4-14-2023

Classification of Land Cover on Sand Dunes

Heleyna Tucker
Hope College

Micah Sterk
Hope College

Hope COLLEGE

Classification of Land Cover on Sand Dunes

Heleyna Tucker and Micah Sterk
Faculty Supervisors: Darin Stephenson and Brian Yurk

INTRODUCTION

Understanding sand dune dynamics is important for both scientific knowledge and land management. In recent years, aerial drones have led to efficient and precise ways of gathering data and modeling sand dune shape and movement.

This project has the following objectives:

- Develop machine learning algorithms to create and enhance digital elevation models and land cover classifications of local dune regions based on remotely-sensed imagery.
- Use machine learning and mathematical processes to automatically align images taken using different cameras at different elevations and angles.
- Explore how image alignment can enhance machine learning classification processes and streamline protocols for more effective data collection in the field.

LOCATIONS

The Hope College Dune Group studies dune complexes at three West Michigan Locations.

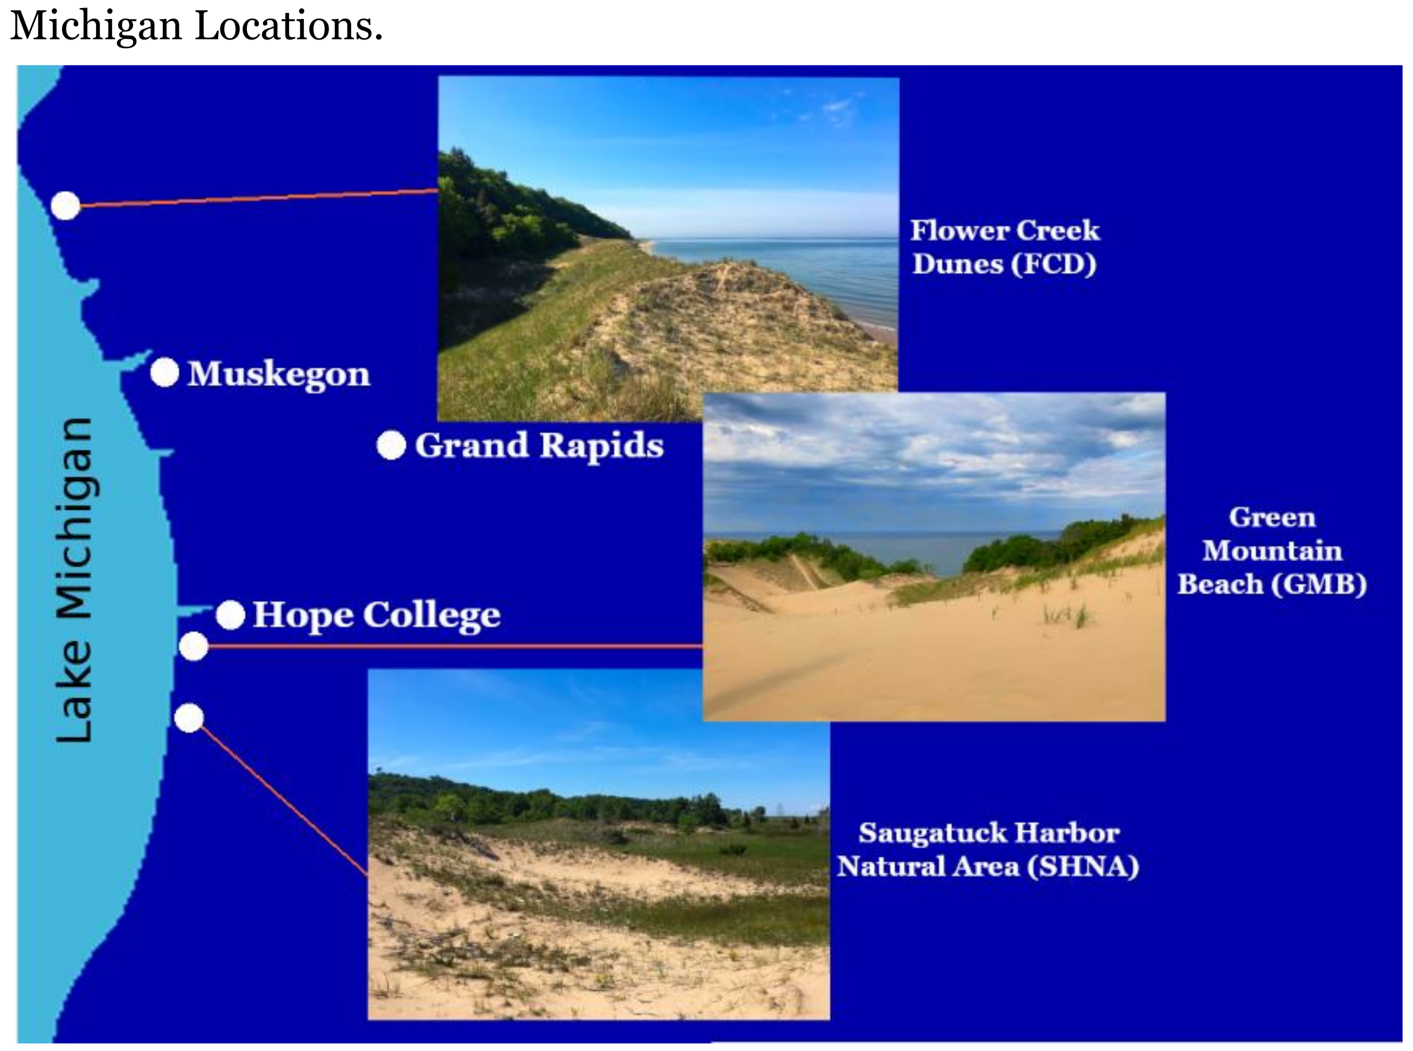

METHODS

Data is collected at these locations using both visible light and multi-spectral cameras attached to drones.

- A flight over an entire dune complex at a height of 120 meters can result in a point cloud of data consisting of over 100 million points. Higher resolution photos taken from lower altitudes can provide more precise data for training models that can be used for prediction across entire complexes.
- Data can be stitched and orthorectified to produce raster datasets containing precise position information and multi-spectral data including visible light, near infrared, red edge, and thermal readings.
