Developing a CNN for automated detection of Carolina bays from publicly available LiDAR data

Mark Lundine[1] and Arthur Trembanis[1]

[1]University of Delaware

November 22, 2022

Abstract

For over a century, the enigmatic Carolina bays have captivated geologists and spurred contentious debate on their origin. These circular to ovate and shallow (median diameter of 222 m, median depth of 2.17 m, median area of 26,249 sq. m) depressions span the Atlantic Coastal Plain (ACP) from northern Florida to southern New Jersey, with total counts ranging between 10,000 and 500,000. Using 1 meter gridded, 1.7 km by 1.7 km LiDAR digital elevation models (DEMs) of Delaware as training images, a convolutional neural network (CNN) was trained to detect Carolina bays. With such a large population size and with such uncertainty around the actual population size, mapping the Carolina bays is a problem that requires an automated detection scheme. Manual detection of bays from LiDAR across the entire Atlantic Coastal Plain would be extremely time intensive and prone to human annotation errors. Using Faster R-CNN within the TensorFlow Python library, a network was trained on 978 LiDAR images for 24 hours (42,450 iterations) on an Intel Core i7-4790K CPU at 4.00 GHz. This network automatically detects bays from LiDAR images with a bounding box and a confidence level. These bounding boxes can then be used to subset and then analyze regions of the DEM for statistics on the bays' three-dimensional shape. Extending this algorithm to DEMs from other areas of the ACP will provide a better understanding of the bays' geographic distribution as well as any differences in morphology between different geographic regions. This method for detecting geomorphic features is a highly efficient process that will provide better means for mapping various types of abundant geomorphic features in the future.

Developing a CNN for automated detection of Carolina bays from publicly available LiDAR data

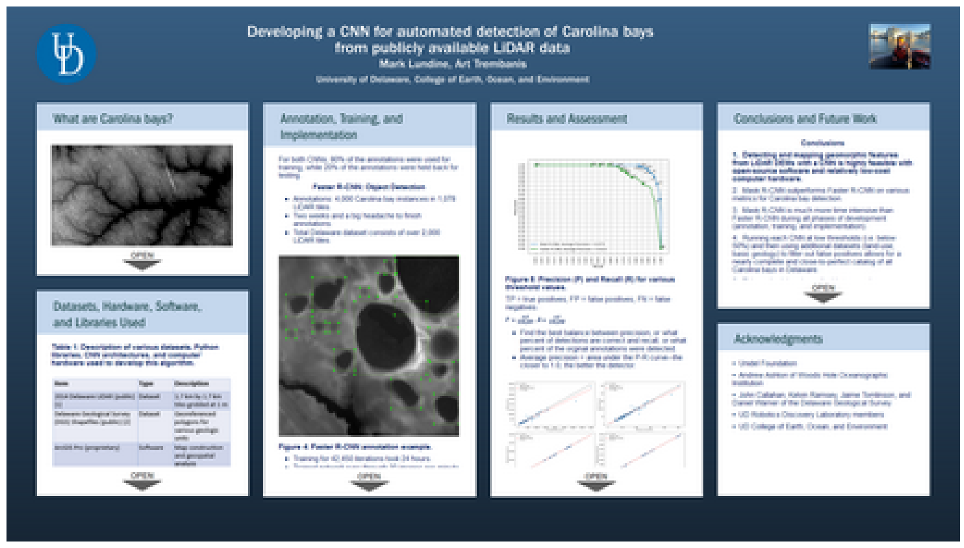

Mark Lundine, Art Trembanis

University of Delaware, College of Earth, Ocean, and Environment

PRESENTED AT:

WHAT ARE CAROLINA BAYS?

