if relying solely on CyAN.

Figures 14 and 15 present additional case studies illustrating the benefits of higher-resolution CyanoHAB detection. In both cases, smaller water bodies and shoreline features are more clearly characterized with our model compared to the standard CyAN output. Figure 14 shows that our model successfully detects a high-risk HAB event localized to a narrow shoreline segment that CyAN fails to resolve. Figure 15 highlights how finer pixel size enables the detection of CyanoHAB in smaller tributaries that would otherwise go undetected when using CyAN.

VI. CONCLUSION

This study demonstrates that deep learning, and specifically convolutional neural networks (CNNs), can effectively predict harmful cyanobacterial blooms (CyanoHABs) using Sentinel-2 satellite imagery. Despite the lack of a chlorophyll *a*-specific spectral band available on Sentinel-3, the model successfully leverages the complex relationships among the available spectral bands, together with spatial context, to classify CyanoHAB risk levels (low, moderate, high) with an overall accuracy of 86.2%. Of particular importance for monitoring applications, the model achieves a specificity of 99.54% for the high-risk




Fig. 14. Images of Swan Lake in Michigan from July 2, 2019 comparing the reference data (CyAN) and the prediction.




Fig. 15. Images of Lake Jordan in North Carolina from July 28, 2019 comparing the reference data (CyAN) and the prediction.

class, indicating that pixels predicted as high-risk are almost always correctly classified.

Bands 4 (665 nm) and 5 (705 nm) are identified as the most critical for CyanoHAB detection. When Band 4 or Band 5 was randomized, overall accuracy dropped substantially to 59.5% and 35.4%, respectively. The effect is even more pronounced for the high-risk class, where randomizing Band 5 reduces accuracy to just 1.2% (from 87.8%). Nevertheless, 11 of the 12 Sentinel-2 bands cause more than a 10% decrease in accuracy for at least one class when randomized, underscoring the importance of the full spectral range for robust detection.

Spatial context also plays a crucial role. When surrounding pixels are replaced with entirely random values, model performance for moderate- and high-risk CyanoHAB dropped to near 0%, indicating that spatial information is indispensable for accurate classification. Together, these findings highlight the value of combining spectral and spatial features for operational CyanoHAB monitoring.

The use of interpolation to increase the spatial resolution of both the CyAN reference data and the Sentinel-2 input bands introduces important limitations in terms of accuracy and reliability. While interpolation allows for approximate

spatial alignment, it does not recover true high-resolution information, which may lead to over-optimistic performance estimates. Ideally, the model should be validated against in situ measurements in future work to confirm the accuracy of the predictions. Even if such validation yields lower quantitative performance, the finding that the full spectrum of Sentinel-2 bands contributes meaningfully to CyanoHAB detection remains an important result for guiding future research.

It is important to note that the model was trained and evaluated exclusively on summer data. Although CyanoHAB events are most prevalent during summer and early fall, they can also occur in other seasons. Future work should therefore explore the temporal adaptability of the model to ensure robust performance across different seasonal conditions.

Another promising avenue for improving model performance is the integration of a virtual orange band, similar to the approach proposed by Niroumand-Jadidi and Bovolo [26]. By explicitly providing a band known to be informative for CyanoHAB detection, the model may be relieved from the need to approximate this spectral feature from the existing bands. Whether this would lead to higher predictive performance or simply reduce training time remains an open question and warrants further investigation.

Overall, this study demonstrates that combining the full spectral range of Sentinel-2 with spatial context enables accurate detection of CyanoHABs at higher resolution than previous satellite-based approaches. The large, automatically generated CyanoHAB dataset - spanning 938,607 images and 1.86 billion labeled pixels across the continental U.S. - provides a geographically diverse resource that supports model development, enables spatially independent training and testing, and facilitates the analysis of all spectral bands. While interpolation and seasonal limitations remain, the results highlight the value of comprehensive spectral inputs and suggest avenues for improvement, including in situ validation, temporal extension, and the potential use of virtual bands. These contributions provide a strong foundation for scalable, high-resolution monitoring of harmful cyanobacterial blooms.

