SNITCHING ON DITCHES: TRACKING SALT MARSH HEALTH USING TRANSFER LEARNING

by

SOPHIA SOMERSCALES

A THESIS

Presented to the Department of Data Science
and the Robert D. Clark Honors College
in partial fulfillment of the requirements for the degree of
Bachelor of Science

April 2023

An Abstract of the Thesis of

Sophia Somerscales for the degree of Bachelor of Science
in the Department of Data Science to be taken June 2023

Title: Snitching on Ditches: Tracking Salt Marsh Health Using Transfer Learning

Approved: *Professor Stephen Fickas*
Primary Thesis Advisor

Coastal salt marshes offer crucial ecological benefits, including carbon sequestration, habitat for many species, and protection against storm surges and erosion. However, human activity has led to significant dieback of these ecosystems on both a national and global scale. Much of the northeastern US salt marshes are experiencing exacerbated loss due to grid ditching, an outdated practice in which standing pools of water were drained by a series of narrow ditches to reduce mosquito populations. Identifying ditches is an important step in tracking salt marsh health, yet ecologists currently lack an efficient method to do so, mostly relying on walking the fields between tides or manually delineating ditches in aerial imagery. This project investigates an alternate workflow for identifying ditches in high-resolution drone imagery captured by the Salt Marsh UAV group at University of Massachusetts Amherst. I implement U-Net, a machine learning that originates from medical imaging, to sift through all the varied water features in a single salt marsh site and classify each pixel in an image as background, ditch, or non-ditch, a process called semantic segmentation. Ultimately, the goal is to produce georeferenced shapefiles that precisely locate ditches on the ground. I use pre-trained versions of U-net and experiment with various parameters to tune the models for optimal results. This is a form of transfer learning, taking models from one domain and repurposing to another. MobileNet-UNet exhibits the highest performance and produces strong ditch segmentation results that ecologists

can utilize with minimal post-processing. Future research should experiment with using multispectral bands like near infrared (NIR) and short-wave infrared (SWIR) or a Digital Elevation Model (DEM) to provide the model with more information. This project provides ecologists an automated method of identifying ditches and demonstrates that transfer learning is a viable alternative to traditional remote sensing water feature extraction methods.

Acknowledgements

I would first like to thank my primary advisor, Dr. Steve Fickas, for mentorship throughout the past four years. Not only did he provide guidance for this project, but he also introduced me to the field of data science and encouraged me to explore environmental applications that aligned with my interests. Without his support, I would not have embarked on this academic trajectory that led me to pursue graduate studies in this field. I would also like to extend my gratitude to Dr. Daphne Gallagher, my CHC representative, for her assistance during the initial planning stages of this project. Her feedback and guidance played a significant role in shaping the aspirations of my thesis.

Furthermore, I would like to express my deepest appreciation to the Salt Marsh UAV team at the University of Massachusetts Amherst. I am immensely grateful to them for proposing this research topic and granting me access to their data and collective knowledge. Their ecological and remote sensing insight was invaluable, and it was a pleasure to be a part of their group this past year.

Finally, I would like to thank my family and friends who supported me through the ups and downs of this journey and tolerated my sudden interest in salt marshes and rants about the model's performance. Thank you, mom and dad, for your love and encouragement. This project would not have been possible without you.

Table of Contents

List of Figures

List of Tables

Introduction

Coastal salt marshes have significant ecological value due to the crucial roles they play in protecting against storm surges and erosion (Donatelli et al. 2018), regulating atmospheric greenhouse gas levels via carbon sequestration (Lockwood and Drakeford 2020), and providing sanctuary for many fish, wildlife, and waterfowl (Kennish 2001). Over the past few decades, these fragile ecosystems have experienced substantial platform dieback from a host of human-caused stressors. Platform dieback occurs through a process called slumping in which sections of the platform banks fragment and collapse into the creek network. It is estimated that over 50% of the original US salt marsh habitat have been lost (Watzin and Gosselink 1992). The most notable human impacts include the transformation of salt marshes for agricultural, residential, and industrial use, sea level rise from global warming, subsidence from groundwater and petroleum extraction, and the practice of grid ditching (Kennish 2001; Jin et al. 2016; Watson et al. 2017).

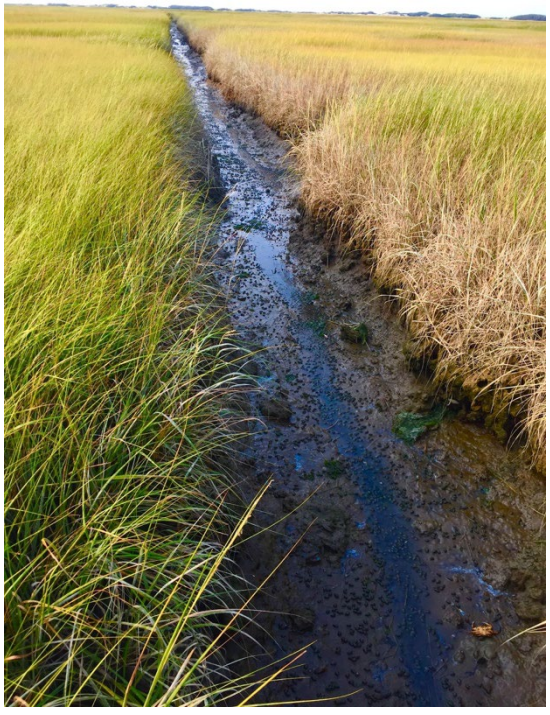

Figure 1. An example of a narrow ditch cutting through a Massachusetts salt marsh. Credit: Ryan Wicks.

Grid-ditching is an approach to managing mosquito populations in which narrow ditches are dug at regular intervals to drain stagnant pools of water where mosquitos are likely to breed. Grid-ditching is extensive in northeastern salt marshes because of efforts by the Civilian Conservation Corps (CCC) to reduce the health impacts of mosquitos and provide employment opportunities as a part of the "New Deal" initiative. It is now an outdated method but the prevailing ditches

impact more than 90% of New England salt marshes by altering the physical and hydrological characteristics of the platform (Kennish 2001).

Tracking salt marsh dieback is an urgent issue since platforms exponentially lose their capability to trap sediment as they recede, meaning dieback triggers a feedback loop that accelerates further loss (Donatelli et al. 2018). To monitor platform changes, identifying ditches is an important step. However, remote sensing scientists do not currently have an efficient way to do this.

Most traditional remote sensing methods of water segmentation assess spectral properties. The normalized difference water index (NDWI) is a multispectral index that combines an image's green band and near-infrared band to differentiate water from vegetation (McFeeters 1996). Similarly, the modified normalized difference water index (MNDWI) combines the green band and short-wave infrared band to extract water features (Xu 2006) and is one of the most widely used (Feyisa 2014). The automated water extraction index (AWEI) considers five spectral bands and suppresses noise from shadowy areas to improve accuracy (Feyisa 2014).

While spectral indices have proved useful, they are not well-suited for the aim of this study. These indices are typically used on satellite imagery, especially Landsat imagery for regional water segmentation, with a moderate spatial resolution of 30 meters. As a result, spectral indices tend to miss small features like ditches that are much thinner than 30 meters. The drone flight images in this study have a spatial resolution of 0.026 meters, more than a hundred times higher than Landsat.