- Machine learning algorithms can be implemented in R and Python to model and make predictions from this data.

LAND COVER CLASSIFICATION

One of the main objectives involves classifying images of dune regions in terms of land cover type. Given one or more images, pixel regions are collected into polygons that are homogeneous relative to land cover type, and then each resulting polygon is labeled by a person using one of a selected set of land cover labels (for example, 'sand', 'live vegetation' and 'dead vegetation'). Pixels are then chosen randomly from these polygons to form training sets and test sets for our classification models. Each pixel will have a coordinate location within the image which can then be mapped to a geolocation if the image is georeferenced. Each pixel will also have data from the camera used including visible light – often initially in RGB format – and sometimes additional data like near infrared, red-edge, or thermal readings. A machine learning algorithm such as a random forest or artificial neural network (ANN) can be constructed and trained to predict the resulting land classification.

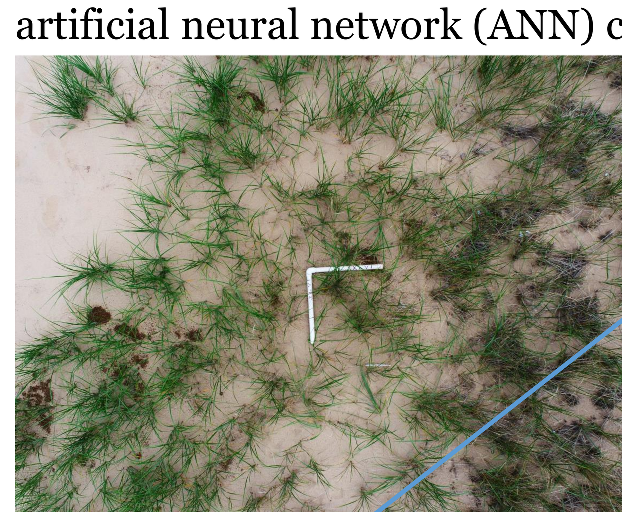

An original image can be covered with polygons by hand to create regions that have uniform land cover (e.g. sand, live vegetation, dead vegetation, or a masked PVC frame that we use for orientation and alignment).

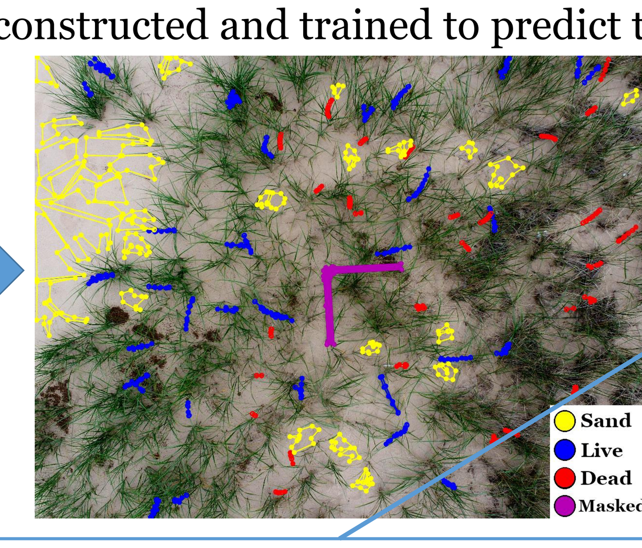

Pixels are chosen randomly from the polygons to provide data sets consisting of pixel labels (land cover type) and information returned from the drone camera about each pixel. Data returned includes visible light information, shown here in RGB format, and also multi-spectral data.