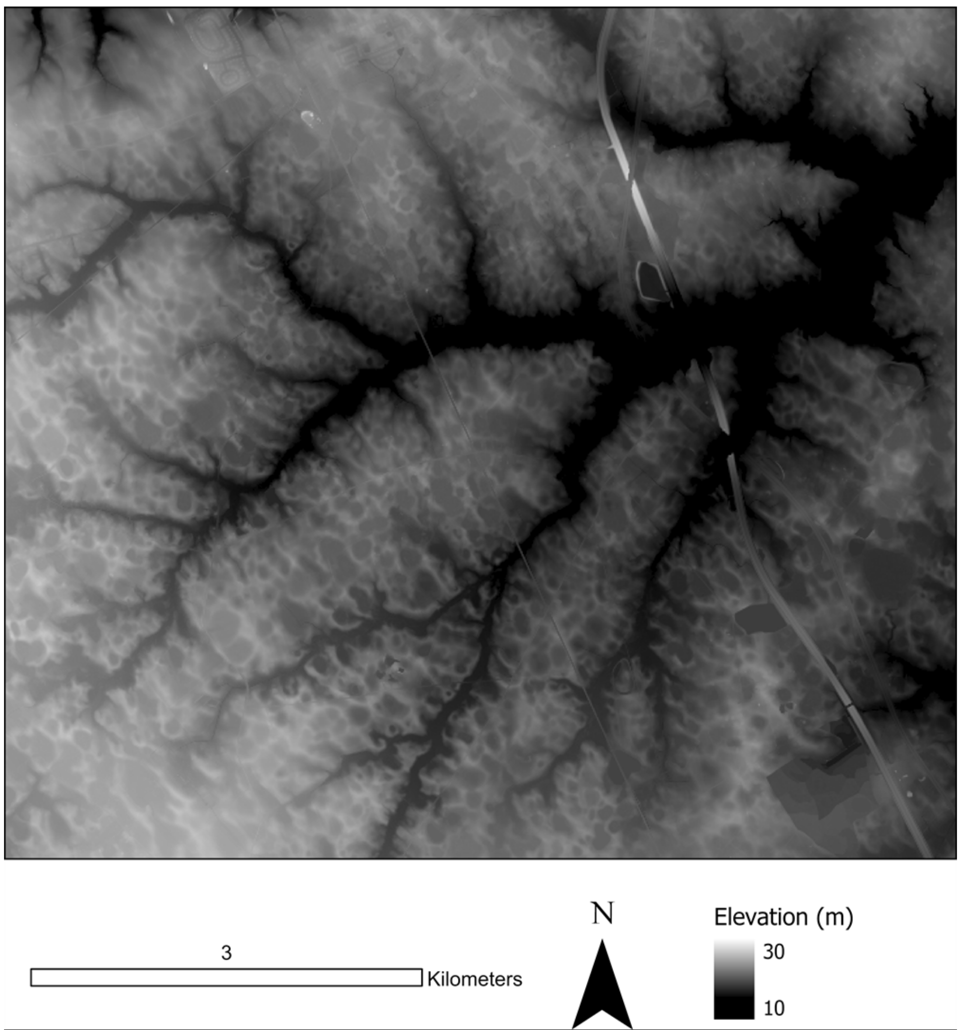

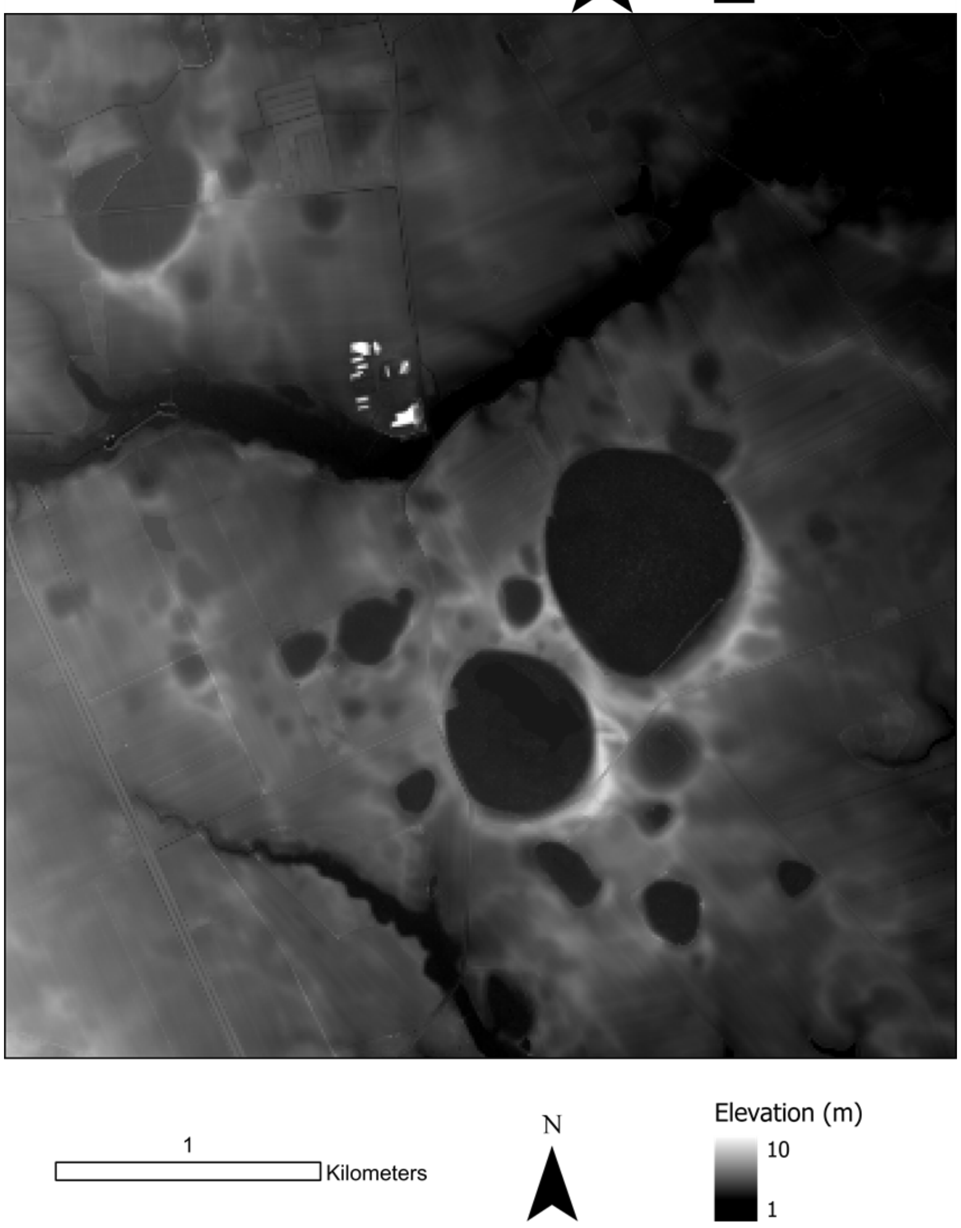

Figure 1: LiDAR digital elevation models (DEMs) of Carolina bays found in central (top) and southern (bottom) Delaware.

- Circular to ovate and shallow depressions with sandy rims, particularly on their southeastern corner.
- Morphologic characteristics: median diameter of 222 m, median depth of 2.17 m, median area of 26,249 m^2 or 2.6 hectares.
- Found throughout the Atlantic Coastal Plain
- Estimates of total count range from 10,000 to 500,000.
- **Too many to map manually.**
- **Here we present a convolutional neural network (CNN) detection scheme for Carolina bays.**

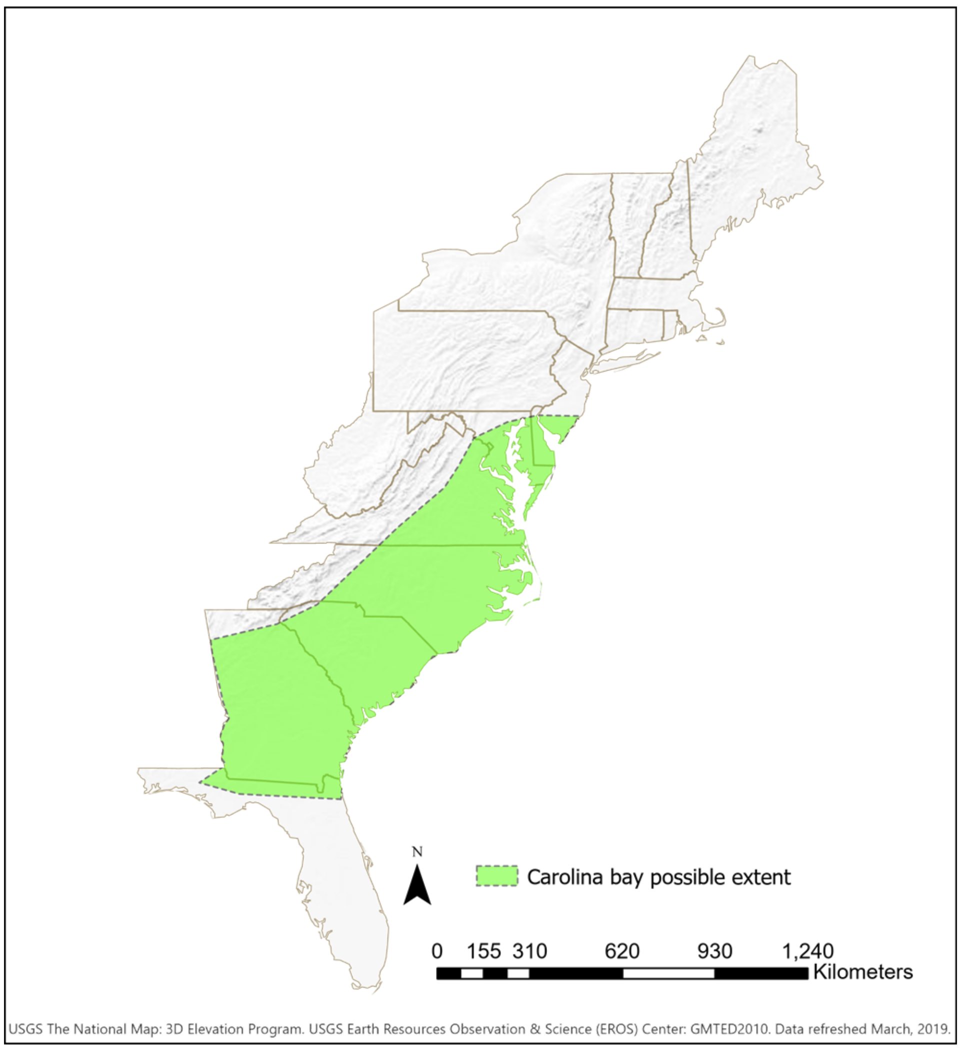

Figure 2: The approximate spatial extent of Carolina bays.

DATASETS, HARDWARE, SOFTWARE, AND LIBRARIES USED

Table 1: Description of various datasets, Python libraries, CNN architectures, and computer hardware used to develop this algorithm.