However, the bigger issue is that spectral indices rely on the distinct spectral signature of water, meaning that all water pixels are classified the same. There is no way to segment an image

into multiple water feature classes using spectral properties only. One alternative is using an object-oriented approach through a software like eCognition. This professional remote sensing system utilizes color, shape, texture, and object size in addition to local neighborhood statistics to perform classification. While eCognition performs well on many remote sensing tasks (Tamta et al. 2015; Xing and Shen 2018; Yang et al. 2018; Xue and Lin 2020) it is not open-source and again has mostly been tested on satellite imagery with moderate spatial resolution.

Background

AI in Geoscience

Earth science is at a critical point of transformation as artificial intelligence (AI) continues to spread throughout the many subdomains and enhance the ability of geoscientists to monitor the Earth's subsystems and respond to environmental changes. Sun et al. (2022) summarize existing applications of AI to all major Earth spheres and find that while Earth AI remains in its beginning phase, recent literature shows promising results on all fronts. The main challenge geoscientists face is the lack of standardized, labeled datasets for training and machine learning (ML) expertise for model development and optimization. Nonetheless, there have been successes in the improved prediction of earthquakes, hurricanes, drought, wildfires, sea ice thickness, groundwater levels, and more.

CNN Water Segmentation

Hydrology is one of the earth sciences fields that has greatly benefited from AI. In addition to research regarding water forecasting, water quality, rainfall runoff, and river sediments and discharge (Sun et al., 2022), there is substantial research on water segmentation (Akiyama et al. 2020; Miao et al. 2018; Singh et al. 2019; Weng et al. 2020). Water segmentation is possible with pixel-wise classification and scholarship in this domain tends to either compare the segmentation results of different models or tinker with an existing model to optimize it for a specific task. Most of these segmentation models are convolution neural networks (CNN), with two of the most popular being SegNet (Badrinarayanan et al. 2017) and U-Net (Ronneberger et al. 2015).

In their comparison of SegNet, U-Net, DeepLab, and DenseNet, Sing et al. (2019) found that all CNN models outperformed the traditional Support Vector Machine (SVM) method for segmenting water from ice. Of the four CNNs, SegNet showed the least improvement over the SVM. DenseNet is a newer, less studied architecture and gave poor quantitative results but showed strong generalization ability on the unlabeled data. DeepLab had poor generalization ability but strong quantitative results. U-Net performed the best overall.

CNNs have also been tested for the segmentation of surface waters as current methods for monitoring lakes and rivers tend to be labor-intensive and have low generalization ability (Akiyama et al. 2020; Weng et al. 2020). When trained on RGB river images of sizes 256 x 256 and 512 x 512, SegNet produced strong results at both resolutions, confirming the potential of the CNN approach (Akiyama et al. 2020). A modified version of the SegNet architecture, SR-SegNet, offers even higher accuracy (Weng et al. 2020). SR-SegNet outperformed traditional methods of identifying surface water with an SVM or an NDWI, as well FCN, DeconvNet, and standard SegNet. Miao et al. (2018) propose RRF DeconvNet, along with a new loss function to sharpen segmentation edges, to segment water bodies in high-resolution Google Earth images.

Despite the encouraging prospect of surface water monitoring using CNNs, scholarship in this area has yet to explore the segmentation of different water features. Current research is focused on binary water classification, i.e. whether each pixel in an image is water or not. However, sometimes it is necessary to identify the types of water features present, where a water feature is a conglomeration of pixels that make up a larger collection like a pond or creek. This is the case when tracking salt marsh health. In this instance, whether a water pixel belongs to a ditch or a different water feature is crucial to monitoring the salt marsh platform. There are studies that use ML to perform multi-class pixel-wise land cover segmentation. Enwright et al.

(2019) found that Random Forest (RF) and various CNN classifiers demonstrated high modeling capacity for such image segmentation, but water was still treated as a single class, as is common in land cover analysis.

This study aims to close this gap in research. Water feature segmentation is a challenging remote sensing task since all water features share similar spectral properties, but CNNs offer a promising alternative.

Dataset

There are currently no labeled salt marsh datasets with delineated water features, so part of this study includes creating one from drone imagery. These salt marsh images are from the Essex Bay site in Ipswich, Massachusetts site at low tide in spring. This site was selected over other salt marsh sites for training and testing since it contains moderate, but not severe, slumping and is the most representative of Massachusetts salt marshes. The low tide spring drone flight was chosen because the minimal vegetation provides the highest visibility of water features for the year. Since ditches are relatively static water features, it is unnecessary to obtain a ditch shapefile layer more than once a year.

Figure 2. An orthomosaic of the Essex Bay salt marsh in Ipswich, Massachusetts at low tide in spring.

We use the LabelMe (Wada 2022) annotation software to label 50 RGB Essex Bay images and create corresponding segmentation masks. All images are size 1500 x 2000 with a spatial resolution of 0.026 meters. LabelMe allows the user to draw boundaries on an image and classify all pixels included in that feature. In this way, we can label images beyond the single pixel level. Each water pixel is labeled as belonging to either the “ditch” class or to the “non-ditch” class. The non-ditch water feature class exists to help the model differentiate ditches from all other water features. Any linear channel that appears to be manmade is considered a ditch and any naturally occurring water feature is considered non-ditch, as shown in Figure 3. Pixels not part of labeled features are classified as background. Table 1 shows the distribution of class samples in the dataset.

The labeled dataset is augmented with rotations and vertical and horizontal flips to produce a dataset of 200 images. This process is illustrated in Figure 4. Finally, an 80-10-10 split is used to divide the dataset into training, validation, and testing data. The training dataset contains 160 images and the validation and testing datasets each contain 20 images.

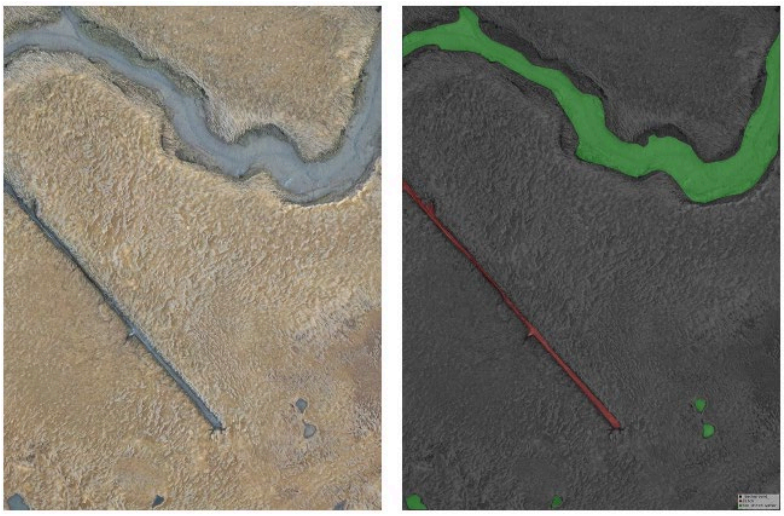

Figure 3. An RGB drone image (left) with its corresponding segmentation map created with LabelMe software (right) where green indicates a non-ditch water feature and red indicates a ditch.

