GEG 133 Lab Exercise: Object Identification

Overview

This project uses GeoAI, or Geospatial Artificial Intelligence to detect palm trees. Geo AI is the integration of artificial intelligence (AI) with geospatial data and technology to analyze spatial information, solve complex geographic problems, and improve decision-making. It uses AI and machine learning (ML) to process vast amounts of location-based data, such as satellite images, GPS coordinates, and map data, to identify patterns, classify features, forecast trends, and build digital twins of the real world.

The project study area is in island of Tongatapu in Oceania where coconuts play an important role in agriculture. Plantations often contain thousands of trees, making manual inventorying a time- and resource-intensive task. Deep learning offers an effective solution for automatically detecting objects in imagery, but building a model from scratch can be laborious and slow. A more efficient option is to use a pretrained model.

In this tutorial, you'll use a pretrained model from the ArcGIS Living Atlas of the World and apply it to high-resolution imagery in ArcGIS Pro to accurately detect palm trees across the landscape. The resulting feature layer can help estimate tree counts, track agricultural practices, and evaluate potential production levels.

Note: This lab is data intensive and requires a GPU with at least 4 GB of dedicated memory. ArcGIS Image Analyst and Deep Learning Libraries for ArcGIS Pro are also needed. Use ArcGIS Pro 3.5x or higher.

Objectives

- Learn how deep learning can automate object detection in imagery.
- Access the Palm Tree Detection model from the ArcGIS Living Atlas of the World.
- Apply this model to high-resolution drone imagery using ArcGIS Pro and the Image Analyst extension.
- Use the Detect Objects Using Deep Learning tool to configure tool parameters such as input raster, output feature class, and model definition.
- Understand and apply optional settings like Non-Maximum Suppression and Batch Size.
- View and inspect the resulting **Detected_Palm_Trees** feature layer.

Acquire and Understand Landsat Data

Data Source Information

For this lab, I have downloaded and prepared a dataset for you from Esri. With ArcGIS GeoAI tools, you can use deep learning pretrained models or train your own models to extract features from raw data, such as detecting trees, digitizing building footprints, or generating land-cover maps.

A .ppkx file is an ArcGIS Pro project package and may contain maps, data, and other files that you can open in ArcGIS Pro.

Instructions

1. If you are using the S-drive and have submitted past labs and received a grade, you will want to delete most old data and labs. If you really like one of your past lab maps, please keep it separately and include it in your GIS portfolio. If you are in the GIST program, it will help you when you take the Capstone course. Please keep your project!
2. Download the packaged project data (.ppkx file) by logging in to [Acronis (https://acronis.monroecc.edu/signin)](https://acronis.monroecc.edu/signin) and navigating to *M-Drive* > *Courses* > *GEO* > *133-Little* > *6_DATA* > *1_GeoAI_Palm_tree*.

   If you are using the S-Drive, place the file on your virtual desktop. In the next step, unzip the file and place it in your student storage.
3. Locate the downloaded file on your computer and double-click Palm_Tree_Detection.ppkx to open the file in ArcGIS Pro. If prompted, sign in with your ArcGIS account.
4. In ArcGIS Pro, you will see a map centered on the village of Kolovai on the Tongan island of Tongatapu. An image layer is displayed on top of the topographic basemap (see Figure 1)