REFERENCES

- [1] Donald Anderson, Per Andersen, Vera Bricej, John Cullen, and J. Rensel. *Monitoring and Management Strategies for Harmful Algal Blooms in Coastal Waters*. Intergovernmental Oceanographic Commission Technical Series No. 59, 01 2001.
- [2] Luděk Bláha, Pavel Babica, and Blahoslav Maršálek. Toxins produced in cyanobacterial water blooms - toxicity and risks. *Interdisciplinary Toxicology*, 2(2), January 2009.
- [3] Isabel Caballero, Raúl Fernández, Oscar Moreno Escalante, Luz Mamán, and Gabriel Navarro. New capabilities of sentinel-2a/b satellites combined with in situ data for monitoring small harmful algal blooms in complex coastal waters. *Scientific Reports*, 10(1), May 2020.
- [4] John M. Clark, Blake A. Schaeffer, John A. Darling, Erin A. Urquhart, John M. Johnston, Amber R. Ignatius, Mark H. Myer, Keith A. Loftin, P. Jeremy Werdell, and Richard P. Stumpf. Satellite monitoring of cyanobacterial harmful algal bloom frequency in recreational waters and drinking water sources. *Ecological Indicators*, 80:84–95, 2017.
- [5] Bing Cui, Haoqing Zhang, Wei Jing, Huifang Liu, and Jianming Cui. Srse-net: Super-resolution-based semantic segmentation network for green tide extraction. *Remote Sensing*, 14(3):710, February 2022.
- [6] EPA. Cyanobacteria Assessment Network (CyAN). Retrieved from <https://www.epa.gov/water-research/cyanobacteria-assessment-network-cyan>, 2023.