Item	Type	Description
2014 Delaware LiDAR (public) [1]	Dataset	1.7 km by 1.7 km tiles gridded at 1 m
Delaware Geological Survey (DGS) Shapefiles (public) [2]	Dataset	Georeferenced polygons for various geologic units
ArcGIS Pro (proprietary)	Software	Map construction and geospatial analysis
gdal (open source) [3]	Python Library	Geospatial file conversion
LabelImg (open source) [4]	Software	Annotation
TensorFlow (open source) [5]	Python Library	Machine/deep learning tasks
Faster R-CNN (open source) [6]	CNN architecture	Object detection framework
Mask R-CNN (open source) [7]	CNN architecture	Instance segmentation framework
Intel Core i7-4790K CPU at 4.00 GHz	Hardware	Computer used for training and implementation

Using a more sophisiticated hardware setup with a GPU would expedite the training and implementation process. Also, open-source alternatives to ArcGIS Pro exist and have many of the same capabilitites as this proprietary software.

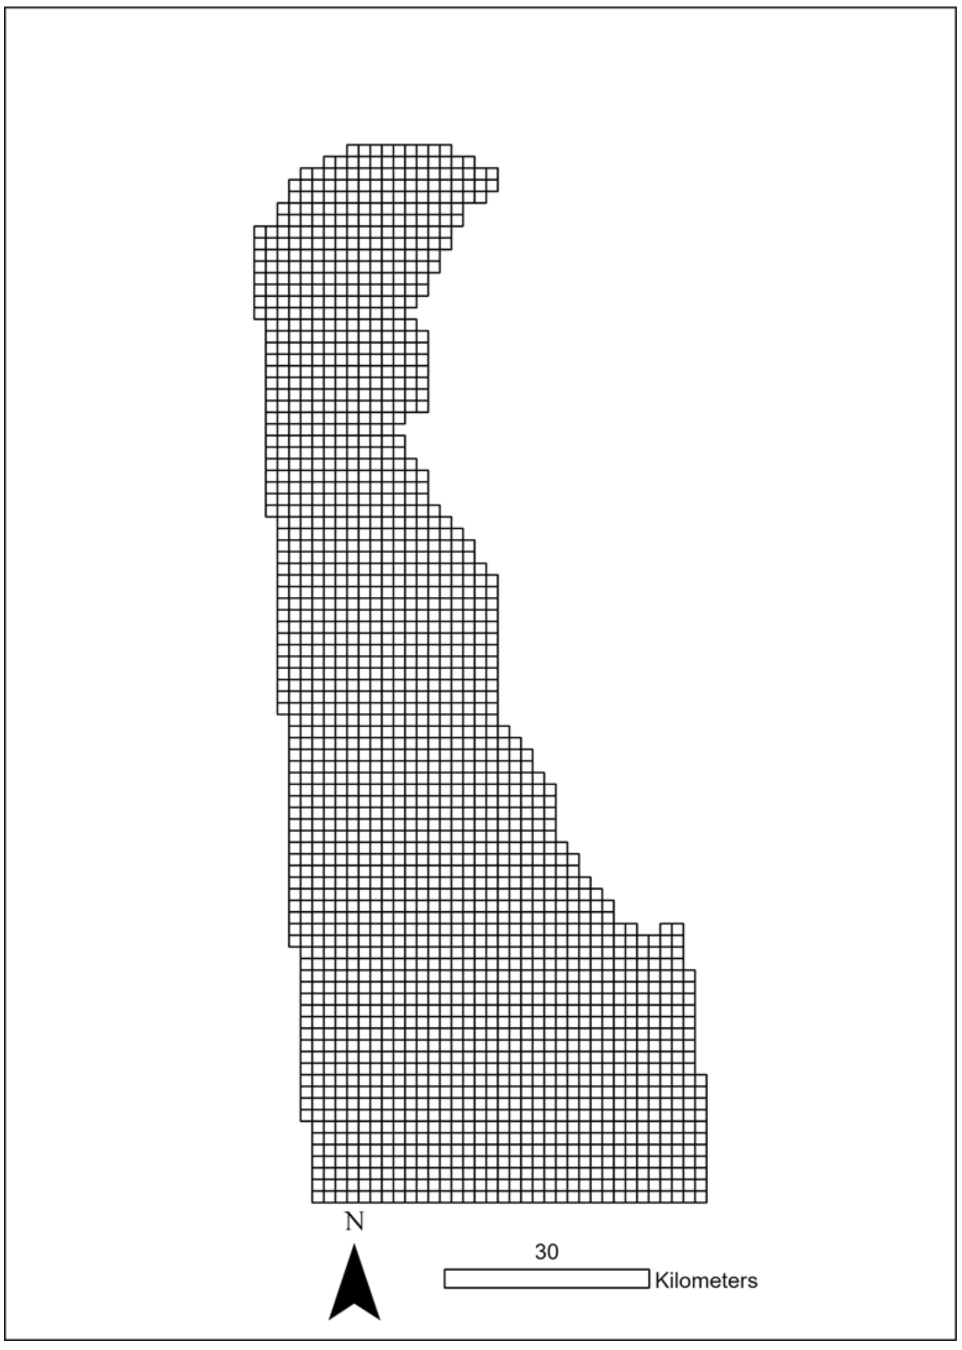

Figure 3: Delaware tile layout. Each box represents the footprint of the LiDAR images used in this study.

ANNOTATION, TRAINING, AND IMPLEMENTATION

For both CNNs, 80% of the annotations were used for training, while 20% of the annotations were held back for testing.

Faster R-CNN: Object Detection

- Annotations: 4,000 Carolina bay instances in 1,078 LiDAR tiles.
- Two weeks and a big headache to finish annotations.
- Total Delaware dataset consists of over 2,000 LiDAR tiles.

