WHAT'S NEW IN ENVI DEEP LEARNING 2.1

May 2023

Zach Norman
Bill Okubo

Agenda

Introduction

Deep Learning

Machine Learning

Notables

Q + A

Speakers Slide

Bill Okubo

Product Manager

bill.okubo@nv5.com

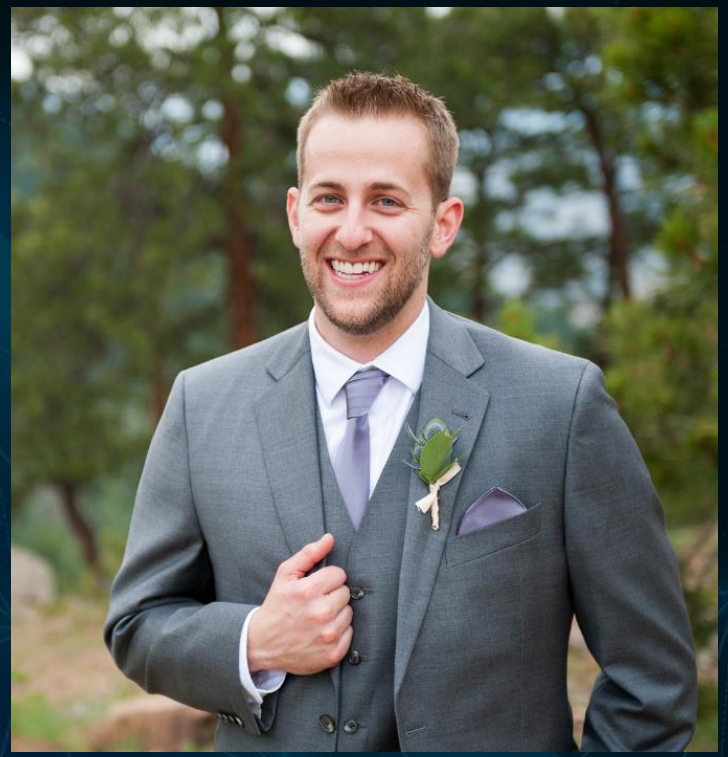

Zach Norman

Product Manager

zachary.norman@nv5.com

Context: The ENVI Deep Learning Module

Applied deep learning for geospatial imagery in ENVI, the leading remote sensing and image analysis software

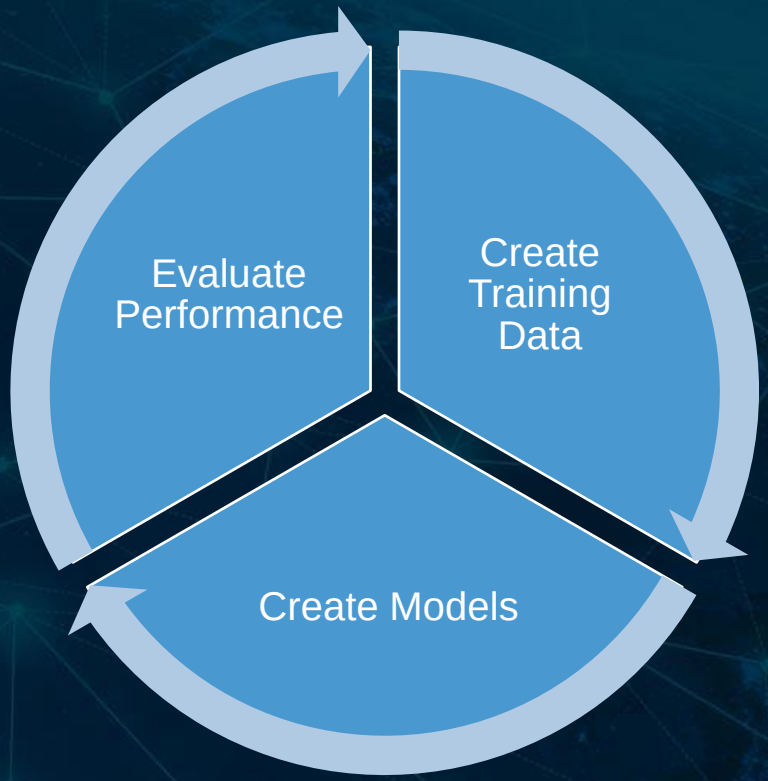

Deep learning workflow in ENVI, built on TensorFlow and Keras

Without needing to program, the capabilities include:

- Segmentation (i.e. cloud masking)
- Object detection (i.e. cars or ships)
- Linear feature extraction (i.e. roads)
- Support for nearly any image format and data modality

Assess building damage after hurricanes and tornadoes

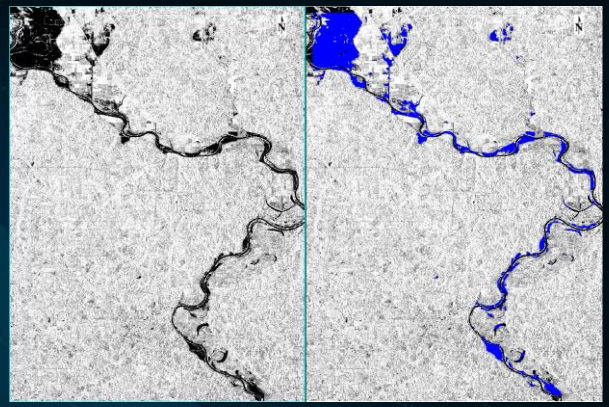

Automated flood detection using SAR

Context: ENVI Machine Learning

Applied machine learning for geospatial imagery in ENVI, the leading remote sensing and image analysis software