Table 1. Distribution of class samples in the dataset

Class	**Percentage of Pixels**
Non-ditch	13.8
Ditch	1.1
Background	85.1

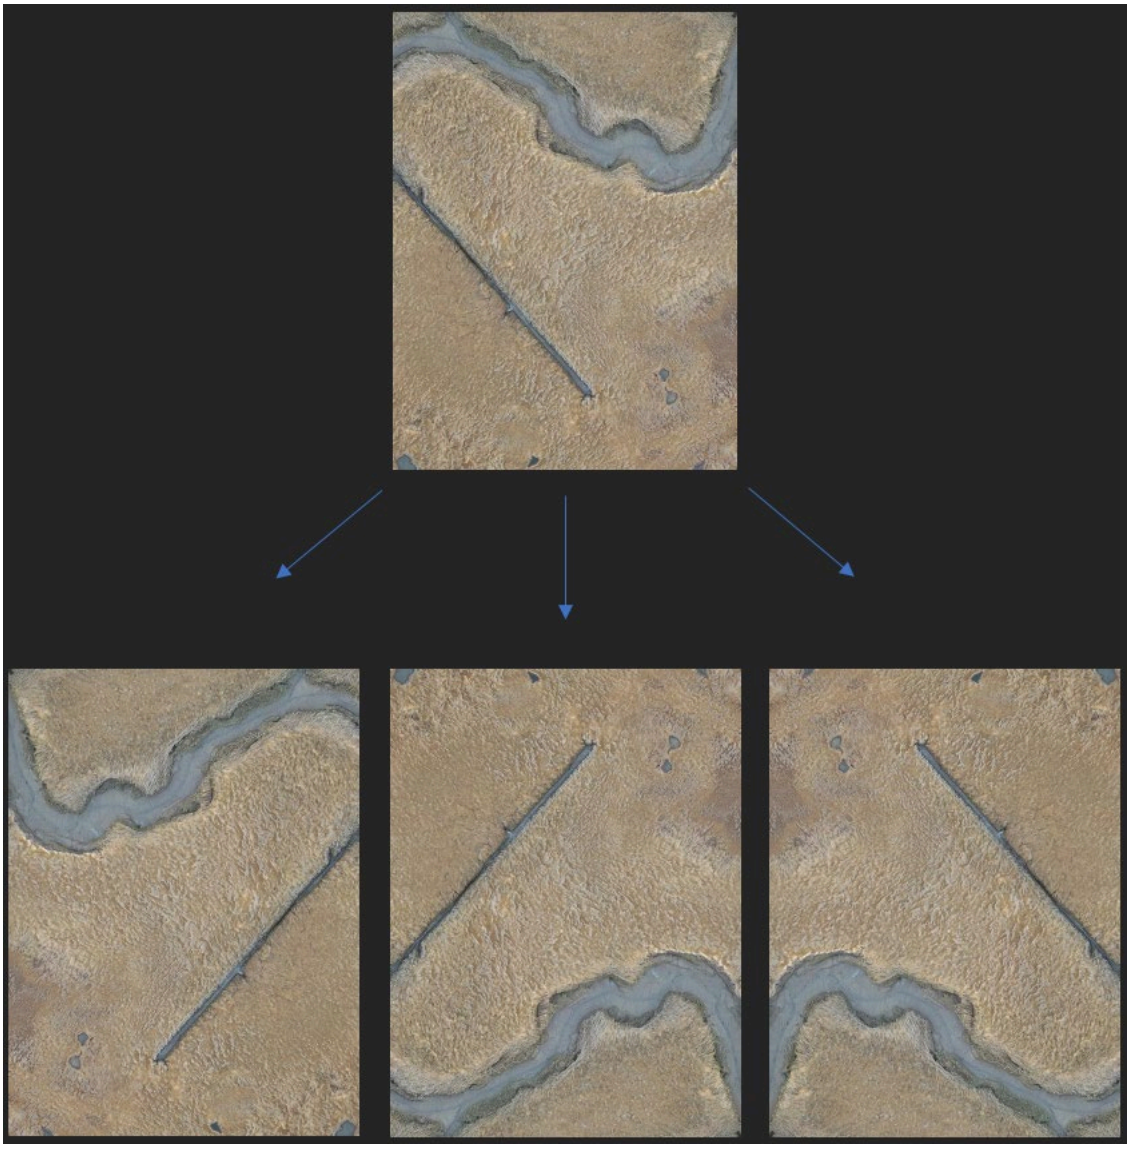

Figure 4. Data augmentation example. The original image (top) is rotated and flipped to produce an additional three unique images (below).

Methods

Performance Metrics

This study uses two standard evaluation metrics for semantic segmentation tasks: Jaccard Index and Dice Coefficient. Both metrics quantify the degree of overlap between the predicted labels and the ground truth labels for a particular class, where 1 indicates perfect overlap and 0 indicates no overlap. Jaccard and Dice scores are more robust measures of model performance than pixel accuracy because they are less impacted by class imbalances in the dataset.

1. Jaccard Index (or Intersection over Union)

$$Jaccard_c = \frac{TP_c}{TP_c + FP_c + FN_c}$$

2. Dice Coefficient (or F1 Score)

$$Dice_c = \frac{2 \cdot TP_c}{2TP_c + FP_c + FN_c}$$

Where TP_c is the number of true positive pixels in class c ∈ C ; FP_c is the number of false positive pixels in c; FN_c is number of false negative pixels in c.

The three classes in our case are ditch, non-ditch , and background.

Model Architectures and Hyperparameters

This study focuses on the U-net architecture due to its success in medical imaging and transferability to remote sensing tasks. Figure 5 displays the traditional U-net architecture as presented by Ronnebeger et al. (2015). The model consists of two main parts: the contracting path and the expansive path. The contracting path uses a series of convolutions, pooling, and downsampling operations to capture increasingly higher-level features from the input image. The expansive path is symmetrical to the contracting path and uses a series of upsampling operations and convolutional layers to recover the spatial resolution of the original image. Low-level features such as edges, colors, and texture are combined with high-level features like object classes or shape using skip connections, shortcuts that feed the output of one layer to layers farther ahead in the network. This preserves spatial information needed for accurate segmentation. The output of the model is a 2-dimensional map where each pixel in the input image is assigned a probability of belonging to one of the three classes. The class with the highest probability is selected for each pixel, resulting in a segmentation mask (Ronnebeger et al. 2015).

