LakeNet: Water Quality Monitoring with Satellite Images and CNNs

Julia Schatz
schat127@umn.edu

Christopher Morse
morse164@umn.edu

Broden Wanner
wanne036@umn.edu

Tanmay Agarwal
agarw139@umn.edu

Abstract

Monitoring the water quality of lakes is a challenging task that can provide significant benefits and insights to environmental conservationists, policy-makers and educators alike. While current methods utilize in-situ measurements to project water quality parameters, such methods are expensive and time consuming. This project proposes a Convolutional Neural Network regressor to predict various water quality metrics from multi-spectral images. Testing results show that this method far outperforms conventional methods of remotely estimating these metrics. In addition, the project provides a new dataset of Minnesota lakes used to train, test, and evaluate this network.

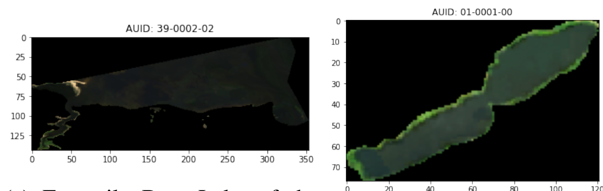

(a) Fourmile Bay, Lake of the Woods

(b) Pine Lake

Figure 1: Sample images of lakes from our synthesized dataset. Note that for visualization, of the 8 channels available for each image, we only use the RGB bands.

1. Introduction

Water is one of the most important natural resources that life on Earth needs to survive. Due to changing land usage, urbanization, pollution, a growing population, and climate change, the access to and the availability of clean water has become a critical issue [24]. According to the WHO and UNICEF, approximately 2.2 billion people around the world do not have access to clean drinking water [28]. Therefore, while monitoring water quality in various regions emerges as an important task, it is difficult to regularly measure water quality for several reasons. Most importantly, it requires experts to take *in-situ* (on-site and real-time) measurements, and many organizations and governments do not have the capacity to monitor water quality at such large scales. Therefore, developing a regional and global capacity to practically and sufficiently measure water quality will help inform policy makers, activists, and water resource managers about anomalies in water quality so they are able to take action to mitigate the threats associated with low-quality water.

In recent decades, researchers have examined the applications of remote sensing data from satellites to estimate water quality. Models based on satellite image data, when evaluated and calibrated on *in-situ* measurements, provide a way to monitor key water quality metrics on a large scale that can be used to detect quality anomalies in water bodies [24]. Water absorbs radiation in the red and near-infrared regions of the electromagnetic spectrum, and optically active water quality variables can therefore be measured from satellite images [18]. Models developed from Landsat, MODIS, and MERIS satellite data have been shown to somewhat accurately estimate various metrics of water quality in specific geographic regions that are optically active [24]. Generally, there are two approaches to water quality metric estimation using remote sensing data: empirical modeling using pure statistics [15] and machine learning models such as Support Vector Regression, Deep Neural Networks, and LSTM networks [20]. However, they suffer from being too simplistic (for statistical methods mostly), having minimal availability of *in-situ* measurements of water quality, or not generalizing well to other geographic areas [24].

To address these issues, we propose using a novel Convolutional Neural Network (CNN) to estimate water quality parameters based on multi-spectral satellite images of lakes. In previous literature, CNNs for regression with water quality metrics have not been studied [24]. However, they offer great potential for capturing nonlinear relationships between satellite image data and water quality measurements, and allowing purely vision-based (hence scalable) and more generalizable approach as the network learns from a large and diverse dataset.

2. Related Work

The use of satellite data for remote water quality monitoring has been well-studied in recent years. Various methods have been researched in this domain, ranging from classical approaches that to fit a model to the data, to deep learning based approaches, that learn the model from a dataset.