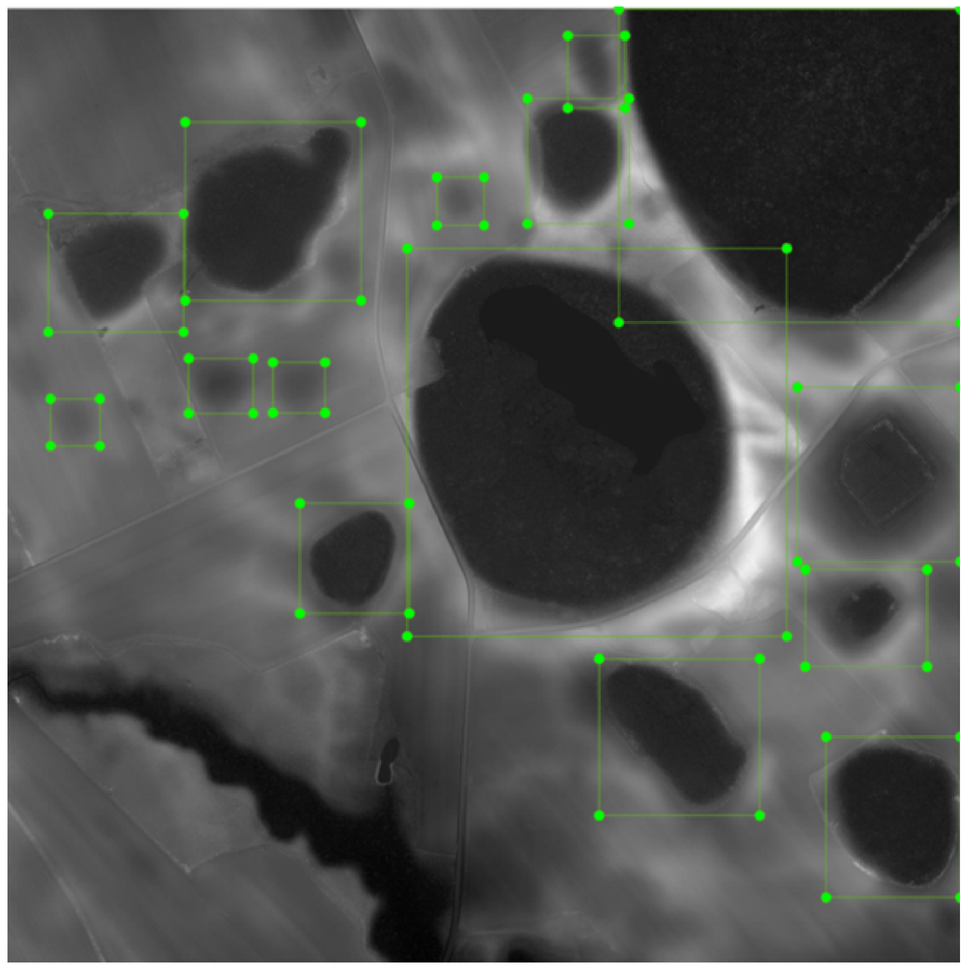

Figure 4: Faster R-CNN annotation example.

- Training for 42,450 iterations took 24 hours.
- Trained network runs through 30 images per minute, one hour for entire dateset.

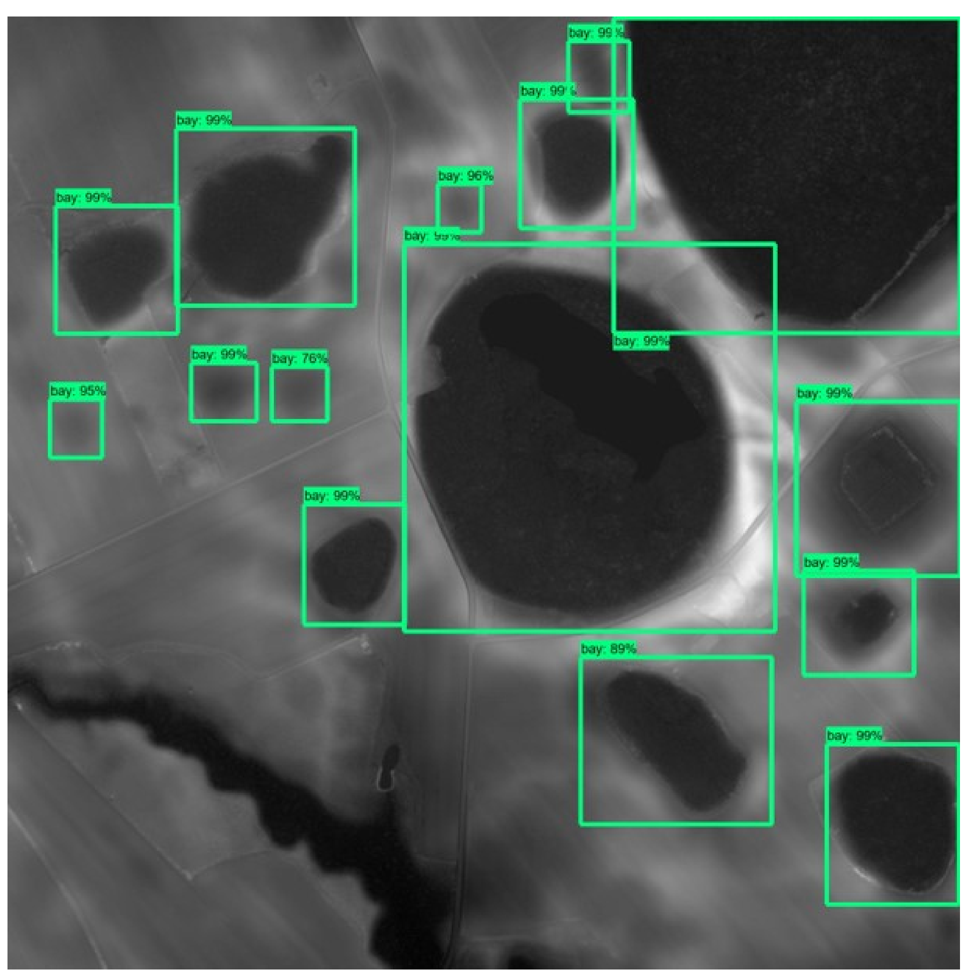

Figure 5: Implementation example for trained Faster R-CNN.

Mask R-CNN: Instance Segmenation

- Bounding boxes are nice, but a real fit would be better.
- Convert Delaware Geological Survey's shapefiles of bays into binary masks.
- Annotate the masks with bounding boxes.
- Final annotation set conists of bounding box annotations, binary maks, and original DEMs.
- Over 1,000 bay instances in 305 images.

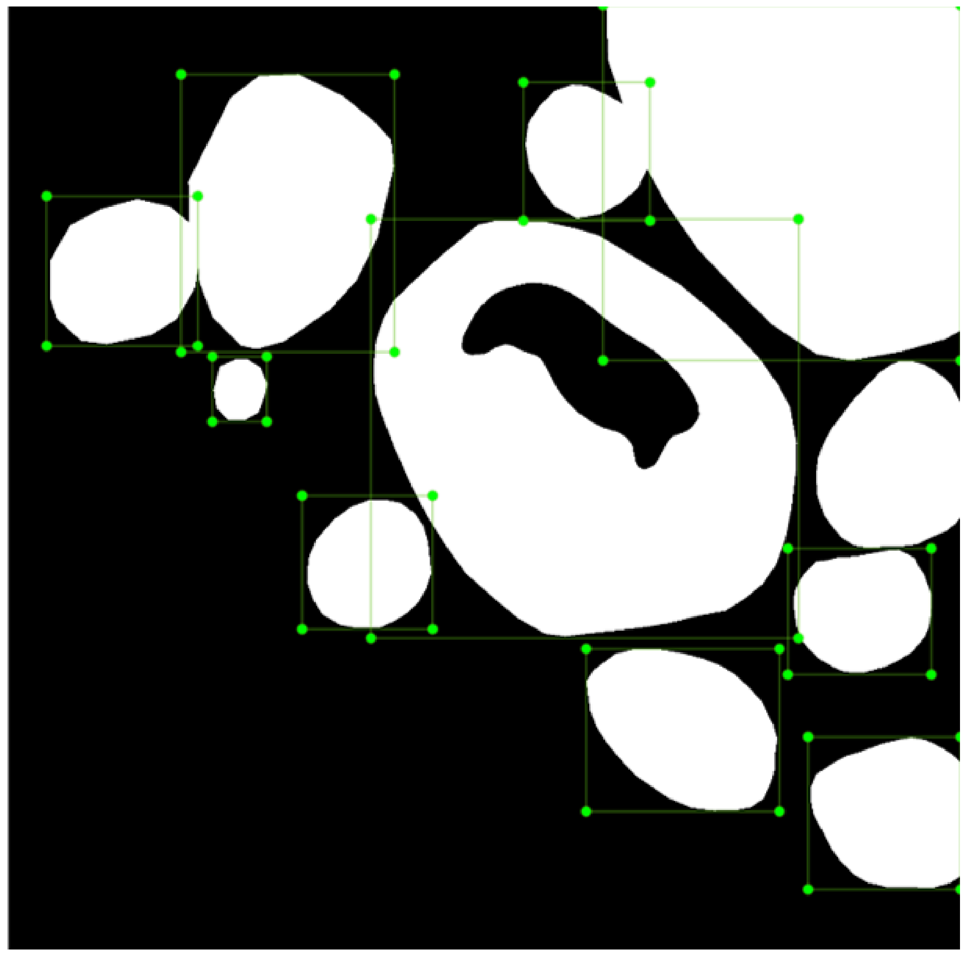

Figure 6: Mask R-CNN annotation example.