   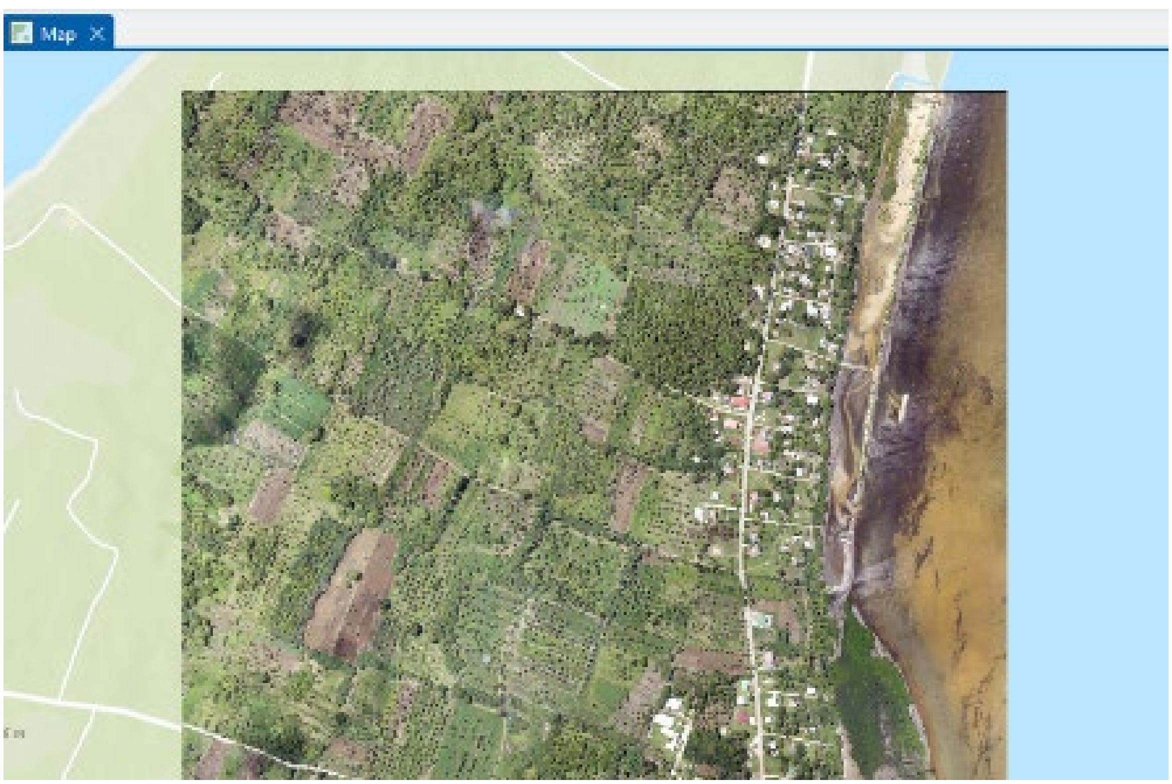

   Figure 1: View of the image layer displayed on top of the basemap in ArcGIS Pro.
5. Zoom in and pan over the map to examine the image.

   This imagery was captured by a drone and orthorectified to remove any distortions. It is high resolution (each pixel represents about 9 by 9 centimeters on the ground) and shows palm trees and other features quite clearly. It is in the TIFF file format with three bands: red, green, and blue, which together form a natural color picture.

Identifying all the palm trees in this image would be very time consuming. Instead, you'll use deep learning to detect them automatically.

Stop. Answer Question 1.

Detect Palm Trees

Deep Learning

You will be using deep learning to detect palm trees in the imagery. If you don't already have a deep learning model available, this first requires training a model from scratch and feeding it large numbers of examples to show the model what a palm tree is. High-performing models can require exposure to tens of thousands of examples. An alternative is to use a model that has already been trained. For this activity, you will retrieve such a model and apply it to your imagery.

Important note: Using the deep learning tools in ArcGIS Pro requires that you have the correct deep learning libraries installed on your computer. If you do not have these files installed, save your project, close ArcGIS Pro, and follow the steps delineated in the "Get ready for deep learning in ArcGIS Pro" instructions. In these instructions, you will also learn how to check whether your computer hardware and software are able to run deep learning workflows and other useful tips. Once done, you can reopen your project and continue with the tutorial.

Instructions

1. Before beginning, read any important messages about this lab in Brightspace.
2. In ArcGIS, navigate to the View tab and select Geoprocessing from the Windows group (see Figure 2).

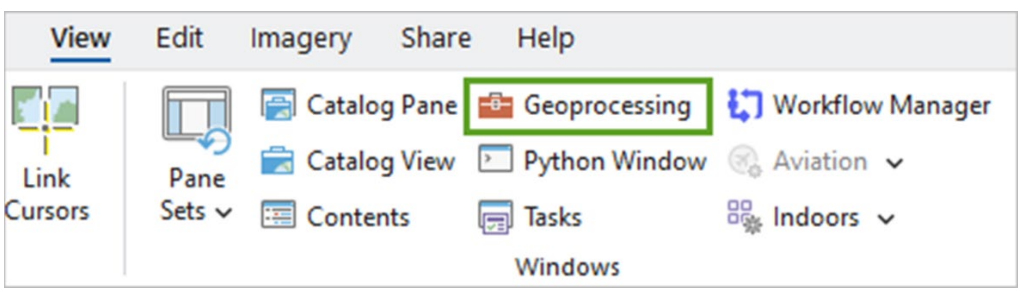

Figure 2: ArcGIS tool ribbon highlighting the Geoprocessing pane selection.