- [7] EPA. The Effects: Dead Zones and Harmful Algal Blooms. Retrieved from <https://www.epa.gov/nutrientpollution/effects-dead-zones-and-harmful-algal-blooms>, 2024.
- [8] ESA. Sentinel-2 Mission Guide. Retrieved from <https://sentinel.copernicus.eu/web/sentinel/missions/sentinel-2>, 2023.
- [9] Sheena M. Feist and Richard F. Lance. Genetic detection of freshwater harmful algal blooms: A review focused on the use of environmental DNA (eDNA) in *Microcystis aeruginosa* and *Prymnesium parvum*. *Harmful Algae*, 110:102124, December 2021.
- [10] F. Garcia-Pichel. Cyanobacteria. In Moselio Schaechter, editor, *Encyclopedia of Microbiology (Third Edition)*, pages 107–124. Academic Press, Oxford, third edition edition, 2009.
- [11] Mohsen Ghanea, Masoud Moradi, and Keivan Kabiri. A novel method for characterizing harmful algal blooms in the persian gulf using modis measurements. *Advances in Space Research*, 58(7):1348–1361, 2016.
- [12] J. F. R. Gower, R. Doerffer, and G. A. Borstad. Interpretation of the 685nm peak in water-leaving radiance spectra in terms of fluorescence, absorption and scattering, and its observation by MERIS. *International Journal of Remote Sensing*, 20(9):1771–1786, January 1999.
- [13] Olivier Hagolle, Mireille Huc, Camille Desjardins, Stefan Auer, and Rudolf Richter. Maja algorithm theoretical basis document, December 2017.
- [14] Paul R. Hill, Anurag Kumar, Marouane Temimi, and David R. Bull. Habnet: Machine learning, remote sensing-based detection of harmful algal blooms. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 13:3229–3239, 2020.
- [15] Neha Joshi, Jongmin Park, Kaiguang Zhao, Alexis Londo, and Sami Khanal. Monitoring harmful algal blooms and water quality using sentinel-3 olci satellite imagery with machine learning. *Remote Sensing*, 16(13):2444, July 2024.
- [16] Hafiza Wajiha Khalid, Rao M. Zahid Khalil, and Muhammad Ateeq Qureshi. Evaluating spectral indices for water bodies extraction in western Tibetan Plateau. *The Egyptian Journal of Remote Sensing and Space Science*, 24(3):619–634, 2021.
- [17] Rabia Munsaf Khan, Bahram Salehi, Masoud Mahdianpari, Fariba Mohammadimanesh, Giorgos Mountrakis, and Lindi J. Quackenbush. A Meta-Analysis on Harmful Algal Bloom (HAB) Detection and Monitoring: A Remote Sensing Perspective. *Remote Sensing*, 13(21):4347, October 2021.
- [18] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. Retrieved from <https://doi.org/10.48550/arXiv.1412.6980>, 2017.
- [19] Chippie Kislik, Iryna Dronova, Theodore E. Grantham, and Maggi Kelly. Mapping algal bloom dynamics in small reservoirs using sentinel-2 imagery in google earth engine. *Ecological Indicators*, 140:109041, 2022.
- [20] Miao Liu, Hong Ling, Dan Wu, Xiaomei Su, and Zhigang Cao. Sentinel-2 and landsat-8 observations for harmful algae blooms in a small eutrophic lake. *Remote Sensing*, 13(21), 2021.
- [21] Jérôme Louis, Bringfried Pflug, Vincent Debaecker, Uwe Mueller-Wilm, Rosario Quirino Iannone, Valentina Boccia, and Ferran Gascon. Evolutions of sentinel-2 level-2a cloud masking algorithm sen2cor prototype first results. In *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, pages 3041–3044, 2021.
- [22] Encarni Medina-López, Gabriel Navarro, Juan Santos-Echeandía, Patricia Bernárdez, and Isabel Caballero. Machine learning for detection of macroalgal blooms in the mar menor coastal lagoon using sentinel-2. *Remote Sensing*, 15(5):1208, February 2023.
- [23] NASA. Satellites on toxic algae patrol. Retrieved from <https://landsat.gsfc.nasa.gov/article/satellites-on-toxic-algae-patrol/>, 2018.
- [24] NASA. CyAN File Search. Retrieved from https://oceansdata.sci.gsfc.nasa.gov/api/cyan_file_search/, 2023.
- [25] NASA. Cyanobacteria Assessment Network (CyAN). Retrieved from <https://oceanscolor.gsfc.nasa.gov/about/projects/cyan/>, 2023.
- [26] Milad Niroumand-Jadidi and Francesca Bovolo. Deep-learning-based retrieval of an orange band sensitive to cyanobacteria for landsat-8/9 and sentinel-2. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, PP:1–10, 01 2023.
- [27] Shi Qiu, Zhe Zhu, and Binbin He. Fmask 4.0: Improved cloud and cloud shadow detection in landsats 4–8 and sentinel-2 imagery. *Remote Sensing of Environment*, 231:111205, 2019.
- [28] Blake Schaeffer, Keith Loftin, Richard Stumpf, and P. Werdell. Agencies Collaborate, Develop a Cyanobacteria Assessment Network. *Eos*, 96, November 2015.
- [29] Blake A. Schaeffer, Sean W. Bailey, Robyn N. Conmy, Michael Galvin, Amber R. Ignatius, John M. Johnston, Darryl J. Keith, Ross S. Lunetta, Rajbir Parmar, Richard P. Stumpf, Erin A. Urquhart, P. Jeremy Werdell, and Kurt Wolfe. Mobile device application for monitoring cyanobacteria harmful algal blooms using Sentinel-3 satellite Ocean and Land Colour Instruments. *Environmental Modelling and Software*, 109:93–103, November 2018.
- [30] G. Schmetterer. *Cyanobacterial Respiration*, page 409–435. Springer Netherlands, 1994.
- [31] Richard P. Stumpf, Timothy W. Davis, Timothy T. Wynne, Jennifer L. Graham, Keith A. Loftin, Thomas H. Johengen, Duane Gossiaux, Danna Palladino, and Ashley Burtner. Challenges for mapping cyanotoxin patterns from remote sensing of cyanobacteria. *Harmful Algae*, 54:160–173, 2016. Global Expansion of Harmful Cyanobacterial Blooms: Diversity, ecology, causes, and controls.
- [32] U.S. National Office for Harmful Algal Blooms. Socio-economic Impacts. Retrieved from <https://hab.whoi.edu/impacts/impacts-socioeconomic/>, 2019.
- [33] W.F. Vincent. Cyanobacteria. In Gene E. Likens, editor, *Encyclopedia of Inland Waters*, pages 226–232. Academic Press, Oxford, 2009.
- [34] T. T. Wynne, R. P. Stumpf, M. C. Tomlinson, R. A. Warner, P. A. Tester, J. Dyble, and G. L. Fahnenstiel. Relating spectral shape to cyanobacterial blooms in the Laurentian Great Lakes. *International Journal of Remote Sensing*, 29(12):3665–3672, 2008.
- [35] Timothy T. Wynne, Richard P. Stumpf, Michelle C. Tomlinson, and Julianne Dyble. Characterizing a cyanobacterial bloom in Western Lake Erie using satellite imagery and meteorological data. *Limnology and Oceanography*, 55(5):2025–2036, August 2010.
- [36] Kai Yan, Junsheng Li, Huan Zhao, Chen Wang, Danfeng Hong, Yichen Du, Yunchang Mu, Bin Tian, Ya Xie, Ziyao Yin, Fangfang Zhang, and Shenglei Wang. Deep learning-based automatic extraction of cyanobacterial blooms from sentinel-2 msi satellite data. *Remote Sensing*, 14(19):4763, September 2022.
- [37] Han Zhai, Hongyan Zhang, Liangpei Zhang, and Pingxiang Li. Cloud/shadow detection based on spectral indices for multi/hyperspectral optical remote sensing imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 144(1):235–253, 2018.
- [38] Xin Zhao, Rongjie Liu, Yi Ma, Yanfang Xiao, Jing Ding, Jianqiang Liu, and Quanbin Wang. Red tide detection method for hy - 1d coastal zone imager based on u-net convolutional neural network. *Remote Sensing*, 14(1):88, December 2021.