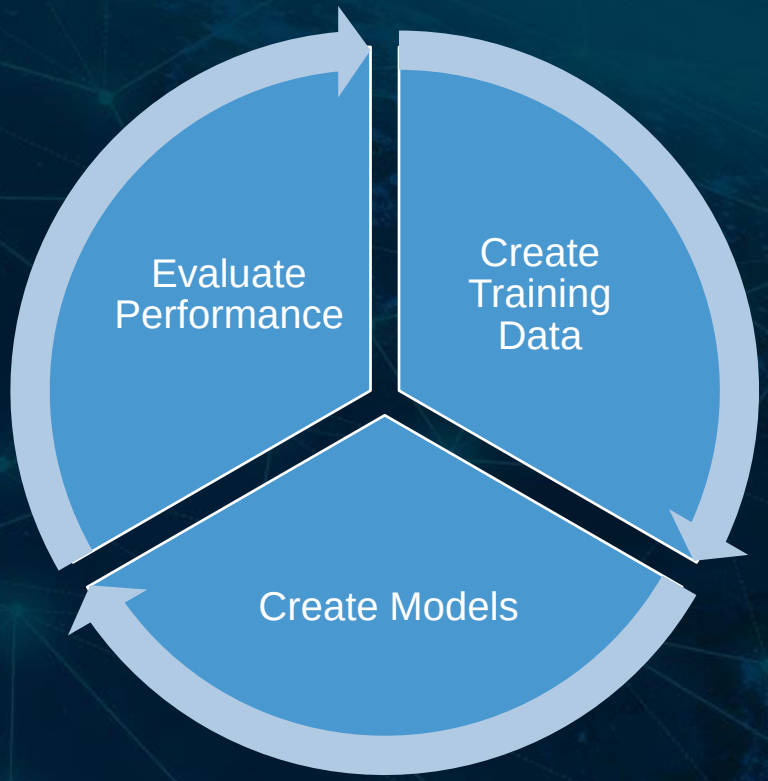

Machine learning follows the same workflow as deep learning, but takes a lot less time

Without needing to program, you have access to:

- Supervised classification (i.e. Random Forest)
- Unsupervised classification (i.e. K-Means)
- Anomaly detection (i.e. Isolation Forest)
- Supports nearly any image format, data modality, or number of bands

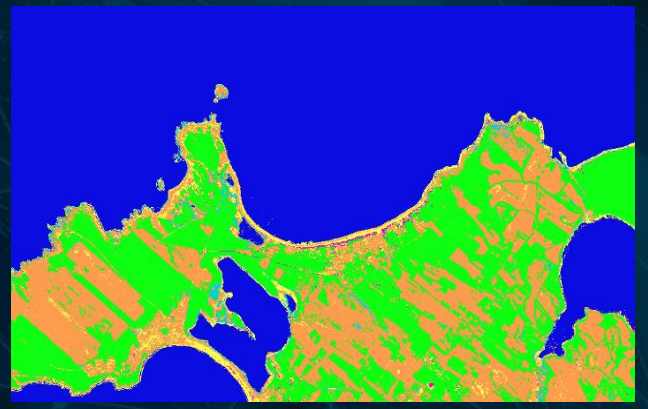

Get answers in minutes, not hours using machine learning

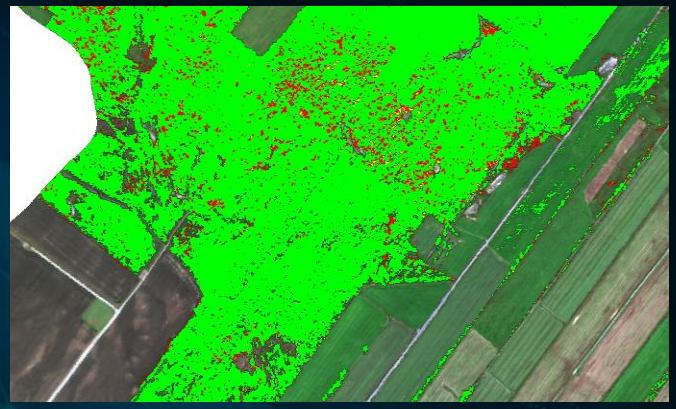

Use machine learning to process hyperspectral imagery with any number of bands

Machine Learning vs Deep Learning: Hardware and Access

MACHINE LEARNING

- **FREE and no additional license required**, packaged with "ENVI Deep Learning"
- Available starting with ENVI 5.6.3
- Uses the CPU
- Recommended 16 GB RAM
- Libraries optimized for Intel, but AMD also supported
- Some algorithms have significant performance improvements with Intel

DEEP LEARNING

- Need to purchase the "ENVI Deep Learning" module to use
- Latest version with be 2.0 and will have machine learning
- Requires a GPU
- Minimum 8 GB GPU RAM
- Only NVIDIA GPUs are supported

DEEP LEARNING

Deep Insights Using TensorBoard

- TensorBoard is TensorFlow's visualization engine for deep learning model performance
- TensorBoard allows us to:
  - Visualize performance metrics (accuracy, precision recall, loss, etc.)
  - Display metrics from separate training sessions in one chart
  - Create visualizations for other types of data, such as images, audio, or text