- Training for 40,479 iterations took two weeks.
- Trained network goes through two images per minute.
- Almost 17 hours to run through the whole state.
- Having an exact fit instead of a box makes it worth the wait.

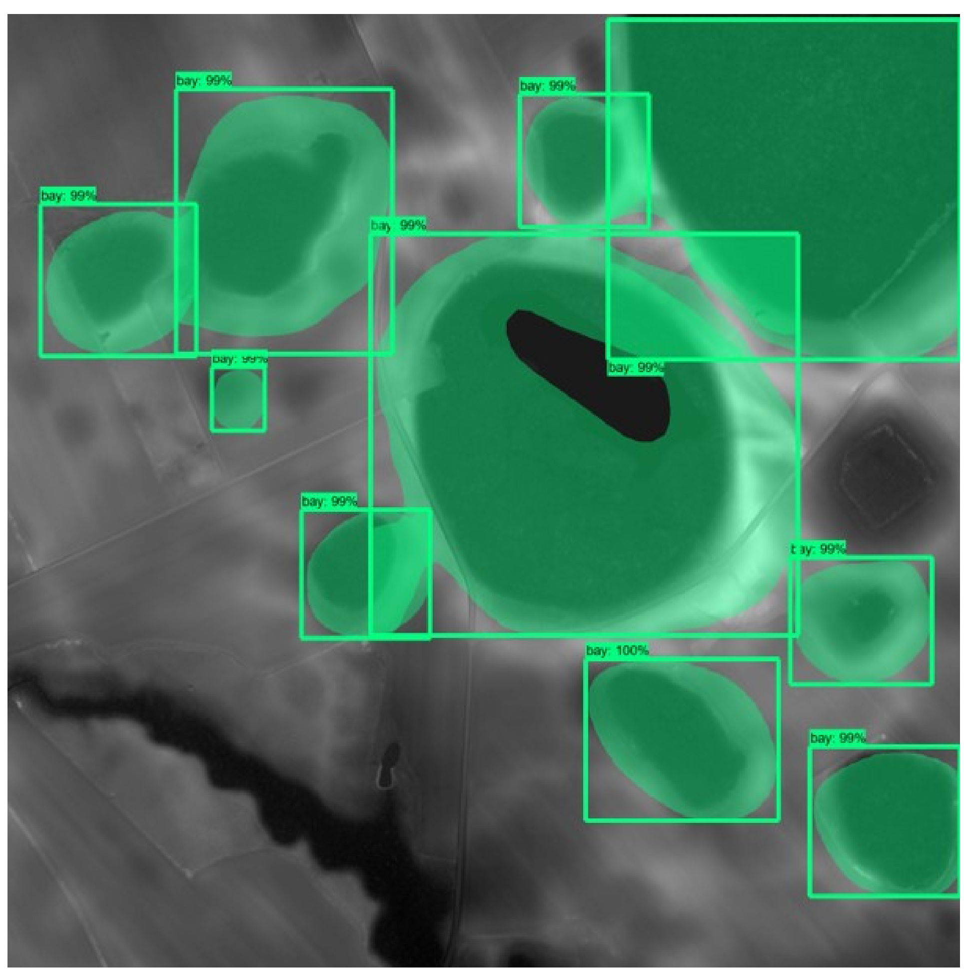

Figure 7: Implentation example for trained Mask R-CNN.

RESULTS AND ASSESSMENT

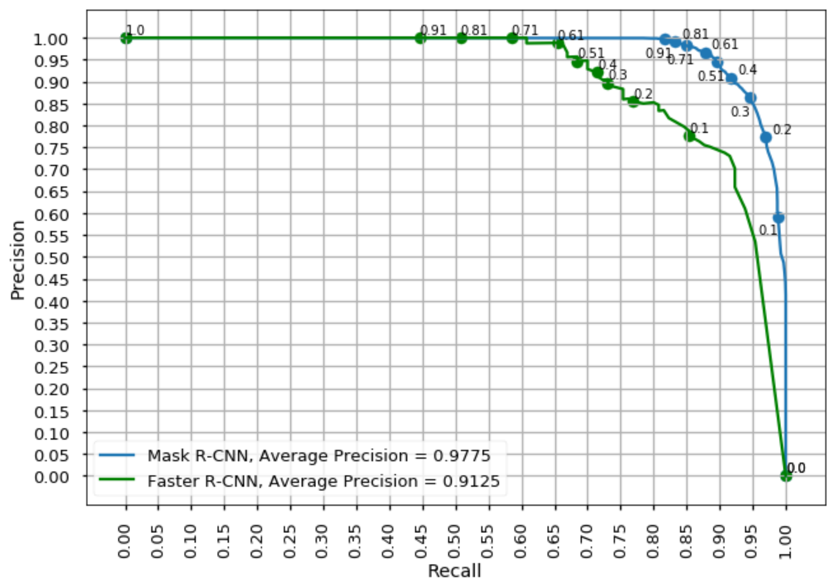

Figure 8: Precision (P) and Recall (R) for various threshold values.

TP = true positives, FP = false positives, FN = false negatives.

$$P = \frac{TP}{TP+FP}, R = \frac{TP}{TP+FN}$$

- Find the best balance between precision, or what percent of detections are correct and recall, or what percent of the orginal annotations were detected.
- Average precision = area under the P-R curve--the closer to 1.0, the better the detector.

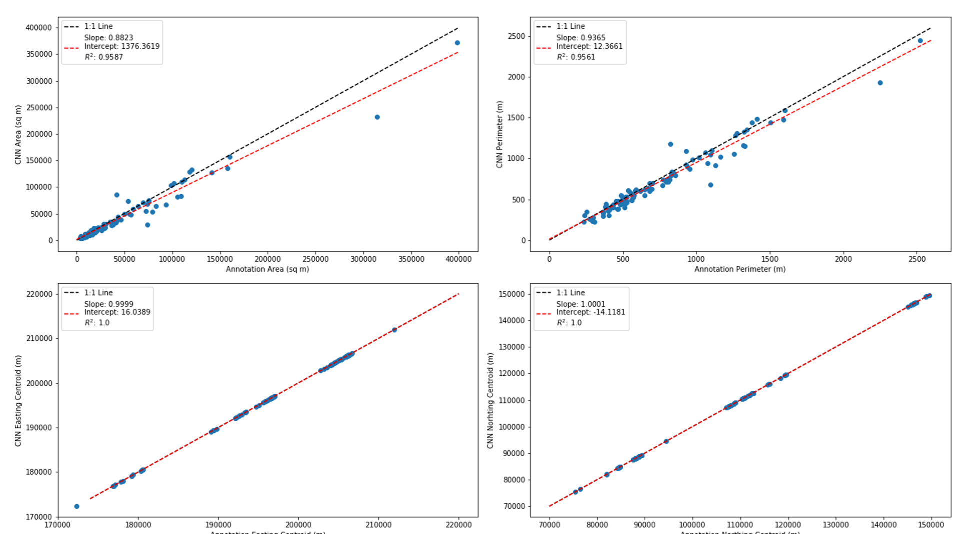

Figure 9: One to one comparisons of area, perimeter, x-centroid, and y-centroid between the Faster R-CNN detected bays and the annotated bays from the test dataset.