2.1. Statistical Regression

While many researchers have found some success in the application of statistical regression models to draw correlations between satellite imagery and *in-situ* data, these models are often not generalizable. This means that the relationship derived from the training data is not often transferable to new data. While statistical regression models (e.g. linear regression and SVR) are relatively easy to implement, nonlinear relationships between variables make simple linear regression less effective. Studies have shown that SVRs are better equipped to capture these nonlinear relationships in the data, but both lack the transferability and accuracy of deep learning models [24].

2.2. Deep Learning for Regression

More recently, with the availability of training data and computing power, the focus of this area of research has shifted towards deep learning for remote water quality monitoring [24]. Many studies have shown Long Short Term Memory (LSTM) neural networks to perform well on tasks of predicting water quality from time series data [34, 10, 30, 37]. However, these methods face a drawback in that their models take on-site measurements of water quality of a lake as input, making it extremely difficult to scale and generalize the prediction methodology. In order to develop methods that can easily provide insights about water quality without having to take in-situ measurements, recent studies use multi-spectral satellite images to analyze and predict the levels of various substances in a water body [24, 36, 29, 32]. Satellite images have been used to monitor water quality in the Tenmile Lake in Oregan [31], estimate levels of chlorophyll-A in Lake Atitlán, Gautemala [4], and Lake Taihu, China [22, 14], and detect water turbidity [13]. More recently, small datasets have been created that associate satellite data of lakes to in-situ measurements of water quality so as to facilitate the development of deep learning methods for water quality prediction [23]. Owing to these developments, we propose a novel CNN-based method that can learn to estimate various parameters of water quality from our dataset.

3. Dataset

Our approach plans to estimate five common quality metrics: visible depth (as measured by Secchi disk method), dissolved oxygen (O_2), pH, specific conductance, and chlorophyll-A concentration (Chl-A). These statistics are gathered from data available from the Minnesota Pollution Control Agency (MN PCA)[2] and cover 9,149 assessed lake-month pairs in the state. The data is from both volunteer-gathered metrics and lab analysis of the water. These metrics were normalized and transformed using the Box-Cox method to transform them into a gaussian shape.

Geographic shapefiles covering these lakes are also provided by the MN PCA[1]. These shapefiles were used to collect wideband satellite images from the US Geographical Survey Landsat-8 using Google Earth Engine [6]. The Landsat data includes 8 bands covering shortwave infrared to ultrablue portions of the electromagnetic spectrum, with a resolution of 30 meters. These data have been atmospherically corrected using LaSRC (Land Surface Reflectance Code) [25] and includes a cloud, shadow, water and snow mask produced using CFMASK (C Function of Mask), as well as a per-pixel saturation mask [5]. We then pad to square and resize the lake images to 128x128 resolution.

From these combined data, 9,149 lake-time combinations were chosen by associating water quality data to satellite images taken from the same month, between May and September in the years 2013 (when Landsat-8 was launched) to 2019. Using only summer months was chosen to minimize the impact of ice on the satellite images. The dataset includes 922 unique lakes.

3.1. Applicability of Data

The allowed values for the chosen water quality metrics vary greatly between regions, individual lakes, and for what purpose the lake is being evaluated, so it is difficult to give a single metric on whether this data spans both acceptable and unacceptable values. However, because these lakes are sampled equally from all of Minnesota, including lakes considered to have one or more impairments by the MN PCA, we can reasonably conclude that they are representative of values for these metrics in Minnesota. Therefore, any model trained on this dataset may have difficulty estimating these values for lakes in other regions where the normal range of values is different.

4. Baseline Methods

To determine the limitations of existing models, we implemented and applied multiple methods from recent literature. The results of the best methods are summarized in Table 1. While these methods are the best available in the literature, they provide poor estimations of our water quality measurements. For the case of dissolved oxygen, these results show that none of the regressors were able to establish a correlation with the available data.