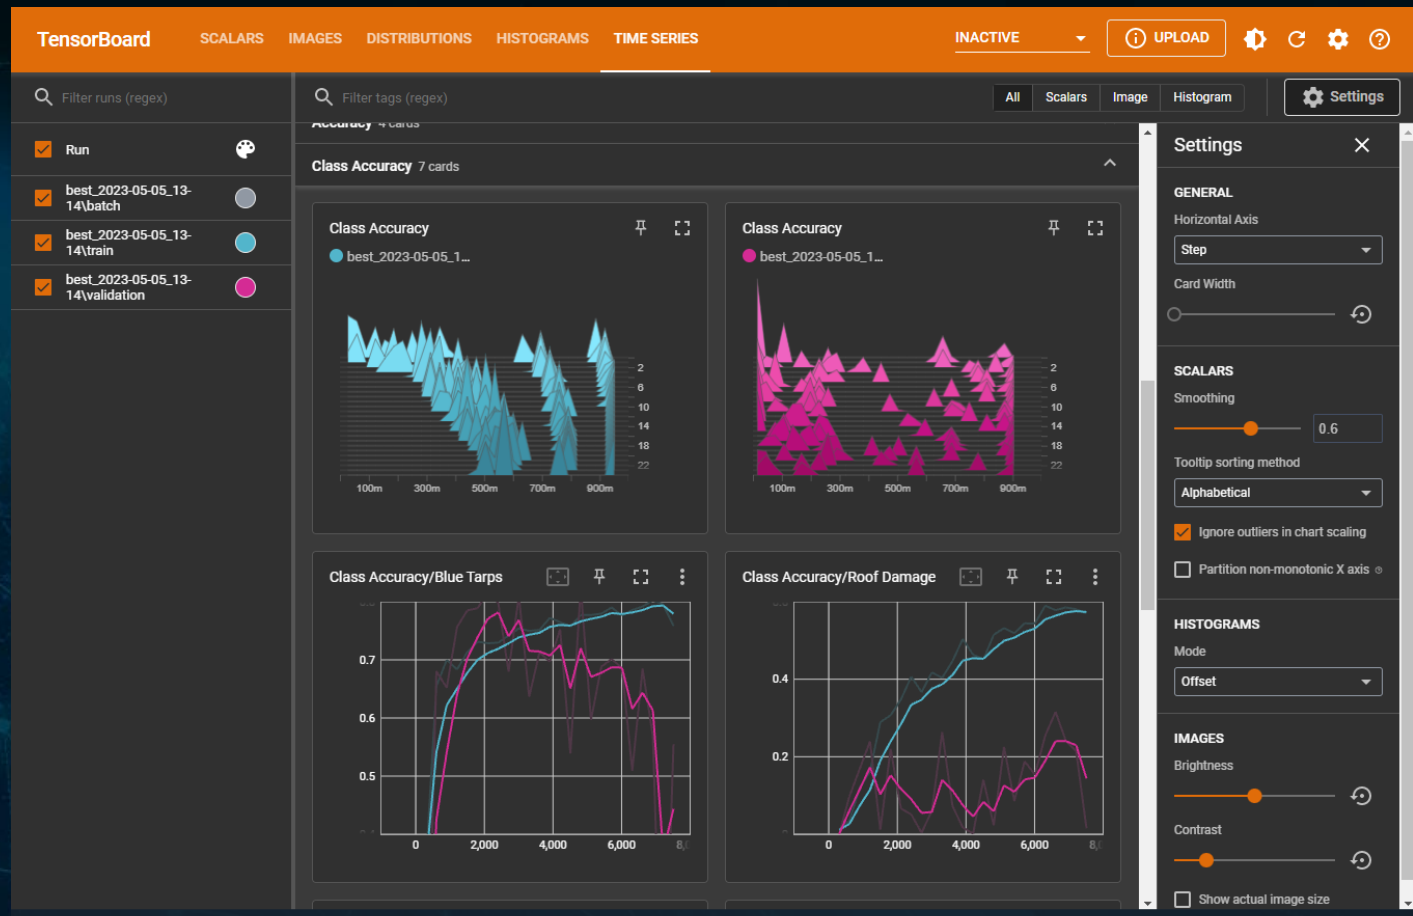

Screenshot of TensorBoard in ENVI Deep Learning, showing some new metrics that are now reported