- Location and size of the detection bounding box should be assessed as well.
- Faster R-CNN centers the boxes correctly, but fits a tighter box than the annotations.

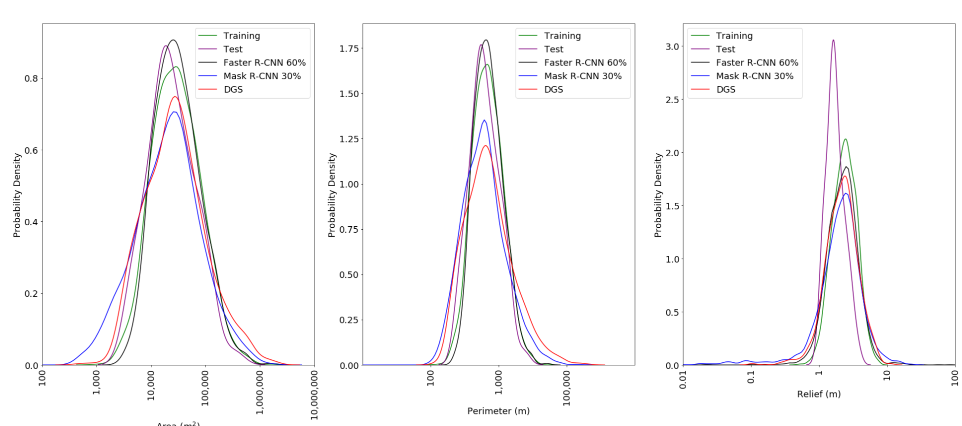

Figure 10: Comparing various detection and annotation datasets.

- Look at morphologic parameters from various datasets and compare.
- Pretty good agreement between all annotation and detection datasets.
- DGS dataset and Mask R-CNN detection datasets are expected to have more spread, some of the annotations lump many bays into one feature.