3. In the Geoprocessing pane search box, type "Detect Objects Using Deep Learning." In the list of results, click on the box labeled "Detect Objects Using Deep Learning" to open the tool (see Figure 3).

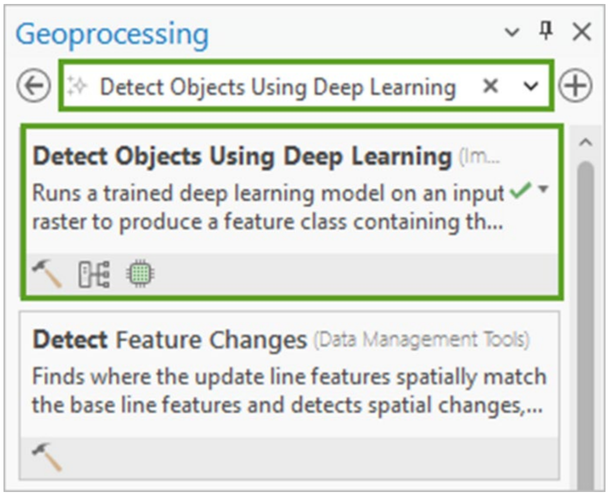

Figure 3: Screenshot of the Geoprocessing pane in ArcGIS. The search bar contains text "Detect Objects Using Deep Learning" and the tool with that name has been selected in the search results.

4. Set the following parameter values in the Detect Objects Using Deep Learning tool (see Figure 4):
   - Input Raster: select Kolovai_Imagery.tif from the drop down menu
   - Output Detected Objects: type Detected_Palm_trees in the form field
   - Model Definition: Browse by selecting the folder icon next to the form field

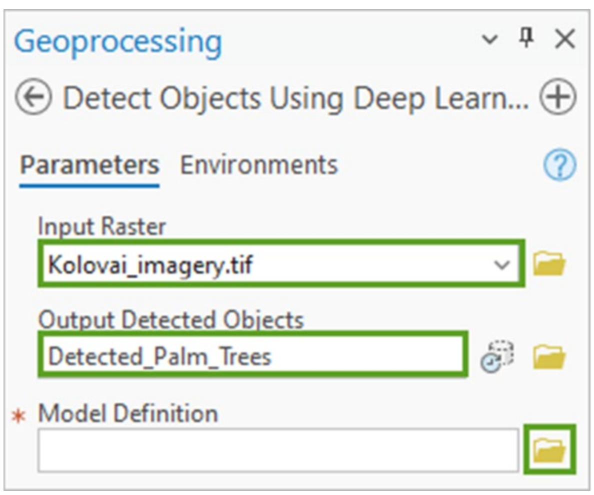

Figure 4: Detect Objects Using Deep Learning tool menu showing the Parameters tab with fields for Input Raster, Output Detected Objects, and Model Definition filled with the parameters described in Step 4.

5. Save.
6. In the Model Definition Browse window, select *Portal* > *Living Atlas* in the navigation menu. Then type “Palm Tree Detection” into the search box, select the “Palm Tree Detection” file from the results and click OK (See Figure 5).

   Once you select the deep learning model, the model arguments load automatically.