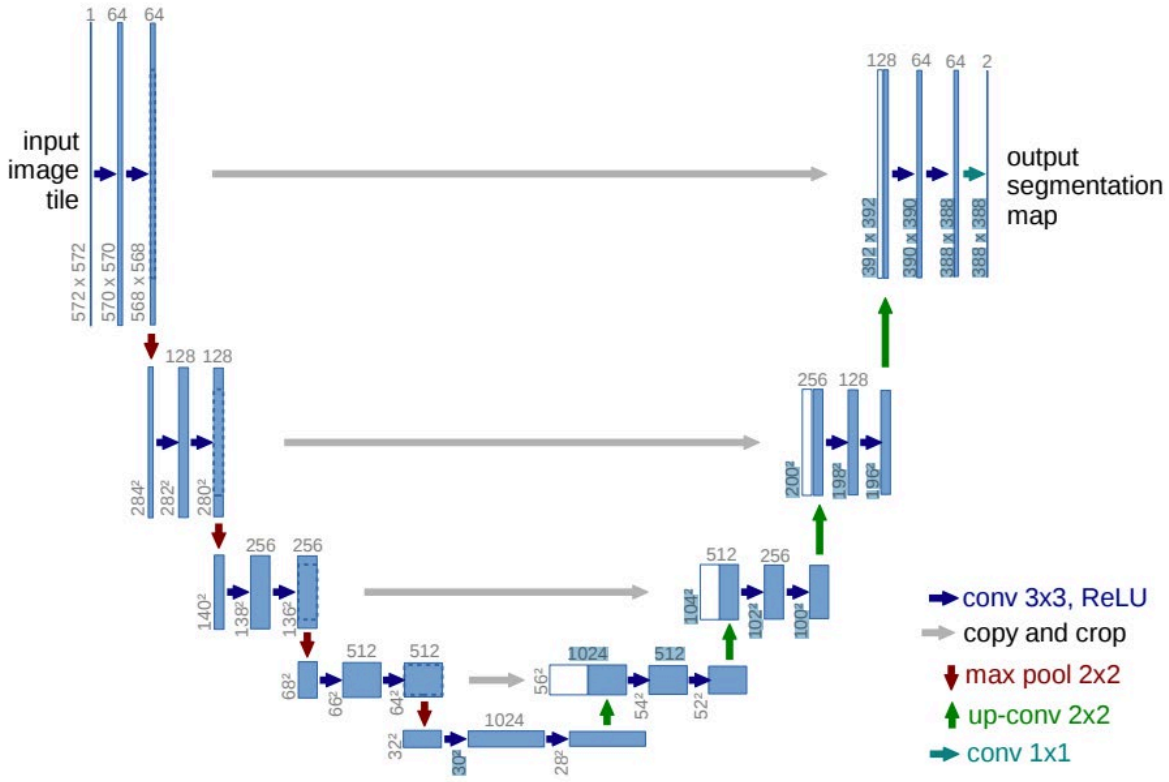

Figure 5. Example of U-net architecture for 32 x 32 image.

Because there are five pooling layers that reduce the resolution by a factor of 2 in the contracting path, input images are required to have dimensions that are multiples of 32. To accommodate for this standard, I downsize my images from 1500 x 2000 to 1472 x 1984.

I also evaluate three convolutional encoders in place of the contracting path: VGG-16, Resnet50, MobileNet. The VGG-16 encoder, originating from image classification, uses very small filters like the classic U-net (Simonyan and Zisserman, 2014). The Resnet50 encoder has a deep residual network that simplifies the training of deep models (He et al. 2016), and the MobileNet encoder features fewer parameters and computations as it was designed for mobile devices (Howards et al. 2017). I use weights for each of these three encoders that are pretrained

on the ImageNet dataset and fine tune the models during training. ImageNet, with over 14 million images, is a promising candidate for remote sensing transfer learning because it includes natural landscapes and bodies of water (Deng et al., 2009).

All models are compiled with the Adam optimizer and a 0.001 learning rate. They are trained with a batch size of 4 and 40 steps per epoch, ensuring a full pass through the training data each epoch. The small batch size is due to computational constraints. The validation batch size is 4 and is used to compute Jaccard and Dice metrics for early stopping.

Loss Functions

This study uses two loss functions as provided in the segmentation_models package: Jaccard and Dice loss. I use these loss functions rather than categorical cross entropy since they translate more directly to mask overlap.

1. Jaccard Loss

$$L = 1 - \frac{1}{3}\sum_{c=1}^{3}\frac{TP_c}{TP_c + FP_c + FN_c}$$

2. Dice Loss

$$L = \frac{1}{3}\sum_{c=1}^{3}\frac{(1+\beta^2)\cdot TP_c}{(1+\beta^2)\cdot FP_c + \beta^2 \cdot FN_c + FP_c}$$

Where gt_c is the ground truth and pr_c is the prediction for class c $\in$ C.

Environment

I use Colab Pro+, a Jupyter notebook service hosted by Google Research. All models are trained using a runtime with 55 GB of available RAM and premium GPUs.

Shapefile Extraction

The final step in the workflow involves converting the segmentation maps output from the best-performing model from raster to vector format. This process transforms the water feature segmentations from collections of pixels into delineated polygons, resulting in shapefiles.

To obtain shapefiles for the Essex Bay ditches, I first use the output segmentation maps (.jpg extension) and their corresponding world files (.pgw extension) to create GeoTIFF images with the Geospatial Data Abstraction Library (GDAL). A GeoTIFF is a georeferenced TIFF image, which is the standard format for raster imagery used in Geographic Information Systems (GIS). I then vectorize the GeoTIFFs using GDAL to create the shapefiles.

Next, I employ GeoPandas to merge all shapefiles into a single file. In this final shapefile, polygons belonging to the background class have a value of 0, the ditch class has a value of 1, and the non-ditch class has a value of 2.

Results

Quantitative Results

Table 2. Performance of all models on the 20 test images using Jaccard loss during training.

	Ditch		Non-ditch		Background			
Model	**IoU**	**Dice**	**IoU**	**Dice**	**IoU**	**Dice**	**Mean IoU**	**Mean Dice**
unet	0.38	0.55	0.83	0.91	0.97	0.99	0.73	0.82
vgg16_unet	0.36	0.52	0.86	0.92	0.98	0.99	0.73	0.81
resnet50_unet	0.41	0.58	0.71	0.83	0.96	0.98	0.69	0.80
mobilenet_unet	0.39	0.57	0.88	0.94	0.98	0.99	0.75	0.83

Table 3. Performance of all models on the 20 test images using Dice loss during training.

	Ditch		Non-ditch		Background			
Model	**IoU**	**Dice**	**IoU**	**Dice**	**IoU**	**Dice**	**Mean IoU**	**Mean Dice**
unet	0.21	0.45	0.79	0.88	0.97	0.98	0.66	0.77
vgg16_unet	0.36	0.53	0.84	0.91	0.97	0.99	0.72	0.81
resnet50_unet	0.14	0.25	0.76	0.87	0.96	0.98	0.62	0.70
mobilenet_unet	**0.45**	**0.62**	0.86	0.93	0.97	0.99	0.76	0.85

As shown in Tables 2 and 3, the models with encoders pretrained on ImageNet (VGG-UNet, ResNet50-UNet, and MobileNet-UNet) generally outperform the U-Net model without pretraining. This demonstrates the advantage of transfer learning even when fine-tuning with much larger images than those used for pretraining. The ImageNet dataset contains images of size 256 x 256 on average (Deng et al., 2009), which are considerably smaller than the 1500 x 2000 salt marsh images used in this study.

MobileNet-UNet consistently outperforms the other models in terms of mean IoU and mean Dice scores, regardless of the loss function used. It also achieves the highest overall performance on the ditch class, obtaining an IoU of 0.45 and a Dice coefficient of 0.62 when trained with Dice loss. The performance of other models on the ditch class remains relatively low with Dice loss.