Table 1: Summary of baseline method r^2 values

Method	Depth	O_2	pH	Cond.	Chl-A
Linear	0.33	-0.04	0.33	0.34	0.32
Ridge	0.29	0.01	0.31	0.35	0.29
AdaBoost	0.34	-0.14	0.37	0.35	0.09
GradientBoost	0.37	-0.03	0.13	0.15	0.27
SVR	0.38	0.11	0.34	0.49	0.41
ANN	0.33	0.01	0.41	0.36	0.37

4.1. Statistical Regressors

Our implementation is inspired by [17] and applies a separate regression model for each parameter. All models use a linear kernel. Before providing the data to the trainer, the lake image is vectorized and principal component analysis is applied to reduce the dimensionality of the data. For Support Vector Regression (SVR), we achieved much lower r^2 values (on shared parameters) than either [17] or [24]. We also tried simple linear regression and ridge regression, as well as two ensemble regressors (i.e. Adaptive Boosting and Gradient Boosting), but none performed as well as the SVR model. Furthermore, as we increased the size of the dataset, these prediction accuracies were not improved. This suggests that these regressors are not capable of accurately extracting relationships from our robust dataset.

4.2. Artifical Neural Network

Our implementation uses a neural network with 5 hidden layers of 100 neurons each, inspired by [24]. A separate model was used to estimate each parameter. Before providing the data to the trainer, the lake image is vectorized and principal component analysis is applied to reduce the dimensionality of the data. All models used the tansig activation function. This method achieved much lower r^2 values (on shared parameters) than reported by [24].

5. Method

Convolutional Neural Networks (CNN) have been shown to perform extremely well on tasks of image classification and regression. Since CNNs provide a method to treat the input as an image and consider various spatial dependencies of the data in generating an output, they emerge as an ideal choice in developing a network that can learn from the features of the input satellite images. While extensive work has been in developing CNNs for popular datasets such as ImageNet, CIFAR-10 and CIFAR-100, very limited research exists on developing CNNs for classification and regression for water quality of water bodies. One study by Pu et. al. uses CNNs to classify lakes on the basis of their water quality level using data from the Landsat-8 satellite [21], and while other work has explored using CNN-LSTMs to make short term predictions of water quality [3], they require *in-situ* measurements of various water quality indicators as input, which can be extremely time consuming and expensive to collect. Hence, there is a need to develop CNN models that can predict the levels of various water quality metrics (and not just classify) and take multi-spectral satellite images as input instead of *in-situ* measurements so as to provide a more scalable and cost effective method of water quality monitoring.

5.1. Existing Architectures

Prior to the development of LakeNet, we tested eight CNN architectures that have been shown to perform well on various tasks on image classification and regression. To that end, we train the VGG [26], squeezenet [9], shufflenet [16], alexnet [12], densenet [8], resnet18 [7], resnext50 [33], wideresnet50 [35] and mnasnet [27] on our dataset. We preprocess the input by selecting the RGB channels of the multi-spectral data available, removing small images from the dataset, and resizing all images to a size of 100×100. Since all of these existing models return output of the same shape as the input, we add an additional linear transformation layer to reshape the output to $batch\_size \times 5$, so as to generate a prediction of the five water quality parameters. We then train the model on the average summed L2 loss for each parameter. Results of the L2 loss on the validation dataset for various epochs can be seen in Figure 2.

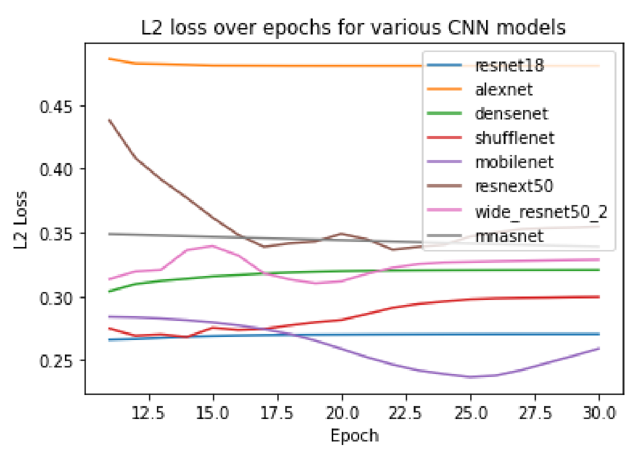

Figure 2: L2 Loss for various epochs on the validation dataset. Note that VGG and squeezenet performed poorly enough that their values are much higher on the *y* axis to not be visualized