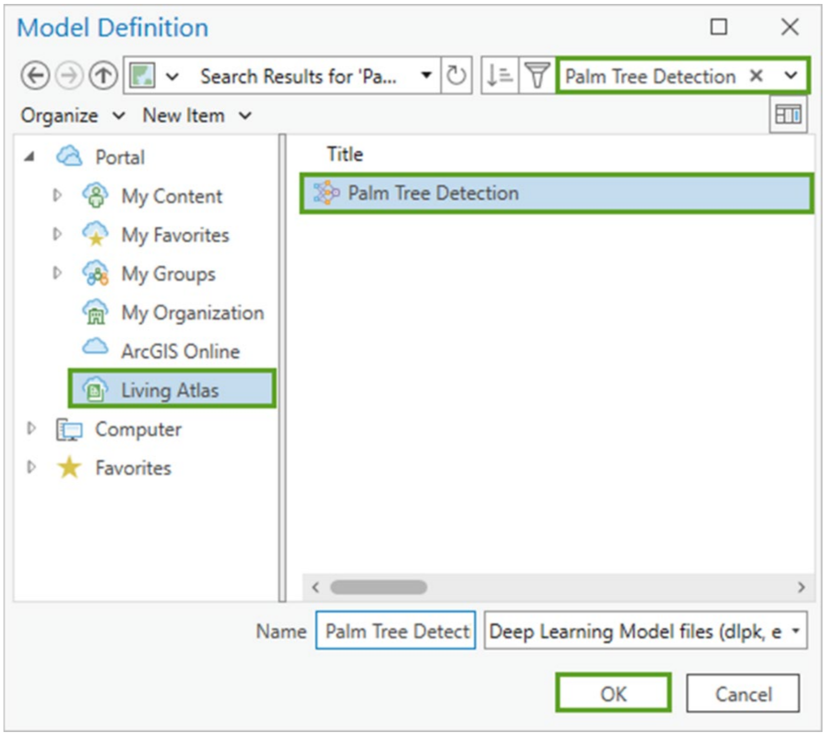

Figure 5: Model Definition browse window with the Palm Tree Detection file selected from search results in the Living Atlas tab.

This will retrieve the Palm Tree Detection pretrained model from ArcGIS Living Atlas of the World.

ArcGIS Living Atlas of the World is Esri's authoritative collection of GIS data. It includes a growing library of over 100 pretrained deep learning models to detect various objects from buildings to ships to agricultural field boundaries.

Note: You can learn more about the Palm Tree Detection model by retrieving it on the [ArcGIS Living Atlas site](). The model is meant to detect palm trees on high resolution imagery (5 to 15 centimeters), such as drone imagery. The imagery is expected to have three bands: red, green, and blue (RGB). This model is a good match for your imagery.

7. Keep all the default values for the Model Definition Arguments.

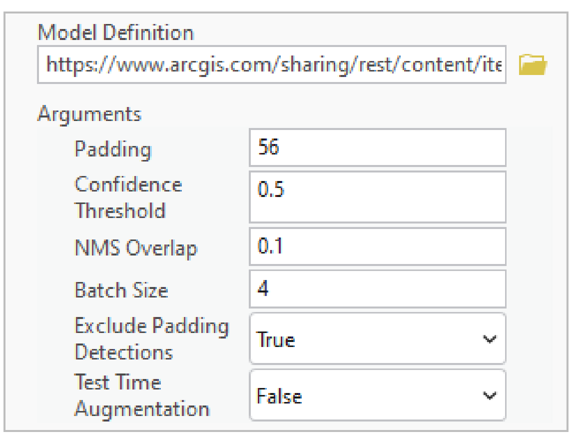

Figure 6: List of Model Definition Arguments (Padding, Confidence Threshold, NMS Overlap, Batch Size, Exclude Padding Detections, and Test Time Augmentation) as specified by the Palm Tree Detection file.

Two arguments are of particular interest:

- **Confidence Threshold:** A threshold of 0.5 means that an object detected will only be added to the output dataset if the model has at least a 50 percent confidence that the object is indeed a palm tree. The optimal threshold value can be found by trial and error.
- **Batch Size:** Deep learning object detection cannot be performed on the entire image in one go. Instead, the tool will cut the image into small pieces known as chips. A batch size of 4 means that the tool will process four image chips at a time. As you run the tool, you may get an out of memory error because your computer doesn't have enough memory for that level of processing. In that case, try decreasing the Batch Size value from 4 to 2 or even 1. Decreasing this value will not affect the quality of the model, only the efficiency of the model's detection process.