When trained with Jaccard loss, the other models show improvement for the ditch class. ResNet50-UNet demonstrates the best results, with an IoU of 0.41 and Dice coefficient of 0.58. In summary, Jaccard loss yields better outcomes for all classes, while Dice loss achieves the overall best results for the ditch class specifically.

Qualitative Results

Figures 6, 7, and 8 display the results of the MobileNet-UNet model trained with Dice loss on several test images. Three key observations regarding the segmentation of the ditch class can be made.

First, although the model does not miss any instances of ditch water features, many ditches appear discontinuous in the predicted segmentation despite being continuous in the ground truth segmentation, as exemplified in Figure 6. This occurs because many ditches are partially obscured by vegetation overhang in the drone imagery, and the model likely relies heavily on color for its ditch class prediction. When vegetation overhang is present, those portions of the ditch become unrecognizable to the model. The penalty associated with connecting such ditches by classifying the obstructing vegetation pixels as ditch instead of background is apparently too high for the model to attempt.

Second, the model tends to misclassify narrow non-ditch water features as ditches, as demonstrated in Figure 7. In these images, the natural creek channel has a width similar to ditches, and the model classifies linear sections of the channel as ditch while classifying curvy sections as non-ditch. This suggests that the model has learned the narrow, linear characteristics of ditches but struggles when these features are mixed with non-ditch water elements.

Lastly, despite these limitations, the segmentation results are still valuable for the purpose of delineating ditches, particularly with some post-processing. The non-ditch and background

classes can be masked out in the shapefile, and any disjointed ditches can be connected using buffers or other GIS methods. This approach enhances the overall usability of the segmentation results for monitoring salt marsh health.

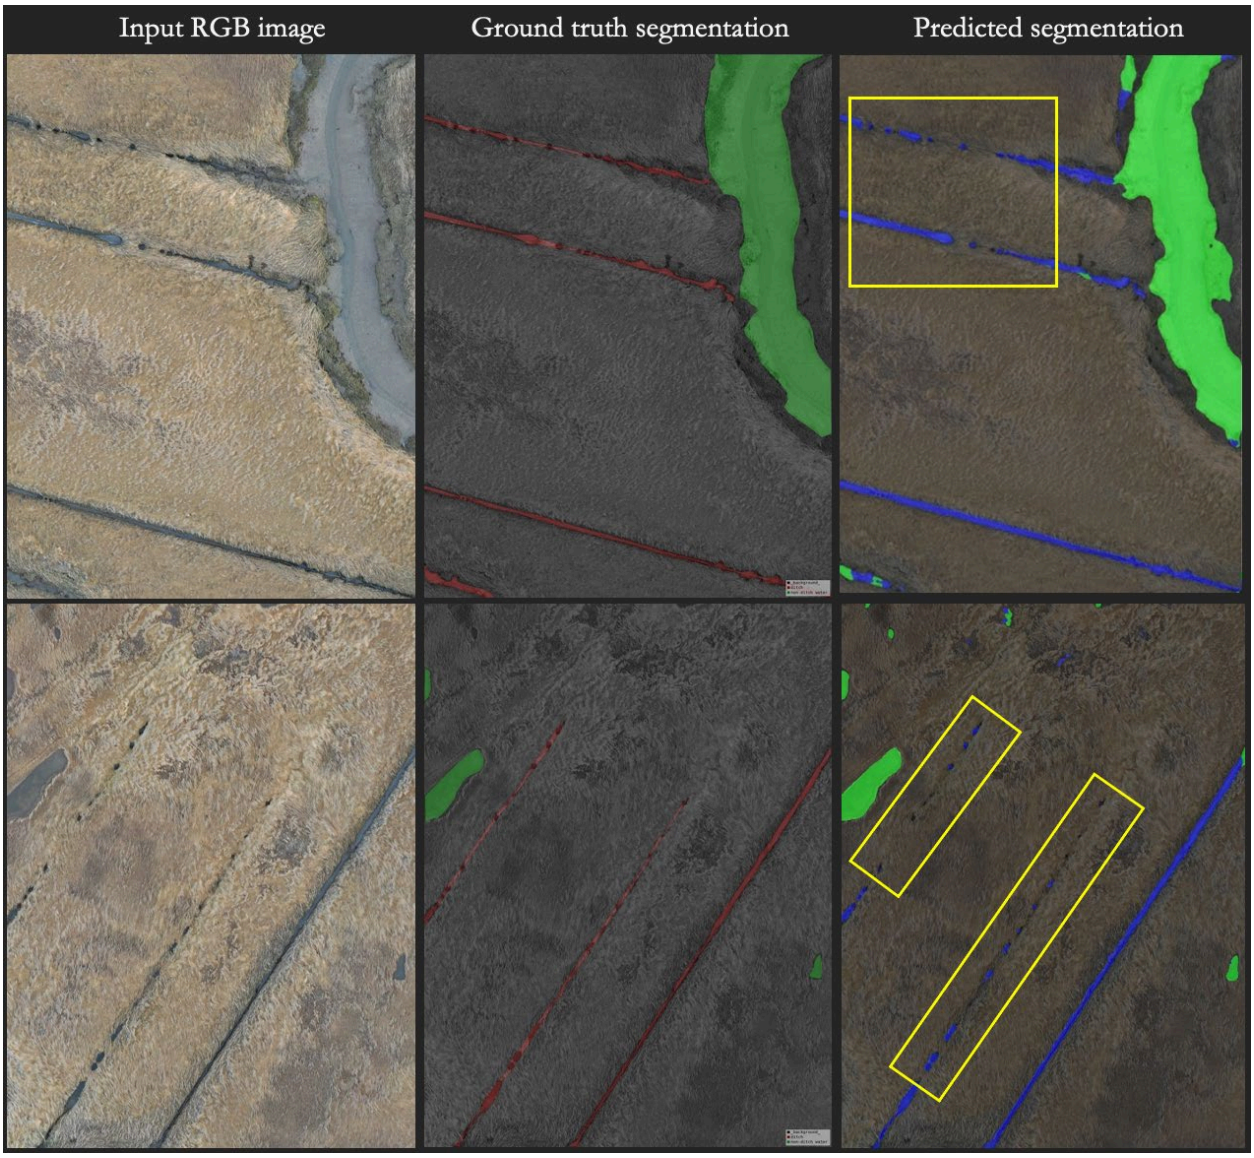

Figure 6. MobileNet-UNet results on a sample of the test images with disjoint ditches boxed in yellow. Ditches are shown in red in the ground truth segmentations and in blue in the predictions. Non-ditch water features are shown in green.