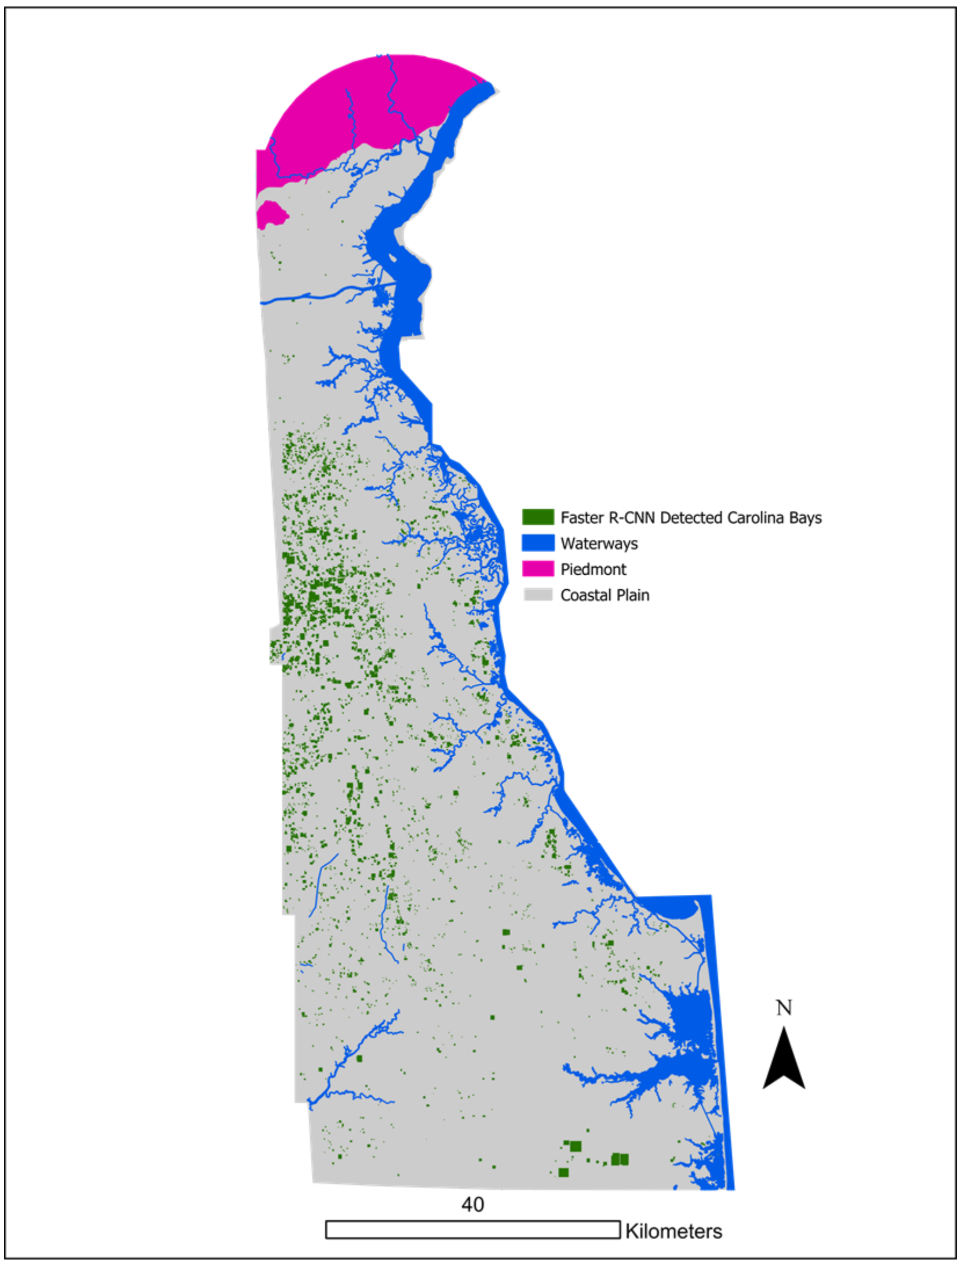

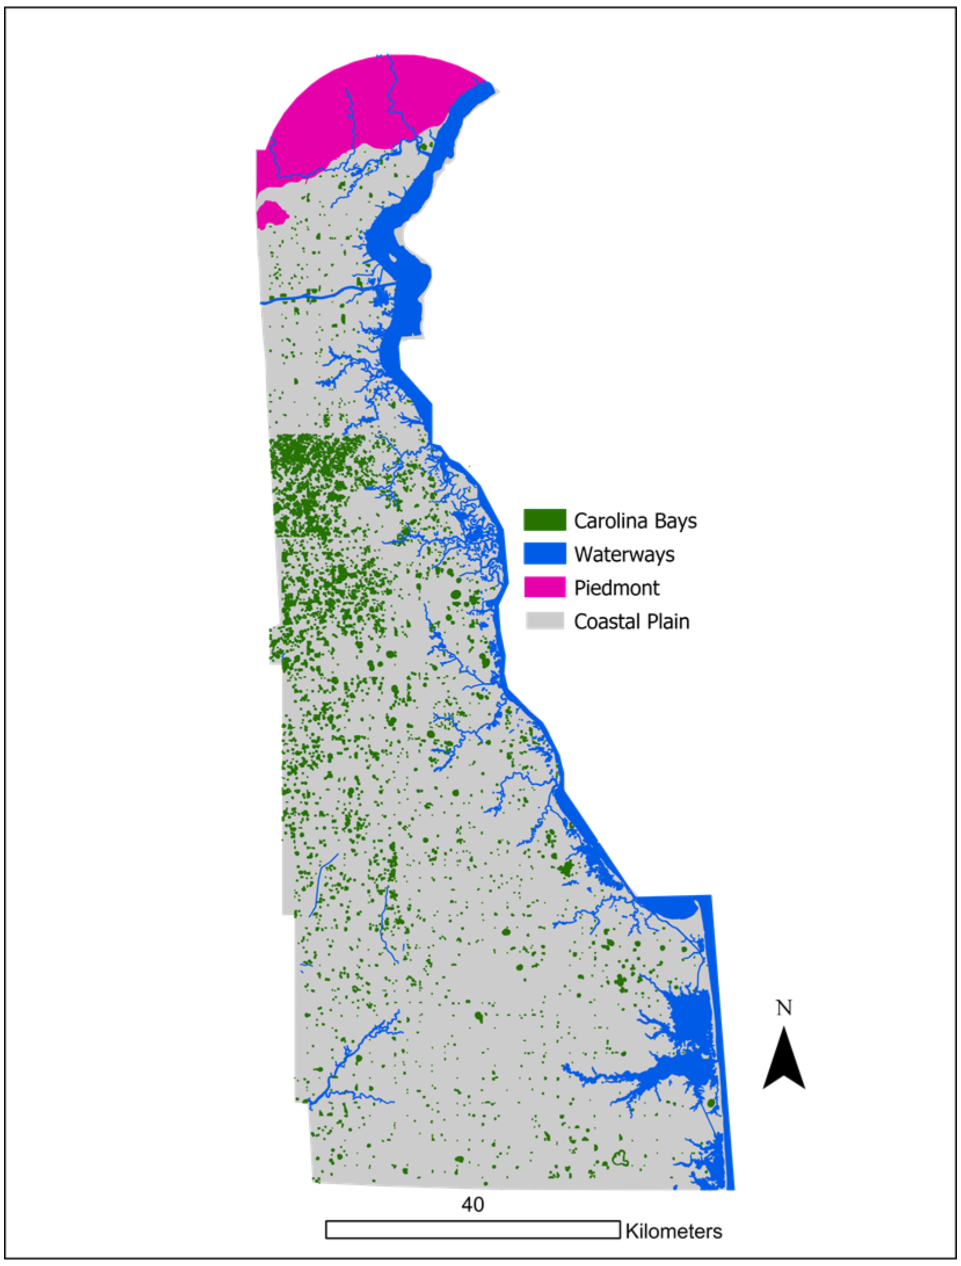

Figure 11: Final maps of Delaware showing Carolina bays, waterways, and general geology.

- The maps shown in Fig. 11 show detections from Faster R-CNN at 60% (top) and Mask R-CNN at 30% (bottom). These detections were then filtered for false positives using land-use data and basic geology.
- For Faster R-CNN, over 4,000 boxes!
- For Mask R-CNN, over 3,000 polygons! Just in Delaware. And we probably missed a few.
- Compare Fig. 11 with Fig. 12, what was already mapped in Delaware.

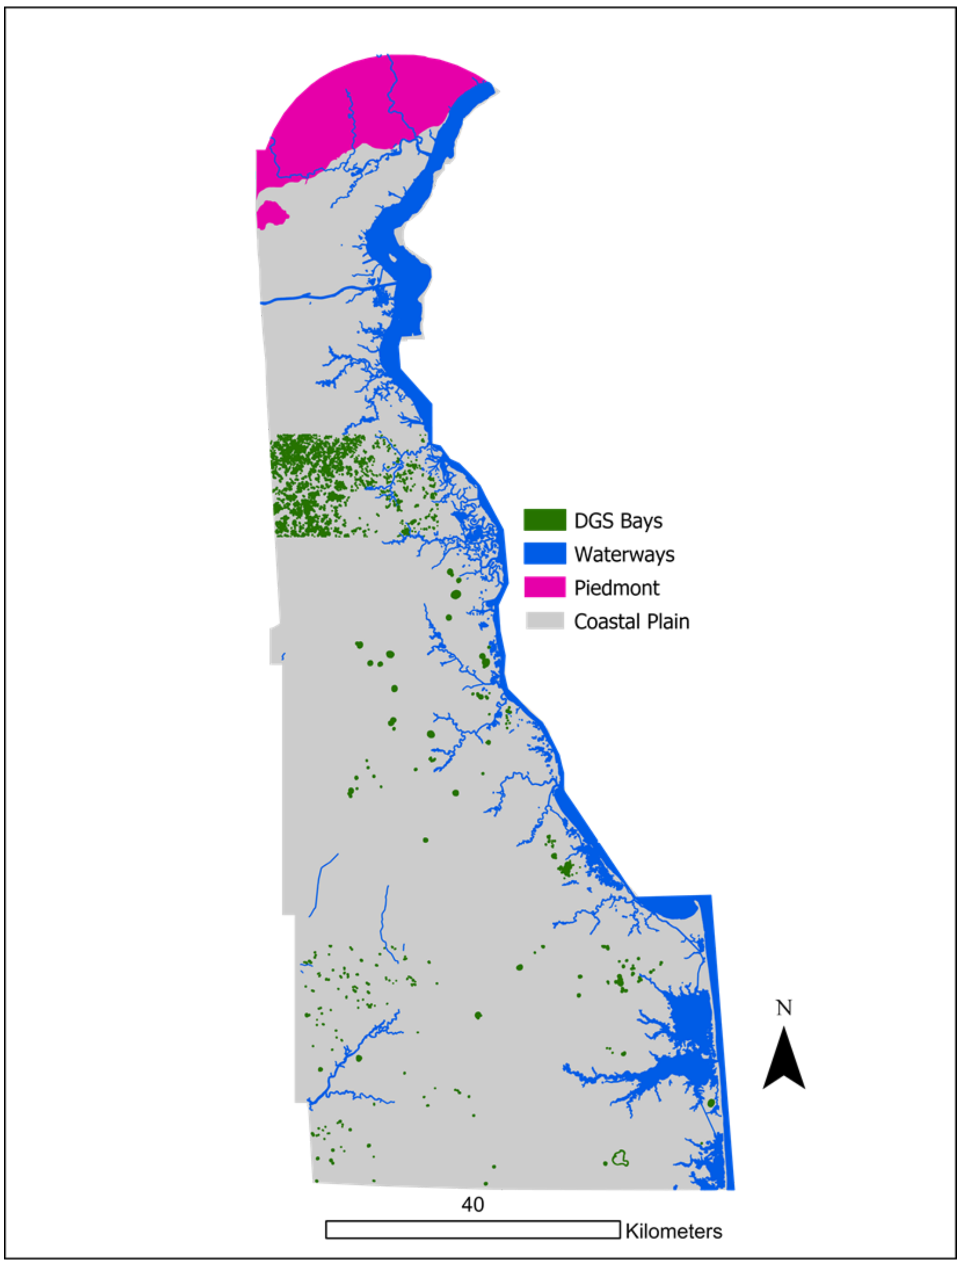

Figure 12: Carolina bays that were already mapped by DGS.

- Need to build large datasets to truly describe a geomorphic feature.
- The uncertainty in variance in various morphologic estimates diminish with larger sample sizes.
- Fig. 13 shows Carolina bay area interquartile range converging to a specific value as the sample size increases. This simulation was run thirty times.