8. Check the box for Non Maximum Suppression and enter “0.5” in the field for Max Overlap Ratio (See Figure 7).

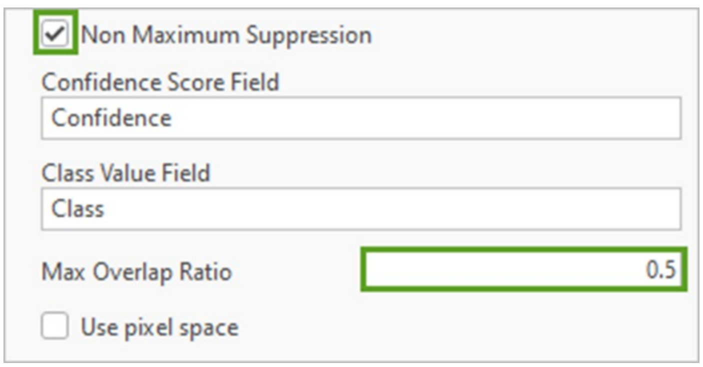

Figure 7: Selection box for Non Maximum Suppression and form field for Max Overlap Ratio.

Sometimes the model detects an object more than once. **Non Maximum Suppression (NMS)** is an optional process that suppresses some of the detected objects when there is duplication. Figure 8 shows an example of this, with one palm tree that has been detected twice. With NMS, only one detected object will be kept. The object that was detected with the highest confidence is kept, the other objects are removed.

Figure 8: One palm tree that has been detected twice. Non Maximum Suppression prevents this from happening.

The **Max Overlap Ratio** parameter determines how much overlap there must be between two detected objects for them to be considered duplicates and for NMS to apply. A value of 0.5 means that the overlap must be of 50 percent or more. Figure 9 shows two detected objects with overlap of less than 50 percent, so they don't represent the same object and NMS will not be applied.

Figure 9: Correct detection of two palm trees, illustrating detection areas with an overlap of less than 50%.

9. Select the Environments tab of the Detect Objects Using Deep Learning tool menu (see Figure 10)

   At this point, you could run the tool as is, and it would proceed to detect palm trees over the entire Kolovai_imagery.tif image, which could take 20 minutes to 1 hour based on your computer's specifications. For the brevity of this tutorial, you'll only detect palm trees in a small subset of the input image.

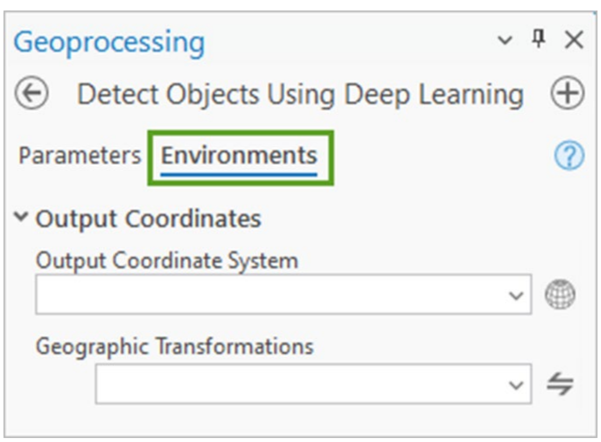

Figure 10: Detect Objects Using Deep Learning tool menu showing the Environments tab with fields in the Output Coordinates section including Output Coordinate System and Geographic Transformations.

10. Navigate to the Map tab, and Open Bookmarks in the Navigate group of the tool ribbon. Choose Detection area from the Map Bookmarks options (see Figure 11).

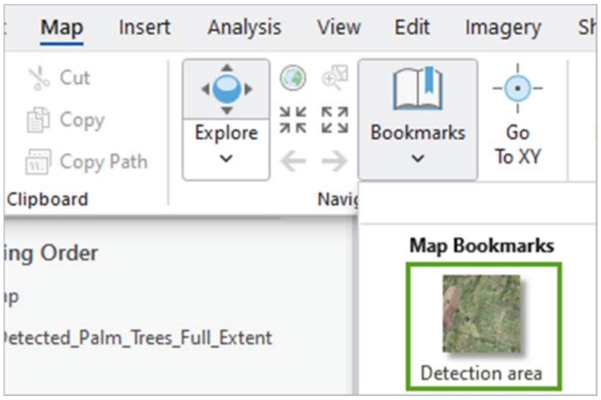

Figure 11: The Map tab of the tool ribbon with the Bookmarks tool open in the Navigate group, showing the Detection area option in the drop down menu.