The two models that performed the best on our dataset are resnet18 and mobilenet, both of which utilize residual blocks in their network architecture. In contrast, the VGG and squeezenet models - which were developed as lower-complexity models with much fewer layers and parameters than other CNN models with similar accuracy - perform extremely poorly on the dataset. Therefore, we deduced that a suitable network architecture for our dataset would be dense and deep, but not so much that a residual block would be required for adequate learning.

5.2. LakeNet

LakeNet is a novel regressor convolutional neural network that consists of three convolutional blocks before a fully connected layer, whose input is a $8 \times 128 \times 128$ input image (a 128×128 image with 8 channels) and whose output is a single parameter estimation. Each *convolutional block* consists of two pairs of convolutional and ReLu layers, followed by a single 2×2 max pooling layer. Except for the first block (where the depth remains constant), each block doubles the input depth. The LakeNet architecture is shown in Figure 3 and the second block is shown in Figure 4. The model was implemented in PyTorch [19] and trained using GPUs.

Prior to training the model, each band of each lake image was normalized to a 0-1 range and each water quality metric was standardized by subtracting the mean and dividing by the standard deviation. Attempts were made to further preprocess the images such as applying blurring effects and supplementing the dataset with rotated images; however, these were not shown to increase the performance of the model.

Instead of using a single LakeNet model to estimate all 5 water quality metrics at once, we train LakeNet on each metric individually, resulting in 5 separately trained models that can be used in tandem for water quality metric prediction. We partitioned our dataset in a randomly stratified manner based on the water quality metrics into 70% being used for training and 30% being used for testing. We then trained the model using the Adam optimizer [11] with a starting learning rate of 0.0001 and mini-batches of size 100 for 150 epochs. Mean-square error (MSE) was found to perform the best for training the model after evaluating multiple loss functions. All code used for the model can be found in Appendix C.

5.3. Other Methods

During the development of LakeNet, other methods were implemented and tested to accumulate a breadth of study and provide our proposed approach with greater significance. Notably, we represented each lake as a 16-dimensional vector containing the means and standard deviations of all eight bands. We then trained each of the baseline regressors on this new representation, and found that the GradientBoost regressor was the only method which provided adequate results. Table 2 shows the r^2 and MSE values for each of the five water quality metrics using the GradientBoost regressor. We hypothesize that this representation may have led better results than the original due to the reduction of noise in the data through this vectorization. However, this method is comparatively worse than LakeNet for every metric except for dissolved oxygen (the results for LakeNet are described in the following section). It is unclear why this method performs so well in this metric, which might be worth further study.

Table 2: Summary of GradientBoost on vectorized method r^2 values

Parameter	r^2	MSE
Secchi Depth	0.558	0.934
O_2	0.521	3.022
pH	0.549	3.712
Conductance	0.372	10164.967
Chlorophyll-A	0.112	103.979