The Old Way

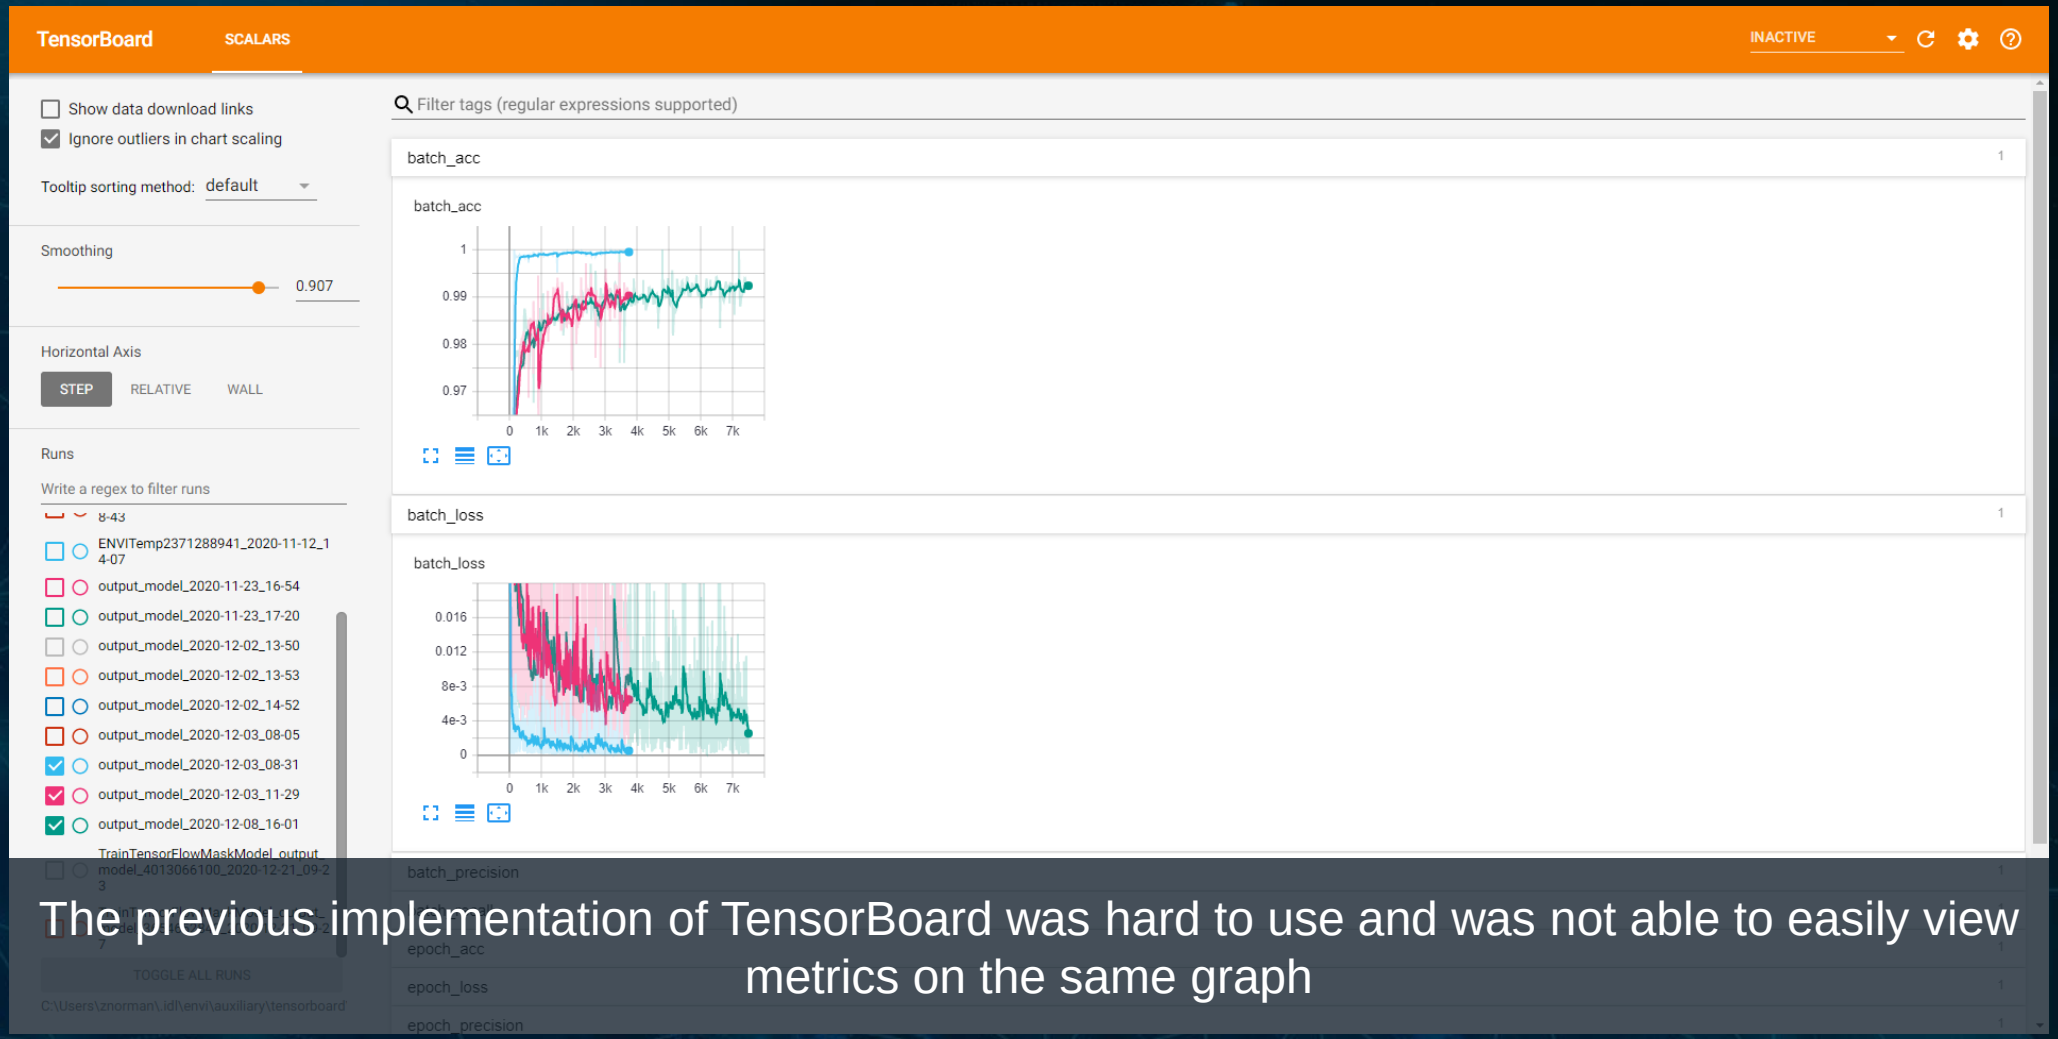

The previous implementation of TensorBoard was hard to use and was not able to easily view metrics on the same graph