11. The map will zoom in to a smaller area of the map of Kolovai. Continue to zoom in to an even smaller area so running the program takes less time and less data (see Figure 12).

Figure 12: A small area of the map of Kolovai, visible after selecting the Detection area tool and zooming in.

12. Return to the Environments tab of the Detect Objects Using Deep Learning tool and expand the Processing Extent section. Select the Current Display Extent button (the map icon in the row of buttons) to view the Top, Left, Right, and Bottom coordinates of your current view of the map (See Figure 13; your numbers might be different).

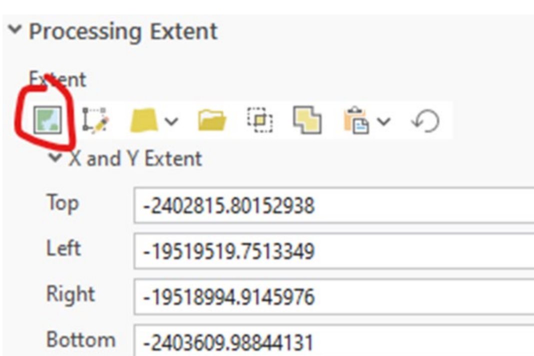

Figure 13: The Processing Extent section of the Environments tab with the Current Display Extent button circled in the Detect Objects Using Deep Learning tool menu.

13. Accept all other default values and click the Run button. It will take a few minutes for the program to finish running. You can monitor progress of the process under the Run button, and select View Details to see more information (see Figure 14).

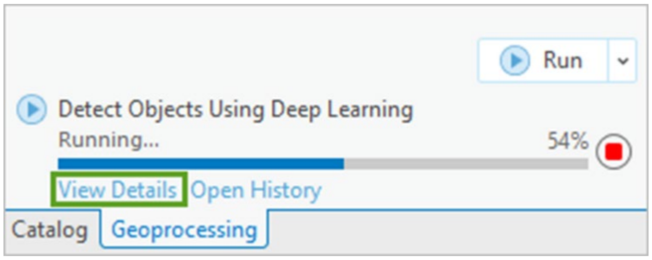

Figure 14: The Run button and loading bar showing percent completion of the program. Options to View Details and Open History are below the progress bar.

Note: This tutorial recommends using an NVIDIA GPU with a minimum of 4 GB of dedicated memory. If you don't have a GPU, the tool will use your CPU, but it will take longer to process the data. If you are not sure whether your computer has a GPU and what its specifications are, see the Check for GPU availability section in the Get ready for deep learning in ArcGIS Pro tutorial.

For more information about the role of the different arguments in the model, see the Palm Tree Detection guide.

14. When the program has finished running you will have successfully detected palm trees in an area of Kolovai using a deep learning pretrained model. ***Save!***

Stop. Answer Questions 2-5.

Style the Result Layer

Next, change the style of the Detected_Palm_Trees layer and review it.

Instructions

1. In the Contents pane, expand the Detected_Palm_Trees option and click the icon to display the Symbology pane (see Figure 15).

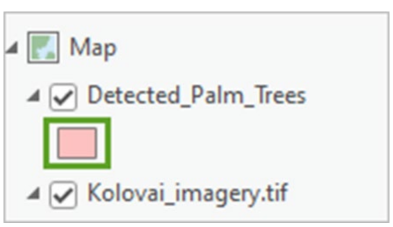

Figure 15: Selecting the icon below Detected_Palm_Trees in the Map list opens the Symbology pane.

2. In the Symbology pane, open the Properties tab (see Figure 16).

Figure 16: The Symbology pane for Detected_Palm_Trees showing the Properties tab under the heading Format Polygon Symbol.