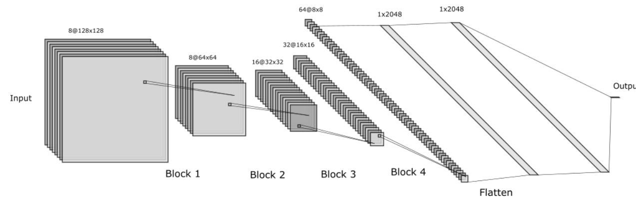

Figure 3: The LakeNet architecture

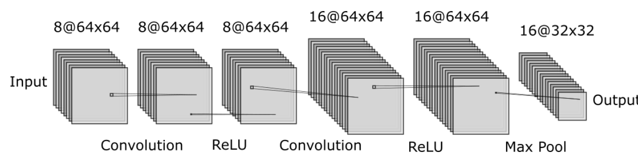

Figure 4: Convolutional block 2

6. Results

The performance of LakeNet on the test set was evaluated based on r^2 and MSE values for all five metrics. These values are provided in Table 3. A plot showing the true and predicted values of conductance (sorted by true conductance value) is in Figure 5. As with the baseline methods, LakeNet predicts dissolved O_2 concentration poorly. This may be because dissolved oxygen is not optically active, leading to little information about it appearing in the surface reflectance [24]. We hypothesize that LakeNet is instead using correlations from other metrics to attempt to predict O_2. Other metrics have possible physical interpretations-Transparency depth can be determined optically, pH and conductance are partially measures of dissolved material in the water that may be optically active, and Chlorophyll-A concentration is related to plant life in the lake.

Table 3: Summary of LakeNet method r^2 values

Parameter	r^2	MSE
Secchi Depth	0.672	0.322
O_2	0.435	0.322
pH	0.563	0.292
Conductance	0.880	0.120
Chlorophyll-A	0.559	0.410

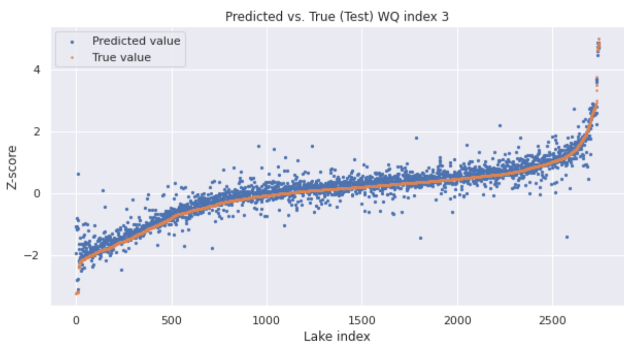

Figure 5: Predicted and true values for conductance in the test dataset

6.1. Intermediate network analysis

Output from intermediate layers of CNNs can be analyzed to understand the rationale behind a network's decisions. As an example, Figure 6 shows the input channels as well as the output channels from the first convolutional block (Conv1) of the trained Secchi depth regressor. Comparing each individual channel between the input and output from Conv1 we find that the network has learned to almost completely ignore the Green, Red, Shortwave, and Panchromatic bands. For these bands, the network ignores pixels in the interior of the lake but only learns to focus on the pixel values near the coast. Subsequently, the model also learns to give importance to the Ultra Blue, Blue and Near Infrared bands the most. For brevity, outputs for deeper layers and other networks are not shown.

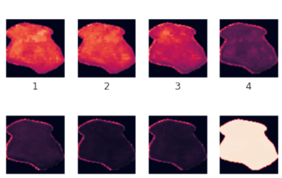

(a) Input Channels

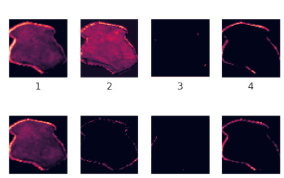

(b) Output from Conv1

Figure 6: Input channels and intermediate outputs

Table 4: Comparison of different loss functions for Secchi Depth regression

Loss Function	r^2
MSE	0.658
L1	0.639

Table 5: Comparison of different optimizers for Secchi Depth regression

Optimizer	r^2	MSE Loss
Adam	0.649	0.344
RMSProp	0.604	0.394
AdaGrad	-0.205	1.202

Table 6: Comparison of number of layers for Secchi Depth regression

Layers	MSE	r^2
2	0.342	0.677
3	0.357	0.662
4	0.351	0.667
5	0.368	0.652

6.2. Ablation Studies

In order to validate the performance of our proposed model, we conducted various ablation studies while varying loss functions, optimizers, and the number of layers in the convolutional network. Results from this ablation for Secchi Depth are shown in tables 4, 5, and 6. Our combination of optimizer, loss function, and number of layers was chosen to be the most general and strongest of these combinations. Particularly in the case of choosing the number of layers, 4 layers provided the most consistently high results between dataset splits while not being prone to overfitting.