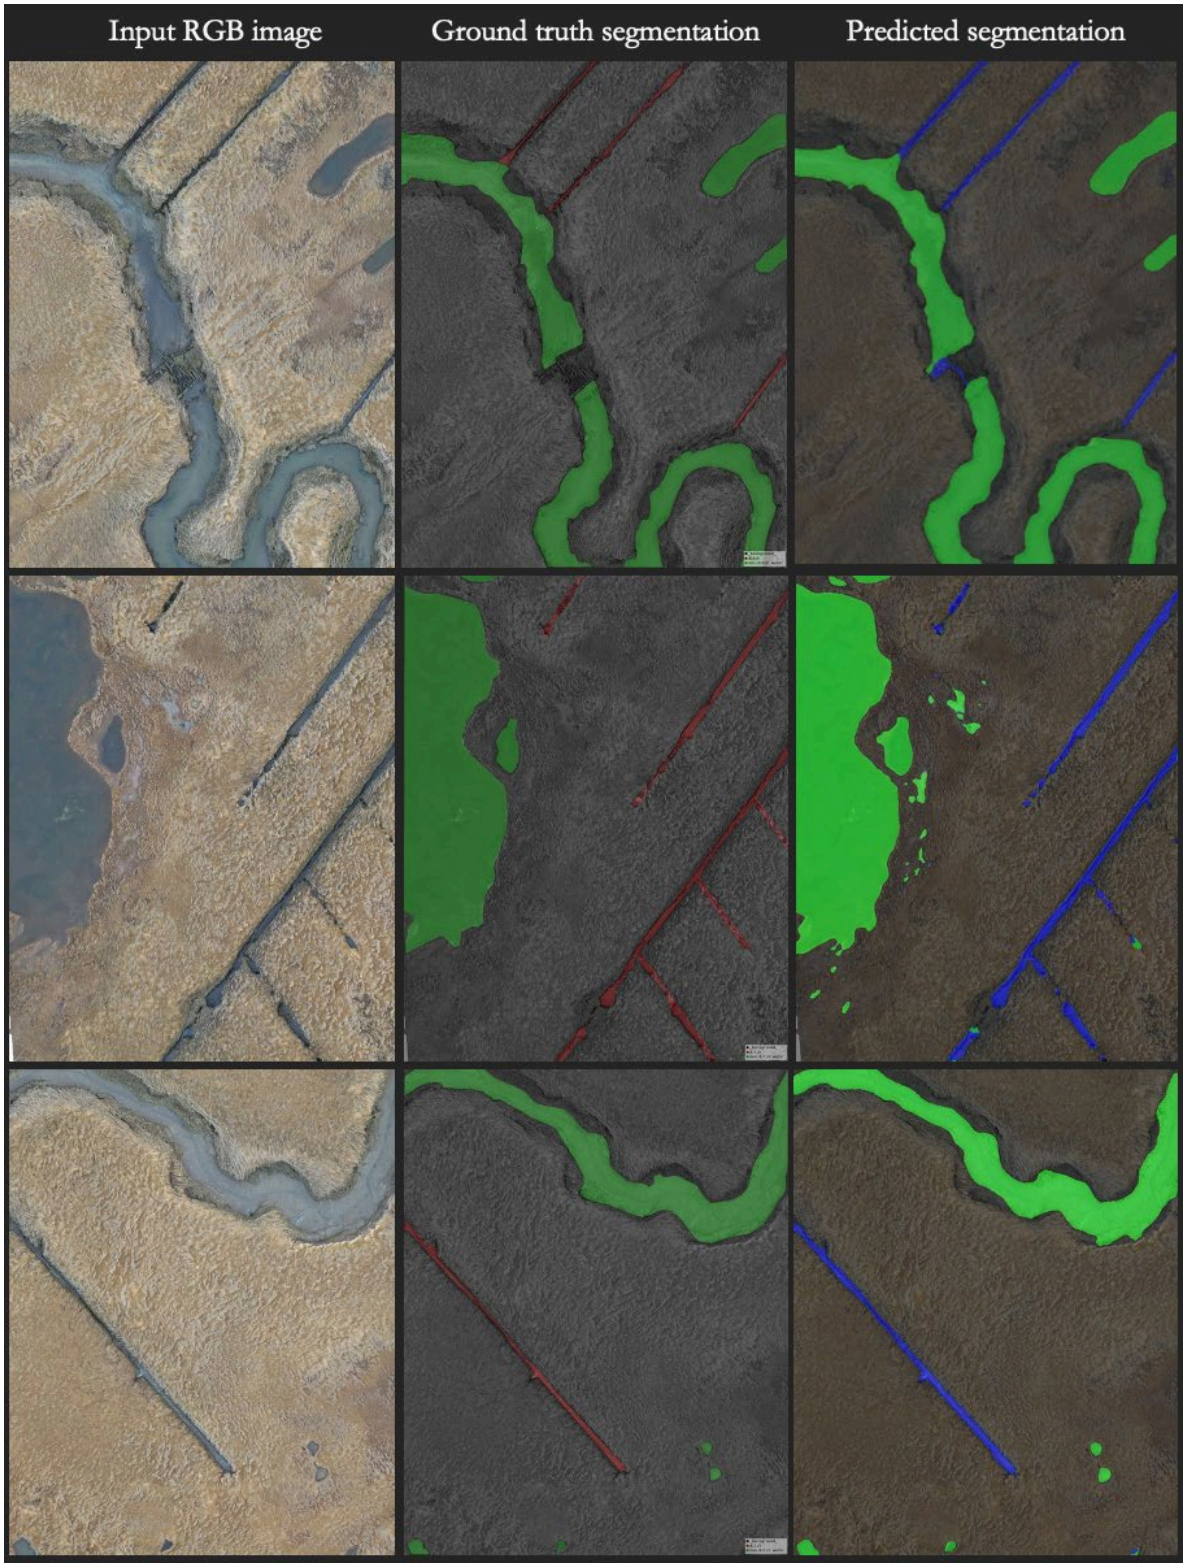

Figure 7. MobileNet-UNet results on a sample of the test images. Ditches are shown in red in the ground truth segmentations and in blue in the predictions. Non-ditch water features are shown in green.