The New Way

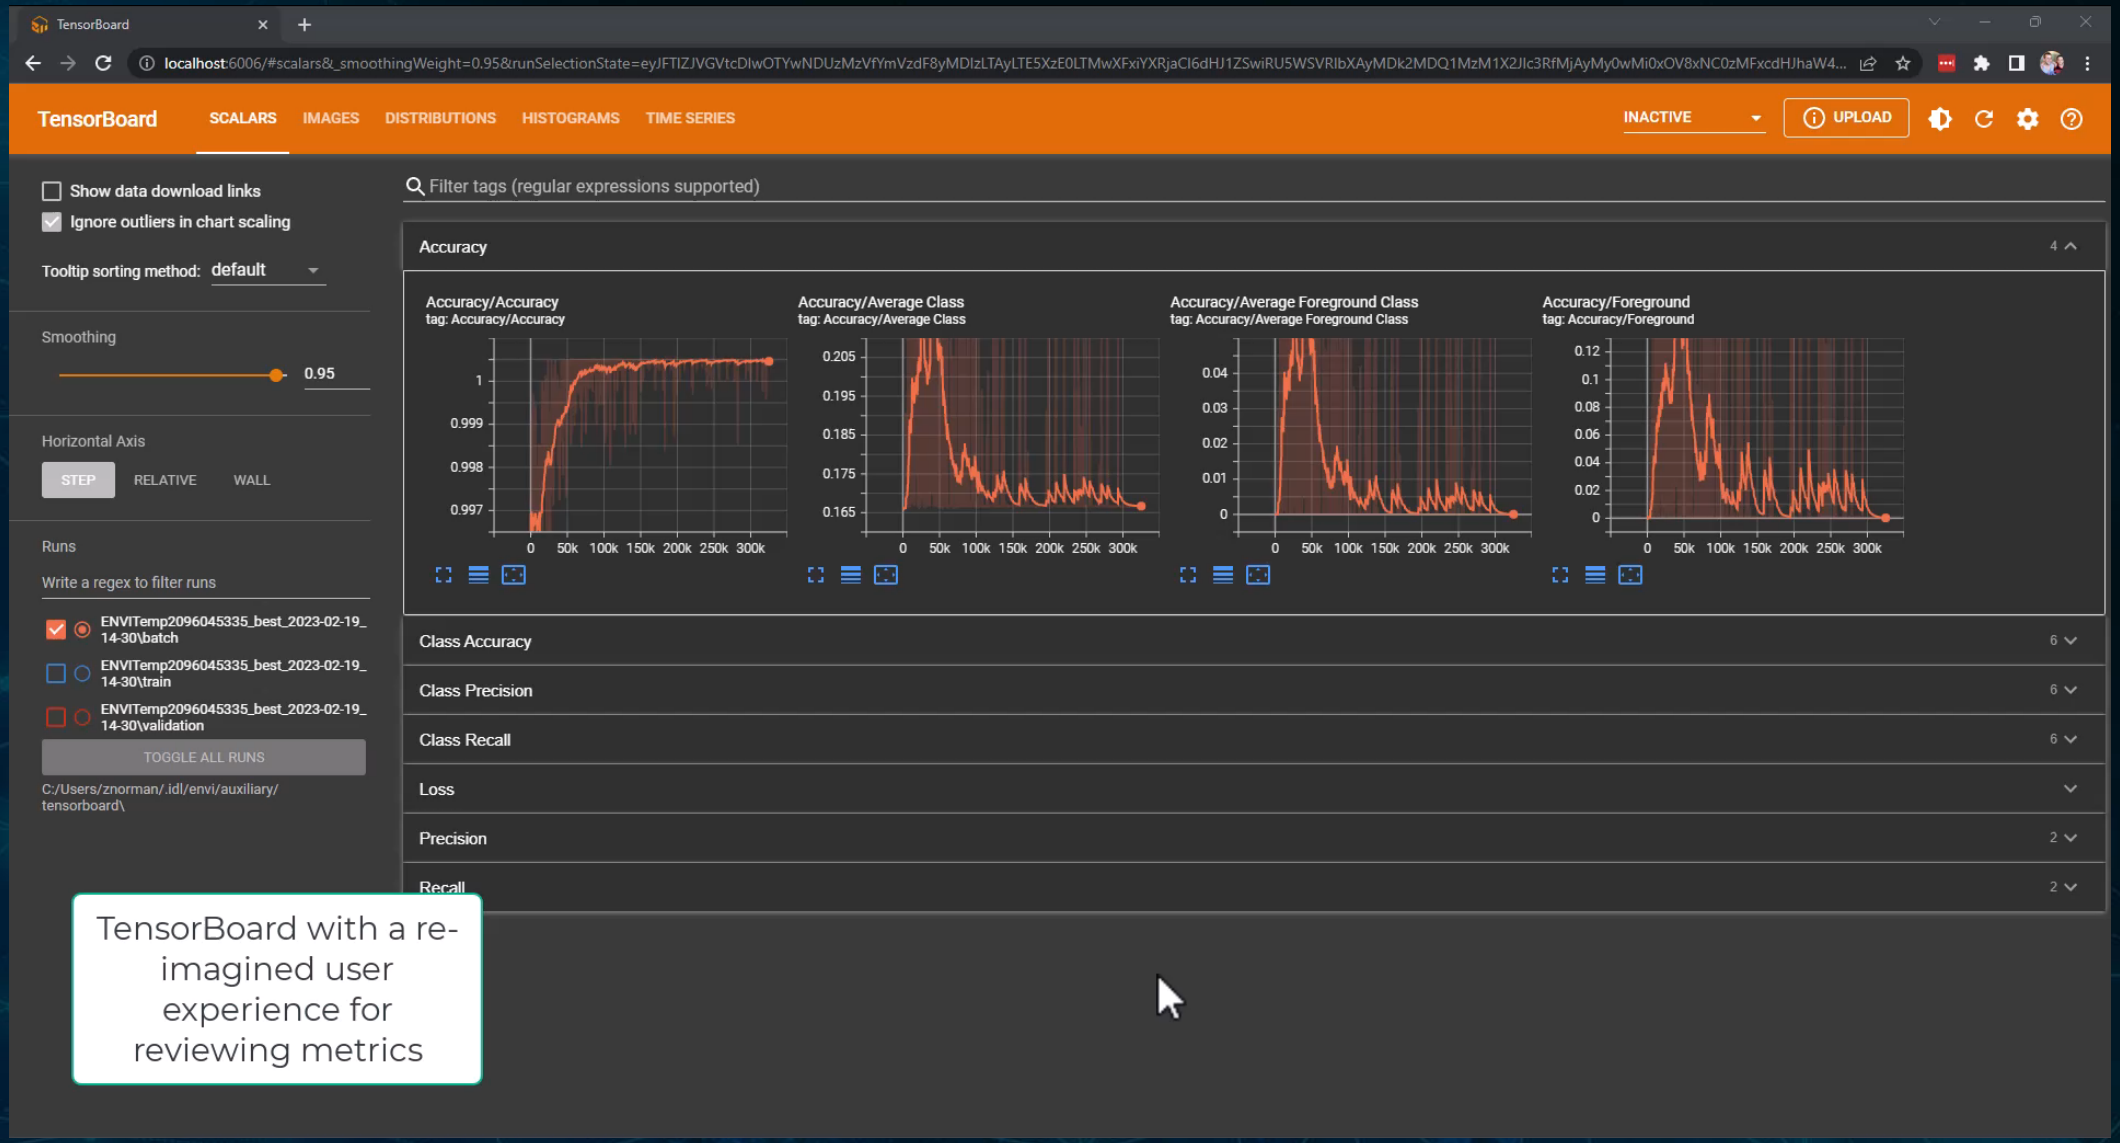

TensorBoard with a re-imagined user experience for reviewing metrics

Notable Changes

- Visualize multiple metrics on one graph!
- X-axis represents number of samples that have been trained on
  - Allows comparisons between training sessions with different batch sizes on different GPUs
- As of ENVI Deep learning 2.0, we have had support for dark mode in TensorBoard

- Each training session get's three TensorBoard logs:
  - "batch" which covers each step of training
    - Updates each step of training
  - "training" which represents the performance of the model on data used directly for training
    - Once per epoch
  - "validation" which represents the performance of the model on data used for validation (i.e. evaluation)
    - Once per epoch

Overall Metrics

- Accuracy
  - How well the model re-creates its input
  - Higher is better
- Precision
  - Measures the accuracy of positive predictions
  - Higher is better
- Recall
  - Measures the completeness of positive predictions
  - Higher is better
- Loss
  - The overall performance of the model
  - Lower is better