7. Conclusion and Future Work

Due to urbanization, climate change, a growing population, and various other factors, it has become increasingly necessary to monitor the quality of water for safe use. Classical methods require experts to gather *in-situ* measurements, but this large-scale procedure is expensive and inefficient. Furthermore, statistical regression models provide low accuracy estimations on our dataset, which limits their viability for proper quality estimation.

Our work provides several contributions. The first is LakeNet, the first regressor convolutional neural network for predicting water quality metrics. LakeNet vastly outperformed all existing methods and the state of the art CNNs for estimating the five water quality metrics on our data. Along with this network, we provide a large dataset of satellite images paired with each of their five water quality metrics. Finally, our quantitative analyses and ablation studies provide insight into comparative abilities and resilience of LakeNet.

Although LakeNet is an improvement on the state of the art, we believe that further improvements can be made. This could include the synchronization of data from other satellites or the use of different network architectures. Moreover, ensuring the quality of the metrics gathered for the training of LakeNet is a vital step to take. Another possible direction of research could be attempting to perform regression on each water pixel instead of each lake. This would capture the complexity of the water quality of larger lakes or water bodies whose metrics might differ in distinct sections. However, to do this in a supervised way would require many complex measurements that could be infeasible to gather.

A. Contributions of Group Members

Our group did not define specific roles. A rough overview of individual contributions is however provided. J. Schatz and B. Wanner collected and refined the dataset. C. Morse and T. Agarwal researched related work. C. Morse, B. Wanner, and J. Schatz implemented and tested baseline methods. T. Agarwal developed the prototype CNN and researched architectures. All group members contributed equally to developing the final LakeNet architecture, writing both reports, and creating both presentations.

B. Comments From Committee

We thank the committee members for their insightful comments during the presentation. A paraphrasing of the comments and our responses are provided.

B.1. Span of collected data

Q: Does the data you have collected span not only metrics for healthy lakes but also unhealthy lakes?

A: Allowable metrics for lake suitability vary greatly between region, individual lake, and the purpose for which the lake is being assessed (aquatic life, recreation, microbial life, etc.) Thus it is difficult to give an overarching or complete answer, however we believe that our data is representative and spans both healthy and unhealthy lakes as described in section 3.1, which we have added in response to this comment.

B.2. Physical basis of results

Q: Why are all methods significantly worse at predicting O_2?

A: Dissolved oxygen is not optically active[24], and thus does not directly affect the image in the available bands. It is possible that our network is using detectable metrics correlated with dissolved oxygen but do not follow exactly, which simple regressors are likely worse at determining. We have added discussion of this to section 6.

B.3. Interpretation of evaluation metric

Q: How do the r^2 values capture the performance of your model?

A: r^2 measures the ratio of the mean square of residuals of the predictions to the variance of the true data:

$$r^2 = 1 - \frac{\frac{1}{N}\sum_i^N (y_{pred,i} - y_{true,i})^2}{\mathrm{Var}(y_{true})}$$

Thus an r^2 of 1 is a perfect prediction (implying zero residuals,) and an r^2 of 0 or lower means that guessing the mean would be a better guess than the model's predictions. In many cases, r^2 can be interpreted as the proportion of variance that is captured by the model and is a common metric used in the literature to evaluate regression models.

C. Dataset and Code

The created dataset and CNN implementation can be found here. See the README document in the folder for more information on the dataset structure and code.