3. Expand the options for Appearance (see Figure 17) and set the following parameters:
    a. Color: no color (transparent)
    b. Outline Color: any bright yellow color, such as Solar Yellow.
    c. Outline Width: 2pt

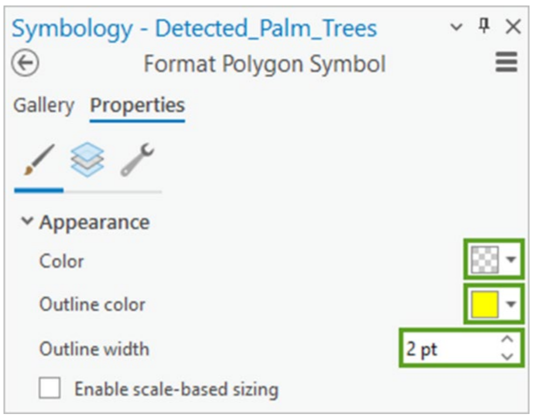

Figure 17: The Properties tab of the Symbology pane with selections in the Appearance section as described in Step 3.

4. Click Apply to update the layer with the new symbology. (See Figure 18)

Figure 18:View of the map after applying the new symbology parameters. Yellow squares now indicate where palm trees have been detected.

5. On the map, zoom in and inspect the Detected_Palm_Trees layer. If the model was successful, it will have identified most of the palm trees and there will be very few cases of false positives—where the model mistakenly found a palm tree where there is none (see Figure 19).

Figure 19: An example of successful palm tree detection on the map.

Note: Running the tool on the entire image will take 40 minutes or longer depending on your computer's specifications. Alternatively, a map showing the result layer for the full extent was included in your project. To view it, go to the Catalog pane, expand Maps, right-click Map_results and choose Open. Over 12,500 palm trees were detected over the entire image.

6. In the Quick Access Toolbar, click the Save Project button to save your project (see Figure 20).

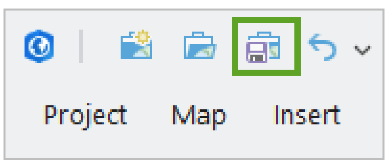

Figure 20: The Save Project button icon in the Quick Access toolbar.

7. Zoom out until you can see all of your map and use the Snipping Tool app on your computer to take a screenshot of your map, including the Contents tab. Use the pen tool to initial the screenshot (see Figure 21).

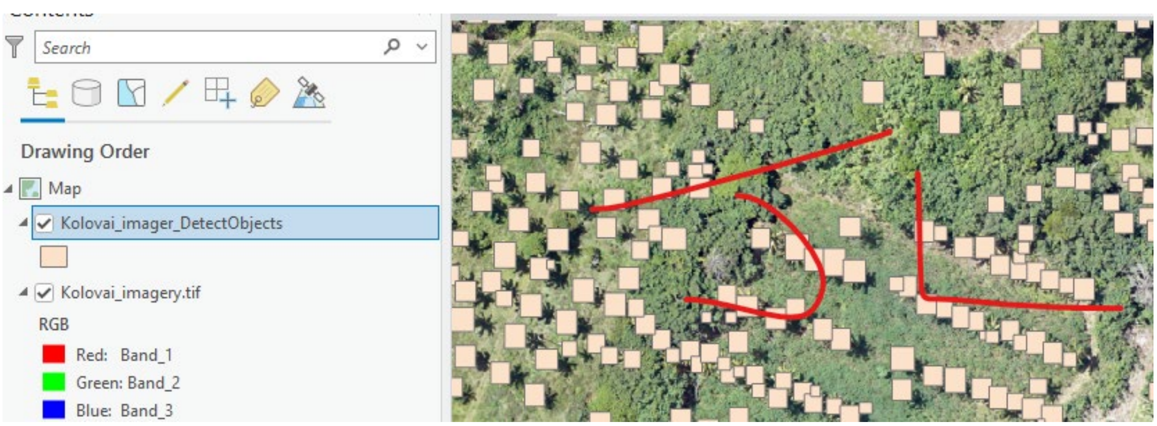

Figure 21: An example screenshot of the map and Contents tab with initials drawn over the map.

Additional Resources

Browse other pretrained models on the Esri ArcGIS Living Atlas of the World Website.

Submission Details

Submit your answer sheet and screenshot for this lab in the dropbox.

Lab Rubric

- Questions = 3 points
- Screenshot of Map = 7 points