label	red	green	blue
live	43	83	49
sand	189	185	176
live	85	111	82
live	44	75	33
sand	184	181	176
live	130	161	127
sand	196	189	181
sand	198	189	180
live	43	79	31
live	89	122	69

Data is used for training and validating a random forest or ANN. The trained model can be used to classify the entire image.

Classified Image

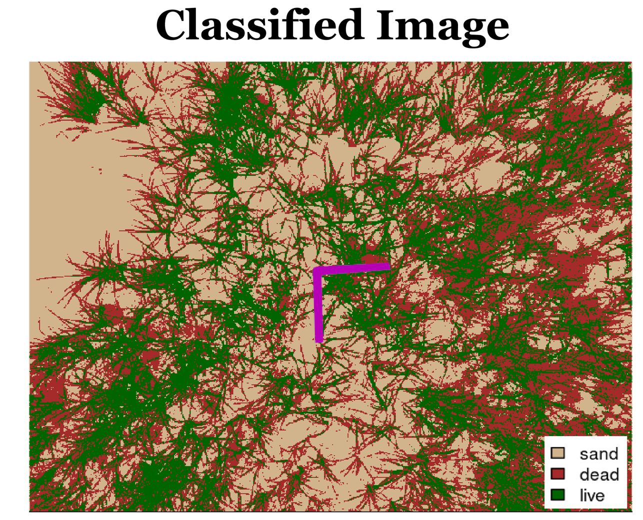

Once they are trained on the training dataset, machine learning models can be checked for accuracy by comparing model predictions on a validation dataset, which was not used in training the model.

In Summer 2022, we worked primarily with images taken at SHNA from an elevation of 4m on July 13, 2021. These were visible-light images that provided RGB data only. Other data features were developed from these, including HSV color layers and 48 texture feature layers at each pixel. We trained a random forest model on each of five chosen images, and then tested the land cover prediction results on pixels chosen from a test set of polygons which had not been used in model selection or training.

Each Quadrat is a chosen section of land that is approximately 2m by 3m.

Quadrat 46 (Image DJI_0010.JPG): Test Accuracy 100.0%
Quadrat 47 (Image DJI_0006.JPG): Test Accuracy 99.0%
Quadrat 48 (Image DJI_0014.JPG): Test Accuracy 99.7%
Quadrat 49 (Image DJI_0018.JPG): Test Accuracy 99.9%
Quadrat 50 (Image DJI_0022.JPG): Test Accuracy 98.0%

These results come from training a different random forest model on each image. One goal that we are working toward is training a single random forest model on data from multiple images that can then be used for prediction on pixels coming from any image in the set.

ACKNOWLEDGEMENTS

Research reported in this publication was supported in part by funding provided by the National Aeronautics and Space Administration (NASA), under award number 80NSSC20M0124, Michigan Space Grant Consortium (MSGC). Additional funding and support provided by Hope College. Thanks to Jack Krebsbach, Blake Harlow, and Theodore Lockett for help on this project.

IMAGE ALIGNMENT

Training robust and flexible models to predict land cover requires massive amounts of data, and these models can benefit from the use of images of the same region taken by different cameras from different elevations. This not only allows for combining data from cameras that return information in different spectra, but also provides opportunities for streamlining workflows when collecting data in the field. Images from multiple sources must be aligned and co-registered to create a usable data set.

Machine learning algorithms can be used to choose a mathematical function that maps one image to another. An unsupervised learning process can be used to choose key points in each image being matched, and then another algorithm can match key points in one image with key points in the other. Such a matching is shown below.