References

[1] Minnesota Pollution Control Agency. Spatial data, 2019. 2

[2] Minnesota Pollution Control Agency. Surface water data, 2019. 2

[3] Rahim Barzegar, Mohammad Taghi Aalami, and Jan Adamowski. Short-term water quality variable prediction using a hybrid cnn–lstm deep learning model. *Stochastic Environmental Research and Risk Assessment*, 34(2):415–433, Feb 2020. 3

[4] Africa I. Flores-Anderson, Robert Griffin, Margaret Dix, Claudia S. Romero-Oliva, Gerson Ochaeta, Juan Skinner-Alvarado, Maria Violeta Ramirez Moran, Betzy Hernandez, Emil Cherrington, Benjamin Page, and Flor Barreno. Hyperspectral satellite remote sensing of water quality in lake atitlán, guatemala. *Frontiers in Environmental Science*, 8:7, 2 2020. 2

[5] Steve Foga, Pat L. Scaramuzza, Song Guo, Zhe Zhu, Ronald D. Dilley, Tim Beckmann, Gail L. Schmidt, John L. Dwyer, M. Joseph Hughes, and Brady Laue. Cloud detection algorithm comparison and validation for operational landsat data products. *Remote Sensing of Environment*, 194:379–390, 2017. 2

[6] Noel Gorelick, Matt Hancher, Mike Dixon, Simon Ilyushchenko, David Thau, and Rebecca Moore. Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 2017. 2

[7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. 3

[8] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks, 2018. 3

[9] Forrest N. Iandola, Song Han, Matthew W. Moskewicz, Khalid Ashraf, William J. Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and ¡0.5mb model size, 2016. 3

[10] Huimin Jia and Xiaofeng Zhou. Water Quality Prediction Method Based on LSTM-BP. In *Proceedings - 2020 12th International Conference on Intelligent Human-Machine Systems and Cybernetics, IHMSC 2020*, volume 1, pages 27–30. Institute of Electrical and Electronics Engineers Inc., aug 2020. 2

[11] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. 4

[12] Alex Krizhevsky. One weird trick for parallelizing convolutional neural networks, 2014. 3

[13] Wen Chin Lee, Chin Lun Lai, and Chun Kai Chang. An integrated water turbidity detection method based on multispectral analysis. In *2011 International Conference on Consumer Electronics, Communications and Networks, CECNet 2011 - Proceedings*, pages 5497–5500, 2011. 2

[14] Junsheng Li, Bing Zhang, Qian Shen, Lei Zou, and Liwei Li. Monitoring water quality of Lake Taihu from HJ-CCD data using empirical models. In *International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 812–815, 2012. 2

[15] Colleen M. Long and Tamlin M. Pavelsky. Remote sensing of suspended sediment concentration and hydrologic connectivity in a complex wetland environment. *Remote Sensing of Environment*, 129:197–209, 2013. 1

[16] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design, 2018. 3

[17] Philipp M Maier and Sina Keller. Machine learning regression on hyperspectral data to estimate multiple water parameters. In *2018 9th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, pages 1–5. IEEE, 2018. 3

[18] Nima Pahlevan, Brandon Smith, John Schalles, Caren Binding, Zhigang Cao, Ronghua Ma, Krista Alikas, Kersti Kangro, Daniela Gurlin, Nguyen Hà, Bunkei Matsushita, Wesley Moses, Steven Greb, Moritz K. Lehmann, Michael Ondrusek, Natascha Oppelt, and Richard Stumpf. Seamless retrievals of chlorophyll-a from sentinel-2 (msi) and sentinel-3 (olci) in inland and coastal waters: A machine-learning approach. *Remote Sensing of Environment*, 240:111604, 2020. 1

[19] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. 4

[20] Kyle T. Peterson, Vasit Sagan, Paheding Sidike, Elizabeth A. Hasenmueller, John J. Sloan, and Jason H. Knouft. Machine learning-based ensemble prediction of water-quality variables using feature-level and decision-level fusion with proximal remote sensing. *Photogrammetric Engineering and Remote Sensing*, 85(4):269–280, apr 2019. 1