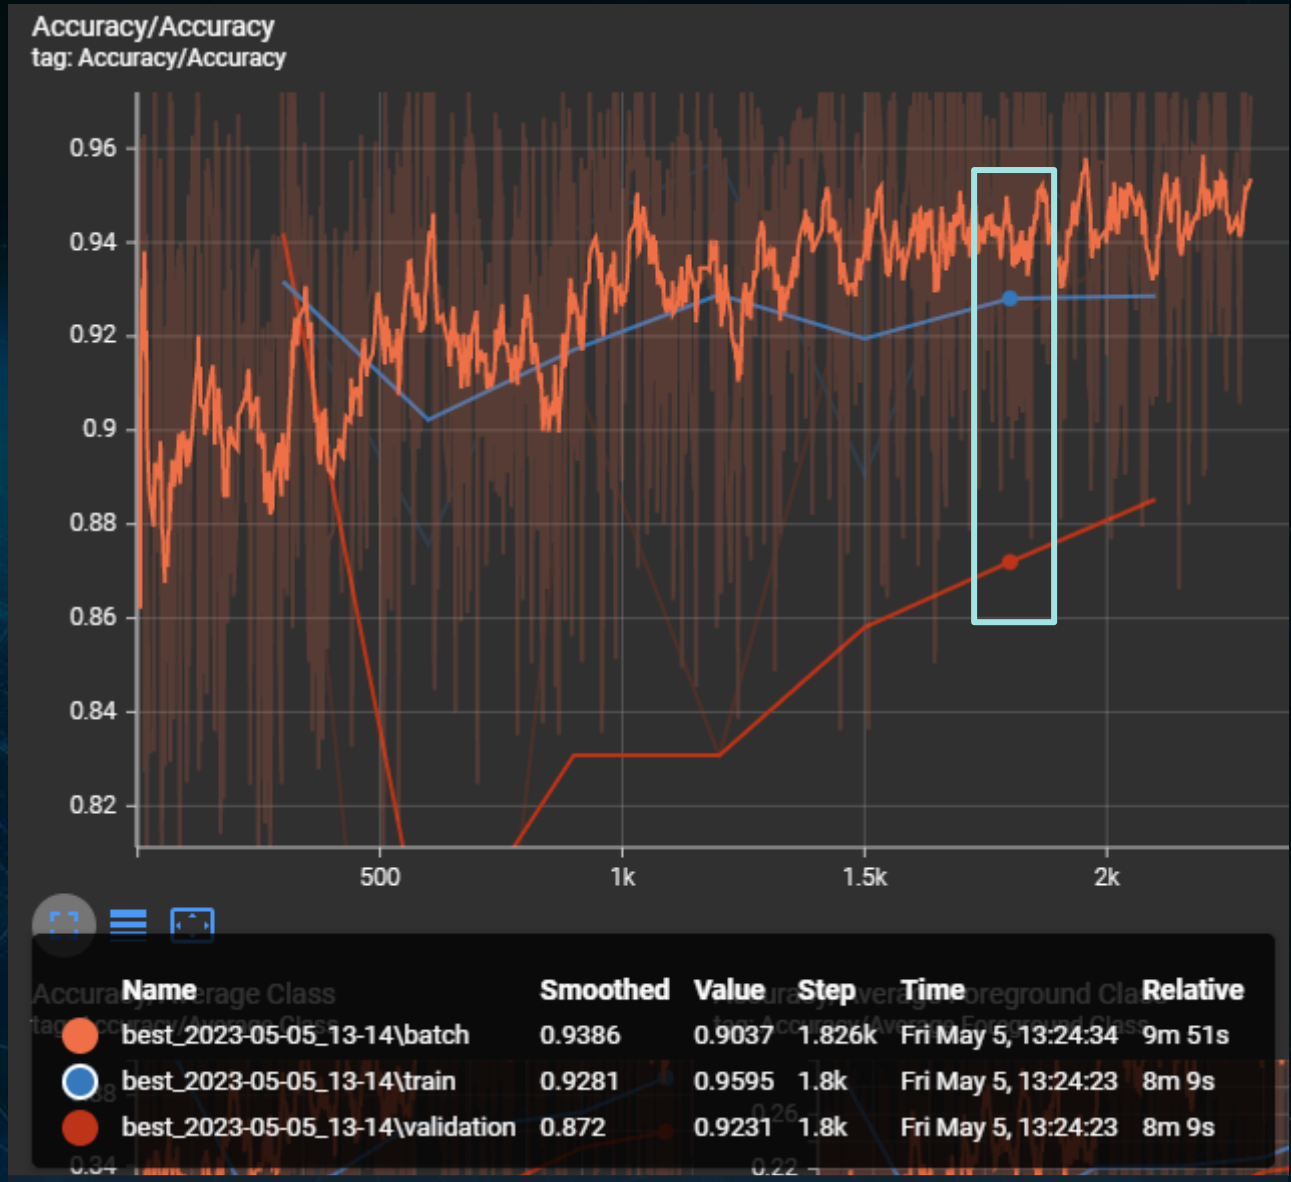

A look at overall accuracy being reported in TensorBoard which allows you to view accuracy during each step, for each epoch, and for validation data in the same plot.

Types of Accuracy

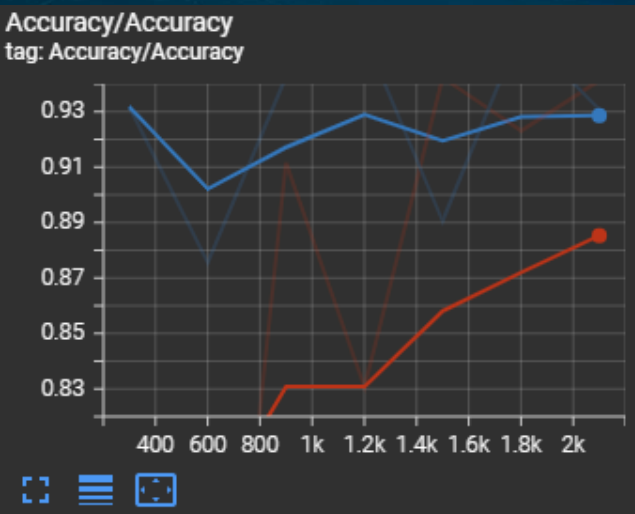

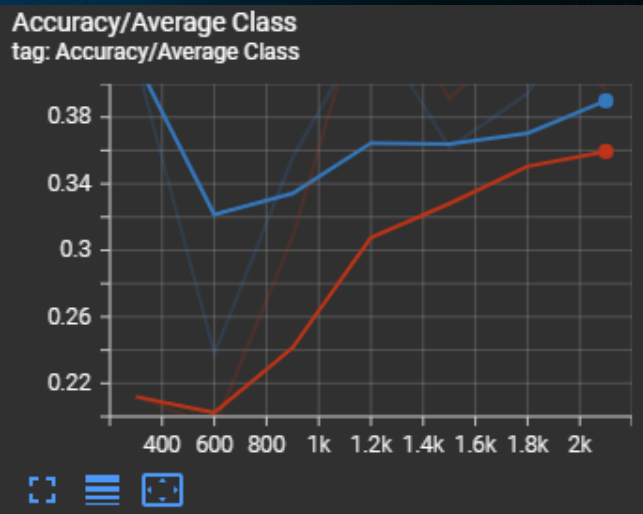

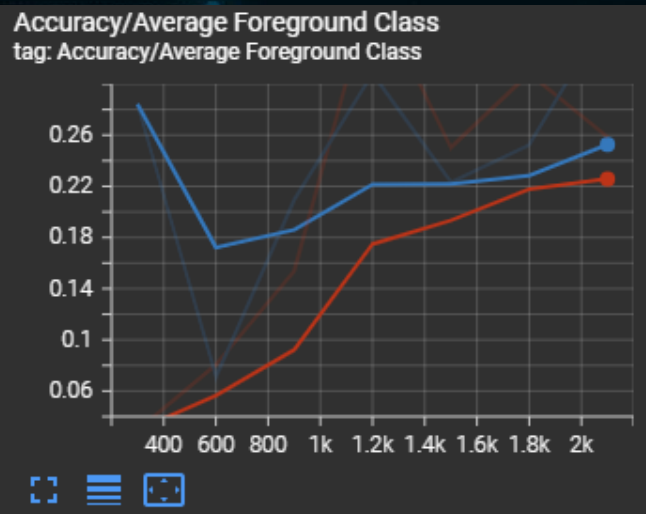

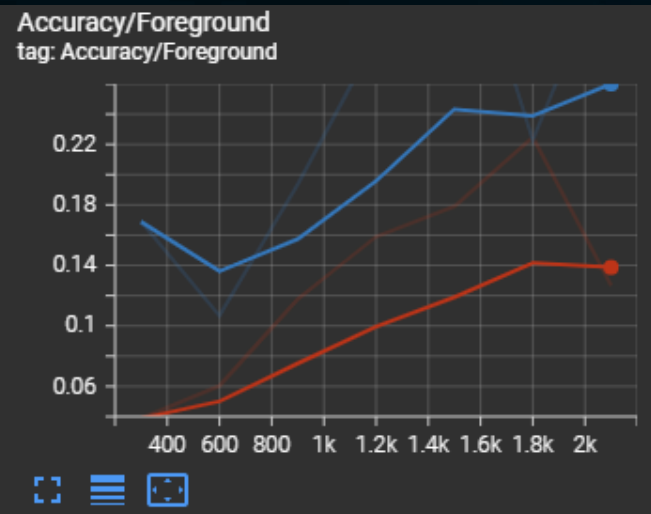

- Accuracy
  - Overall, traditional value
- Average class
  - Accuracy of each individual class, averaged together
- Average foreground class
  - Accuracy of each foreground classed, averaged together (i.e. features you are interested in and have labeled)
  - Useful for rare or small features when background outnumbers foreground
- Foreground
  - Accuracy for non-background features (i.e. features you are interested in and have labeled)
  - Useful for rare or small features when background outnumbers foreground