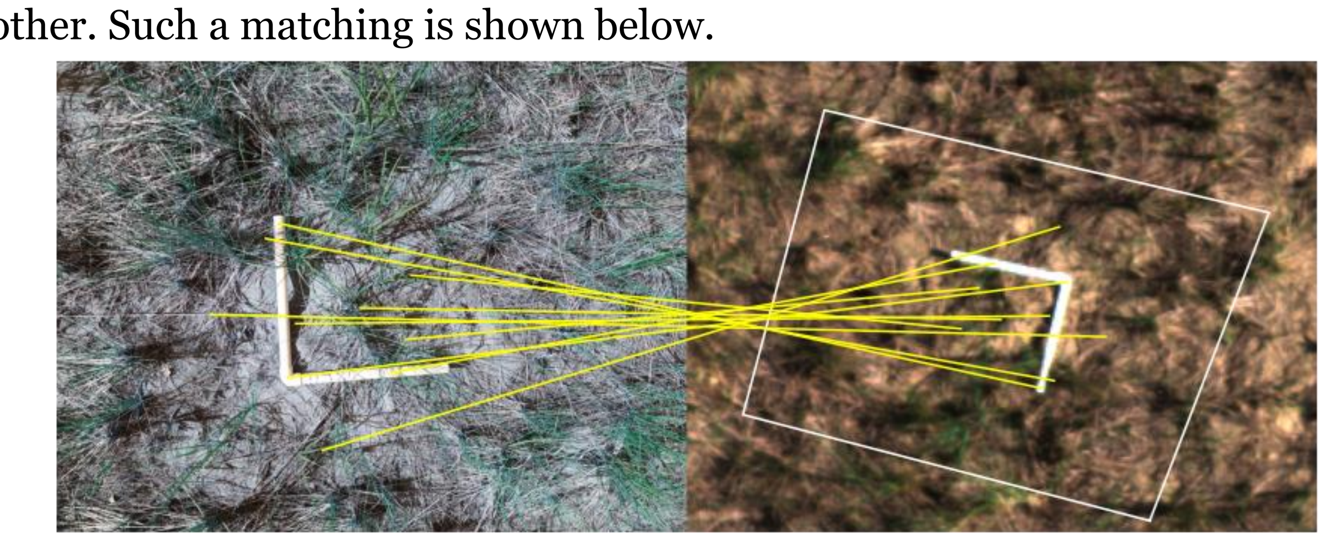

Once a matching of key points is known, an optimization process such as least squares can be used to choose a function mapping one image to the other. One option for such a function is a homography. A homography H maps a point (x,y) of the plane to the output given by

$$H(x,y) = (u/w, v/w) \quad \text{where} \quad \begin{bmatrix} u \\ v \\ w \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

The algorithm selects values of the constants (a,b,c,d,e,f,g,h) that map the key points to one another nearly optimally.

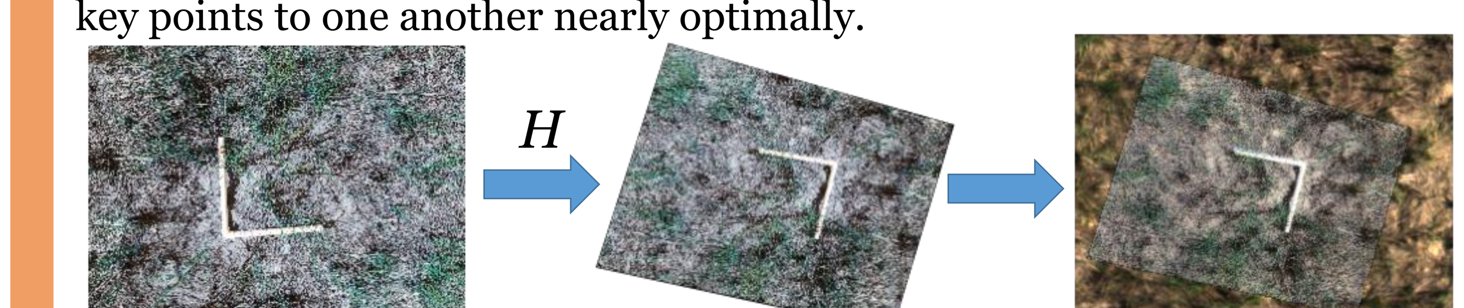

While homographies can give reasonably good matches between images, variation in elevation in the dune topography will likely give rise to cases where more general families of mappings are required. One goal of this work is to understand how to quantify the accuracy of the matching of two images and how to select more general mappings as needed.

FUTURE WORK

We plan to

- Continue to explore ways of aligning images from multiple sources to create a single, more comprehensive data set that can be used to train more robust classification models.
- Explore streamlined processes of data collection in the field that result from the ability to align images taken by different drones and cameras.
- Develop classification algorithms that make use of multiple input images to build models, including models trained by choosing polygons from different images.
- Explore land cover classification procedures that incorporate non-pixel-based methods including local texture analysis, image segmentation, and region/object-based approaches involving histograms of spectra.