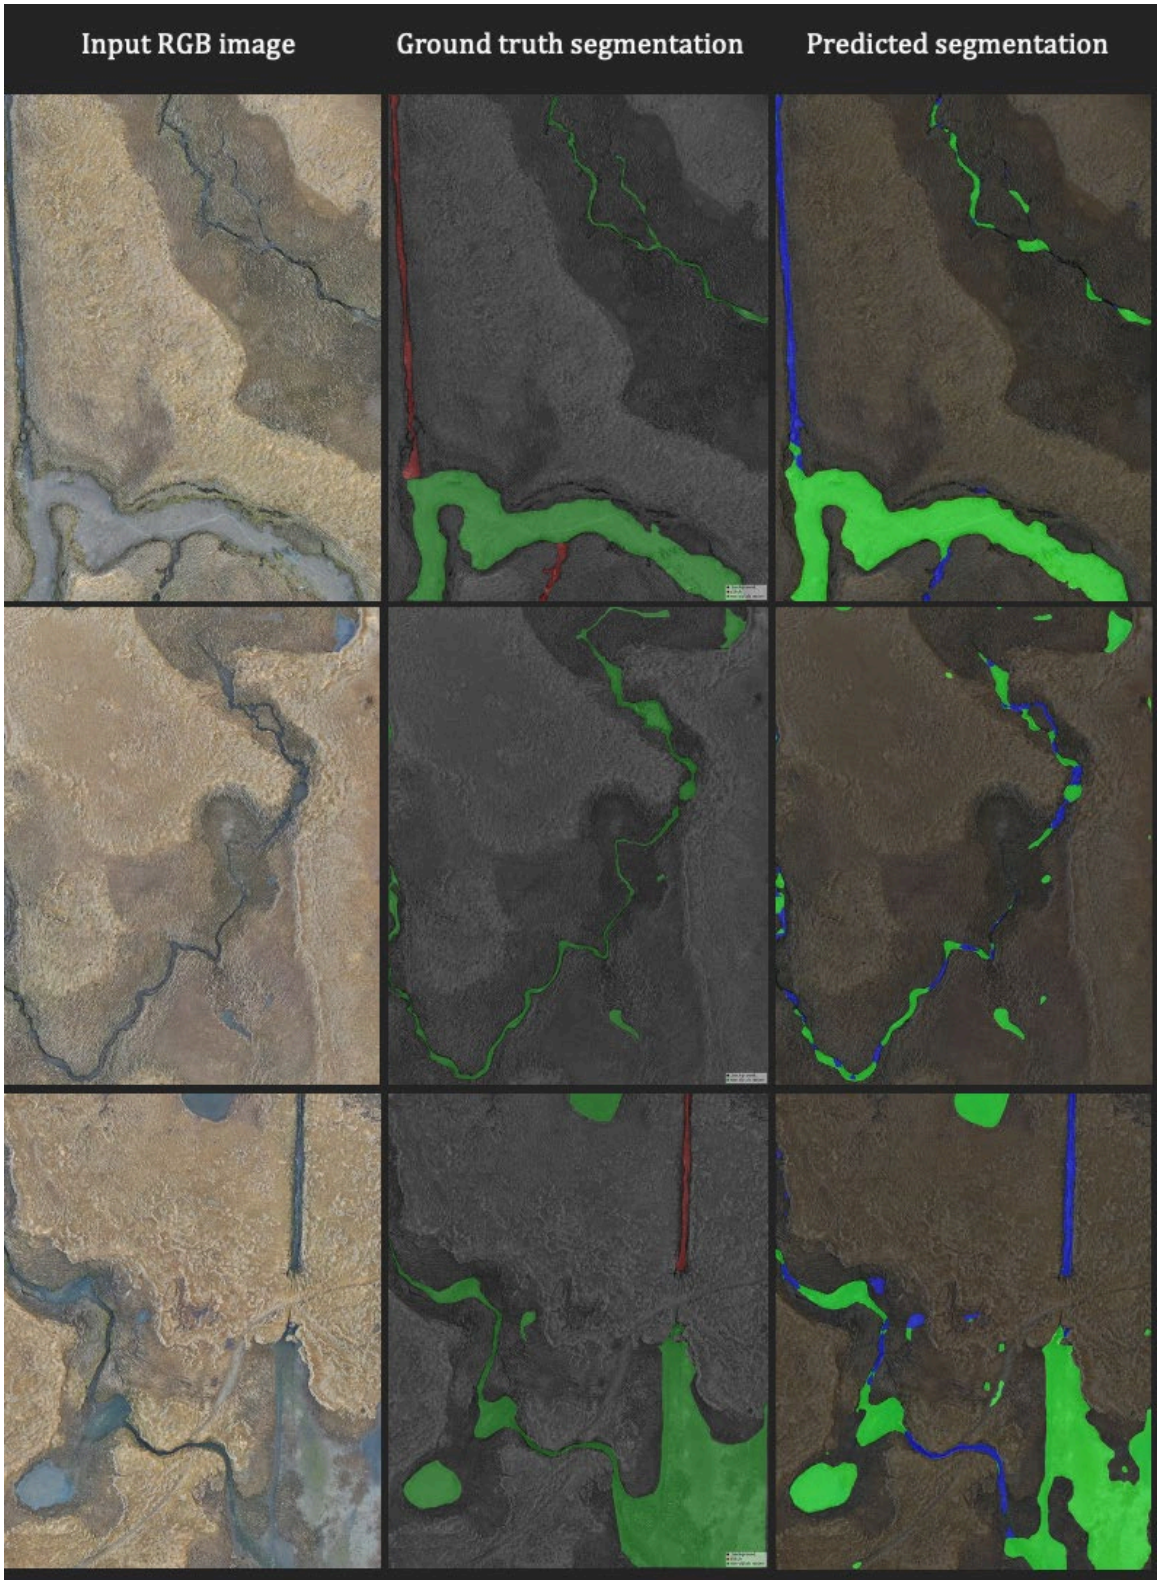

Figure 8. MobileNet-UNet results on a sample of the test images containing narrow creeks (non-ditch) that the model misclassifies as ditches. Ditches are shown in red in the ground truth segmentations and in blue in the predictions. Non-ditch water features are shown in green.

Conclusion and Future Works

This study demonstrates the effectiveness of U-Net-based architectures for semantic segmentation of salt marsh drone imagery. The analysis compared the performance of four different architectures, including the traditional U-Net and three U-Net variations with ImageNet pretrained encoders (VGG-16, ResNet50, and MobileNet). The models were trained using both Dice and Jaccard loss functions to investigate their influence on segmentation performance. The primary focus was on the ditch class to provide ecologists a more efficient method of delineating ditches.

The results indicate that the models with encoders pretrained on ImageNet generally outperformed the traditional U-Net. This highlights the benefits of transfer learning, even when fine-tuning models with larger images than the original pretrained weights were trained on. Among the models, MobileNet-UNet achieved the highest performance in terms of mean IoU and mean Dice scores, regardless of the loss function used.

The use of Jaccard loss led to improved results for all models. This suggests that Jaccard loss is better suited for achieving improved outcomes for the ditch class while maintaining reasonable performance for the other classes. Notably, MobileNet-UNet obtained the best overall performance for the ditch class when trained with Dice loss with an IoU of 0.45 and a Dice coefficient of 0.62. Some limitations were observed, such as the model's inability to identify ditches obscured by vegetation overhang and the misclassification of narrow non-ditch water features as ditches.

Future research could explore ways to address these challenges. The additional use of multispectral bands, like near-infrared (NIR) and shortwave infrared (SWIR), could help water features stand out against the background and make segmentation easier for the model. Similarly,

using images at high tide when water features are fuller might make them easier to segment. The training dataset of 160 images was relatively small, so labeling more images would allow the model more opportunities to learn complex features like narrow creeks. Including images from other salt marsh sites may also enhance its generality and robustness.

This study provides valuable insights into the potential of deep learning-based semantic segmentation for analyzing salt marsh drone imagery. To our knowledge, this is the first project to extract water features under more than one class from salt marshes using CNNs. The results emphasize the importance of considering various factors such as architecture, loss function, and transfer learning for optimal performance. The findings contribute to the growing field of AI in earth observation applications and encourage further collaboration between computer science and the earth sciences.

Bibliography

Akiyama, T. S., Junior, J. M., Gonçalves, W. N., Bressan, P. O., Eltner, A., Binder, F., & Singer, T. (2020). *DEEP LEARNING APPLIED TO WATER SEGMENTATION*. https://doi.org/10.5194/isprs-archives-XLIII-B2-2020-1189-2020.

Badrinarayanan, V., Kendall, A., & Cipolla, R. (2015). *SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation*. Retrieved January 3, 2022, from http://mi.eng.cam.ac.uk/projects/segnet/.