[21] Fangling Pu, Chujiang Ding, Zeyi Chao, Yue Yu, and Xin Xu. Water-quality classification of inland lakes using landsat8 images by convolutional neural networks. *Remote Sensing*, 11(14), 2019. 3

[22] Shen Qian, Zhang Bing, Wu Chuanqing, and Li Junsheng. Retrieval of Chlorophyll-a and suspended matter concentration in water supply sources of Wuxi and Suzhou using multi-spectral remote sensing images. In *2008 International Workshop on Earth Observation and Remote Sensing Applications, EORSA*, 2008. 2

[23] Matthew R. V. Ross, Simon N. Topp, Alison P. Appling, Xiao Yang, Catherine Kuhn, David Butman, Marc Simard, and Tamlin M. Pavelsky. Aquasat: A data set to enable remote sensing of water quality for inland waters. *Water Resources Research*, 55:10012–10025, 11 2019. 2

[24] Vasit Sagan, Kyle T. Peterson, Maitiniyazi Maimaitijiang, Paheding Sidike, John Sloan, Benjamin A. Greeling, Samar Maalouf, and Craig Adams. Monitoring inland water quality using remote sensing: potential and limitations of spectral

indices, bio-optical simulations, machine learning, and cloud computing, jun 2020. 1, 2, 3, 4, 6

[25] Kristi Sayler. Landsat 8 Collection 1 (C1) Land Surface Reflectance Code (LaSRC) Product Guide. Technical report, U.S. Geological Survey, 2020. 2

[26] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2015. 3

[27] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile, 2019. 3

[28] UNICEF and WHO. 1 in 3 people globally do not have access to safe drinking water – UNICEF , WHO New report on inequalities in access to water, 2019. 1

[29] Jinzhe Wang, Junping Zhang, Tong Li, and Xiao Wang. Water quality analysis of remote sensing images based on inversion model. In *International Geoscience and Remote Sensing Symposium (IGARSS)*, volume 2018-July, pages 4861–4864. Institute of Electrical and Electronics Engineers Inc., oct 2018. 2

[30] Yuanyuan Wang, Jian Zhou, Kejia Chen, Yunyun Wang, and Linfeng Liu. Water quality prediction method based on LSTM neural network. In *Proceedings of the 2017 12th International Conference on Intelligent Systems and Knowledge Engineering, ISKE 2017*, volume 2018-Janua, pages 1–5. Institute of Electrical and Electronics Engineers Inc., jul 2017. 2

[31] Michael T Waxter. Analysis of landsat satellite data to monitor water analysis of landsat satellite data to monitor water quality parameters in tenmile lake, oregon quality parameters in tenmile lake, oregon. *Civil and Environmental Engineering Master's Project Reports*, 2014. 2

[32] Yuanfeng Wu, Bing Zhang, Junsheng Li, Hao Zhang, Qian Shen, and Di Wu. A first approach to automatic monitoring prodedure of case II water quality from HJ-1 satellite images. In *International Geoscience and Remote Sensing Symposium (IGARSS)*, volume 2, 2009. 2

[33] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks, 2017. 3

[34] Qiangqiang Ye, Xueqin Yang, Chaobo Chen, and Jingcheng Wang. River water quality parameters prediction method based on lstm-rnn model. pages 3024–3028. Institute of Electrical and Electronics Engineers Inc., 6 2019. 2

[35] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks, 2017. 3

[36] Zhihui Zeng and Hui Lv. Prediction of urban water supply quality based on multi-sensor technology. In *2009 4th IEEE Conference on Industrial Electronics and Applications, ICIEA 2009*, pages 2214–2217, 2009. 2

[37] Huiqing Zhang and Kemei Jin. Research on water quality prediction method based on AE-LSTM. In *Proceedings - 5th International Conference on Automation, Control and Robotics Engineering, CACRE 2020*, pages 602–606. Institute of Electrical and Electronics Engineers Inc., sep 2020. 2