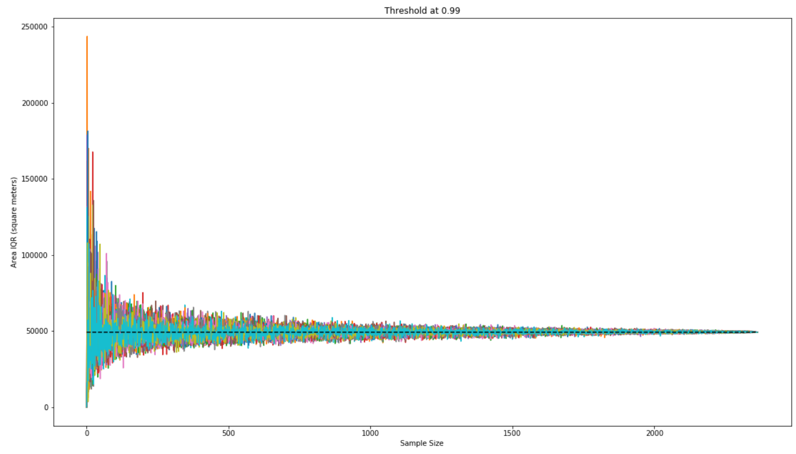

Figure 13: As more Carolina bays are mapped, the variance estimate (IQR) for their area converges.

CONCLUSIONS AND FUTURE WORK

Conclusions

1. Detecting and mapping geomorphic features from LiDAR DEMs with a CNN is highly feasible with open-source software and relatively low-cost computer hardware.

2. Mask R-CNN outperforms Faster R-CNN on various metrics for Carolina bay detection.

3. Mask R-CNN is much more time intensive than Faster R-CNN during all phases of development (annotation, training, and implementation).

4. Running each CNN at low thresholds (i.e. below 50%) and then using additional datasets (land-use, basic geology) to filter out false positives allows for a nearly complete and close-to-perfect catalog of all Carolina bays in Delaware.

5. This methodology is applicable to mapping many other types of abundant geomorphic features.

Future Work

1. Apply both models to LiDAR and similar digital elevation datasets outside of Delaware.

2. Build annotation set for detection of relict parabolic dunes (ex. shown below).

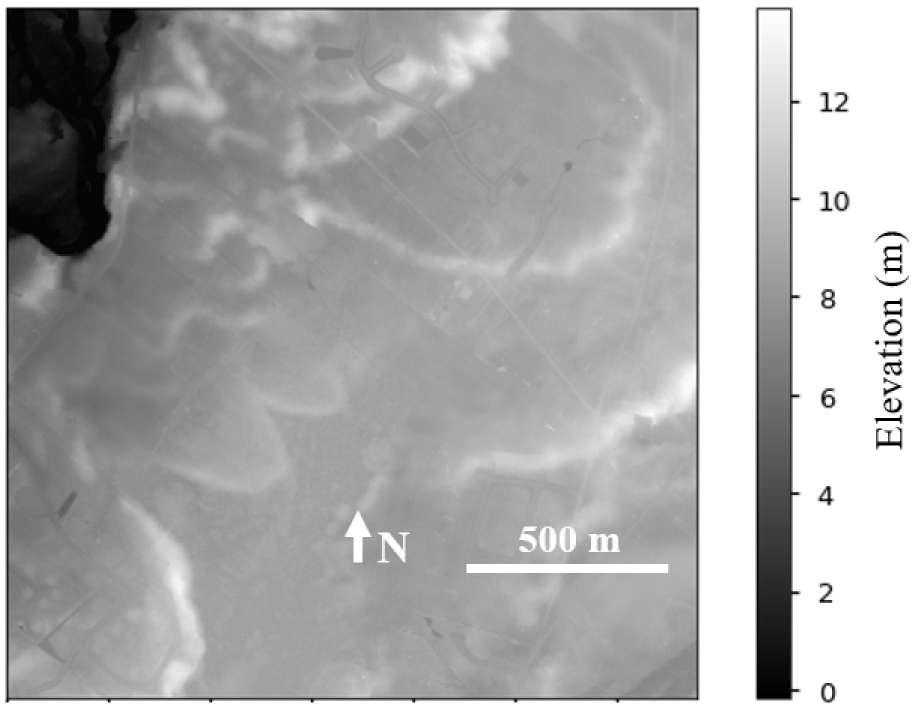

3. Investigate geological, ecological, and paleoclimatic conditions controlling bay spatial distribution and morphology.

4. What is the lowest resolution DEM that can be used to detect bays from? Can we input other topographic metrics (slope, valley depth, etc.) as training data in a multiband image?

ACKNOWLEDGMENTS

- Unidel Foundation
- Andrew Ashton of Woods Hole Oceanographic Institution
- John Callahan, Kelvin Ramsey, Jaime Tomlinson, and Daniel Warner of the Delaware Geological Survey
- UD Robotics Discovery Laboratory members
- UD College of Earth, Ocean, and Environment

Sorry but time is up!

ABSTRACT

For over a century, the enigmatic Carolina bays have captivated geologists and spurred contentious debate on their origin. These circular to ovate and shallow (median diameter of 222 m, median depth of 2.17 m, median area of 26,249 m2) depressions span the Atlantic Coastal Plain (ACP) from northern Florida to southern New Jersey, with total counts ranging between 10,000 and 500,000. Using 1 meter gridded, 1.7 km by 1.7 km LiDAR digital elevation models (DEMs) of Delaware as training images, a convolutional neural network (CNN) was trained to detect Carolina bays. With such a large population size and with such uncertainty around the actual population size, mapping the Carolina bays is a problem that requires an automated detection scheme. Manual detection of bays from LiDAR across the entire Atlantic Coastal Plain would be extremely time intensive and prone to human annotation errors. Using Faster R-CNN within the TensorFlow Python library, a network was trained on 978 LiDAR images for 24 hours (42,450 iterations) on an Intel Core i7-4790K CPU at 4.00 GHz. This network automatically detects bays from LiDAR images with a bounding box and a confidence level. These bounding boxes can then be used to subset and then analyze regions of the DEM for statistics on the bays' three-dimensional shape. Extending this algorithm to DEMs from other areas of the ACP will provide a better understanding of the bays' geographic distribution as well as any differences in morphology between different geographic regions. This method for detecting geomorphic features is a highly efficient process that will provide better means for mapping various types of abundant geomorphic features in the future.

REFERENCES

[1] OCM Partners, 2019: 2014 USGS CMGP Lidar: Sandy Restoration (Delaware and Maryland), https://inport.nmfs.noaa.gov/inport/item/49662.

[2] Delaware Geological Survey, 2019: Delaware Geologic Map Nos. 8,9,10,11,12,13,14,15,16,17,18,19,21,22,23,24, https://www.dgs.udel.edu/data.

[3] Geospatial Data Abstraction Library, 2019: https://pypi.org/project/GDAL/.

[4] Tzutalin, 2015: LabelImg, https://github.com/tzutalin/labelImg.

[5] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Rafal Jozefowicz, Yangqing Jia, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Mike Schuster, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng, 2015: TensorFlow: Large-scale machine learning on heterogeneous systems, tensorflow.org.

[6] Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun, 2015: Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks, https://arxiv.org/pdf/1506.01497v3.

[7] Kaiming He, Georgia Gkioxari, Piotr Dollár, Ross Girshick, 2018: Mask R-CNN, https://arxiv.org/pdf/1703.06870.pdf.

SWITCH TEMPLATE