Bei Xue and Xiaohu Lin. (2020). Water System Segmentation Method of High Resolution Remote Sensing Image Based on eCognition. *Journal of Physics: Conference Series*, *1651*(1), 012162. https://doi.org/10.1088/1742-6596/1651/1/012162

J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 248-255. IEEE, 2009

D. Gupta. Image Segmentation Keras: Implementation of SegNet, FCN, U-Net and other models in Keras. GitHub, 2017. https://github.com/divamgupta/image-segmentation-keras.

Donatelli, C., Ganju, N. K., Zhang, X., Fagherazzi, S., & Leonardi, N. (2018). Salt Marsh Loss Affects Tides and the Sediment Budget in Shallow Bays. *Journal of Geophysical Research: Earth Surface*, *123*(10), 2647–2662. https://doi.org/10.1029/2018JF004617.

Enwright, N. M., Wang, L., Wang, H., Osland, M. J., Feher, L. C., Borchert, S. M., & Day, R. H. (2019). Evaluation of the Potential of Convolutional Neural Networks and Random Forests for Multi-Class Segmentation of Sentinel-2 Imagery. *Remote Sensing 2019, Vol. 11, Page 907*, *11*(8), 907. https://doi.org/10.3390/RS11080907.

Feyisa, G. L., Meilby, H., Fensholt, R., & Proud, S. R. (2014). Automated Water Extraction Index: A new technique for surface water mapping using Landsat imagery. *Remote Sensing of Environment*, *140*, 23–35. https://doi.org/10.1016/J.RSE.2013.08.029

He, Kaiming, et al. "Deep residual learning for image recognition." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770-778.

Howard, Andrew G., et al. "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications." arXiv preprint arXiv:1704.04861 (2017).

Jin, Y., Yang, W., Sun, T., Yang, Z., & Li, M. (2016). Effects of seashore reclamation activities on the health of wetland ecosystems: A case study in the Yellow River Delta, China. *Ocean & Coastal Management*, *123*, 44–52. https://doi.org/10.1016/J.OCECOAMAN.2016.01.013

K. Wada. LabelMe: Image Polygonal Annotation with Python. GitHub, 2022. https://github.com/wkentaro/labelme.

Kennish, M. J. (n.d.). *Coastal Salt Marsh Systems in the U.S.: A Review of Anthropogenic Impacts on JSTOR*. Retrieved May 27, 2022, from https://www.jstor.org/stable/4300224?casa_token=Amtxa4mWessAAAAA%3A1M3DBsvY_HJp2-oB9PJ2PDJGyAWay29hKmhvDlcB54Qep0PfSpfg5OSxsV9n1O5flKGneFR_-oHRyyTZMpsqi3VHQEZWaqVFAlSBjxqi2lEkpq25N0I&seq=1

Lockwood, B., & Drakeford, B. M. (2021). The value of carbon sequestration by saltmarsh in Chichester Harbour, United Kingdom. *Https://Doi.Org/10.1080/21606544.2020.1868345*, *10*(3), 278–292. https://doi.org/10.1080/21606544.2020.1868345

McFeeters, S. K. (2007). The use of the Normalized Difference Water Index (NDWI) in the delineation of open water features. *Https://Doi.Org/10.1080/01431169608948714*, *17*(7), 1425–1432. https://doi.org/10.1080/01431169608948714

Miao, Z., Fu, K., Sun, H., Sun, X., & Yan, M. (2018). Automatic Water-Body Segmentation from High-Resolution Satellite Images via Deep Networks. *IEEE Geoscience and Remote Sensing Letters*, *15*(4), 602–606. https://doi.org/10.1109/LGRS.2018.2794545.

O. Ronneberger, P. Fischer, and T. Brox. U-net: Convo-lutional networks for biomedical image segmentation. InInternational Conference on Medical image computing andcomputer-assisted intervention, pages 234–241. Springer,2015

K. Simonyan and A. Zisserman. Very deep convolutionalnetworks for large-scale image recognition.arXiv preprintarXiv:1409.1556, 201

Singh, A., Kalke, H., Loewen, M., & Ray, N. (2019). River Ice Segmentation with Deep Learning. *IEEE Transactions on Geoscience and Remote Sensing*, *58*(11), 7570–7579. https://doi.org/10.1109/TGRS.2020.2981082.

Sun, Z., Sandoval, L., Crystal-Ornelas, R., Mousavi, S. M., Wang, J., Lin, C., Cristea, N., Tong, D., Carande, W. H., Ma, X., Rao, Y., Bednar, J. A., Tan, A., Wang, J., Purushotham, S., Gill, T. E., Chastang, J., Howard, D., Holt, B., … John, A. (2022). A review of Earth Artificial Intelligence. *Computers & Geosciences*, *159*, 105034. https://doi.org/10.1016/J.CAGEO.2022.105034.

Tamta, K., Bhadauria, H. S., & Bhadauria, A. S. (n.d.). *Object-Oriented Approach of Information Extraction from High Resolution Satellite Imagery*. *17*(3), 47–52. https://doi.org/10.9790/0661-17344752

Vali, A.; Comai, S.; Matteucci, M. Deep Learning for Land Use and Land Cover Classification Based on Hyperspectral and Multispectral Earth Observation Data: A Review. Remote Sens. 2020, 12, 2495. https://doi.org/10.3390/rs12152495.

Watson, E. B., Wigand, C., Davey, E. W., Andrews, H. M., Bishop, J., & Raposa, K. B. (2017). Wetland Loss Patterns and Inundation-Productivity Relationships Prognosticate Widespread Salt

Marsh Loss for Southern New England. *Estuaries and Coasts*, *40*(3), 662–681. https://doi.org/10.1007/S12237-016-0069-1/FIGURES/9

Watzin, Mary C, and James G. Gosselink. *The Fragile Fringe: Coastal Wetlands of the Continental United States*. Baton Rouge: Louisiana Sea Grant College Program, 1992. Print.

Weng, L., Xu, Y., Xia, M., Zhang, Y., Liu, J., & Xu, Y. (2020). Water Areas Segmentation from Remote Sensing Images Using a Separable Residual SegNet Network. *ISPRS International Journal of Geo-Information 2020, Vol. 9, Page 256*, *9*(4), 256. https://doi.org/10.3390/IJGI9040256.

Weng, W., & Zhu, X. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. *IEEE Access*, *9*, 16591–16603. https://doi.org/10.48550/arxiv.1505.04597.

Xing, X., & Shen, J. (2018). Offshore Oil Slicks Extraction by Landsat Data Based on eCognition Software in South China Sea. *ICALIP 2018 - 6th International Conference on Audio, Language and Image Processing*, 144–147. https://doi.org/10.1109/ICALIP.2018.8455845

Xu, H. (2007). Modification of normalised difference water index (NDWI) to enhance open water features in remotely sensed imagery. *Https://Doi.Org/10.1080/01431160600589179*, *27*(14), 3025–3033. https://doi.org/10.1080/01431160600589179

Yang, X., Qin, Q., Grussenmeyer, P., & Koehl, M. (2018). Urban surface water body detection with suppressed built-up noise based on water indices from Sentinel-2 MSI imagery. *Remote Sensing of Environment*, *219*, 259–270. https://doi.org/10.1016/J.RSE.2018.09.016