Per Class Metrics

- At the end of each epoch of training, we now report the accuracy, precision, and recall of each class
- You can view this for training and validation, but not for each batch
- Depending on how rare your features are, and how much training data you have, these values may jump around a bit

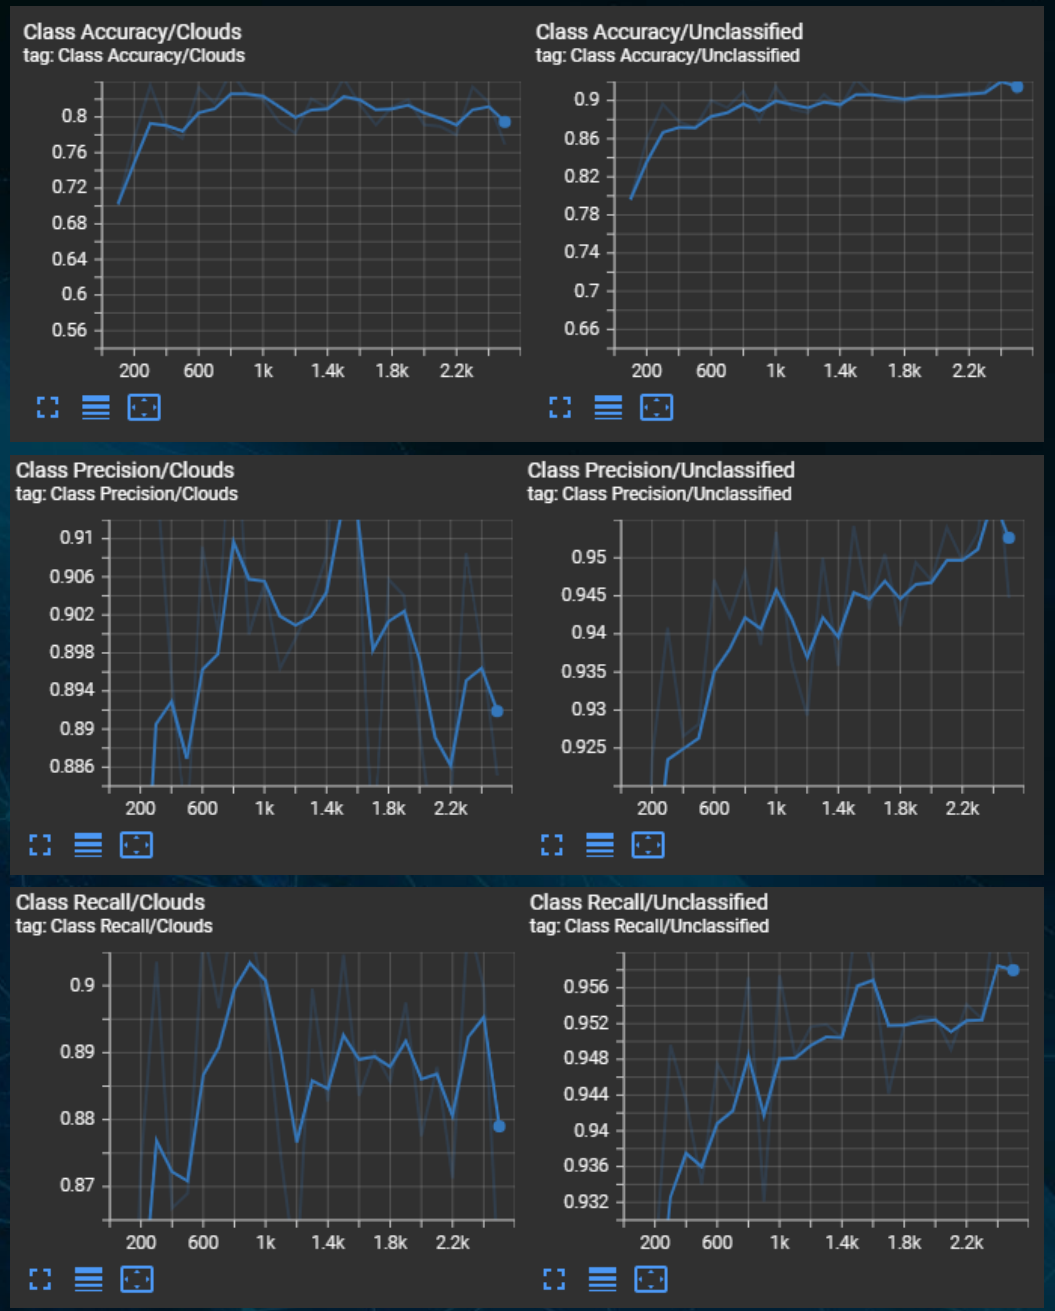

Multi-Class Class Metrics

Example of accuracy, precision, and recall for foreground classes

Blue represents training and **red** is validation

Here's the legend you'll see in TensorBoard

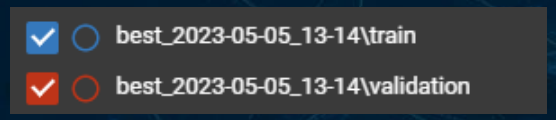

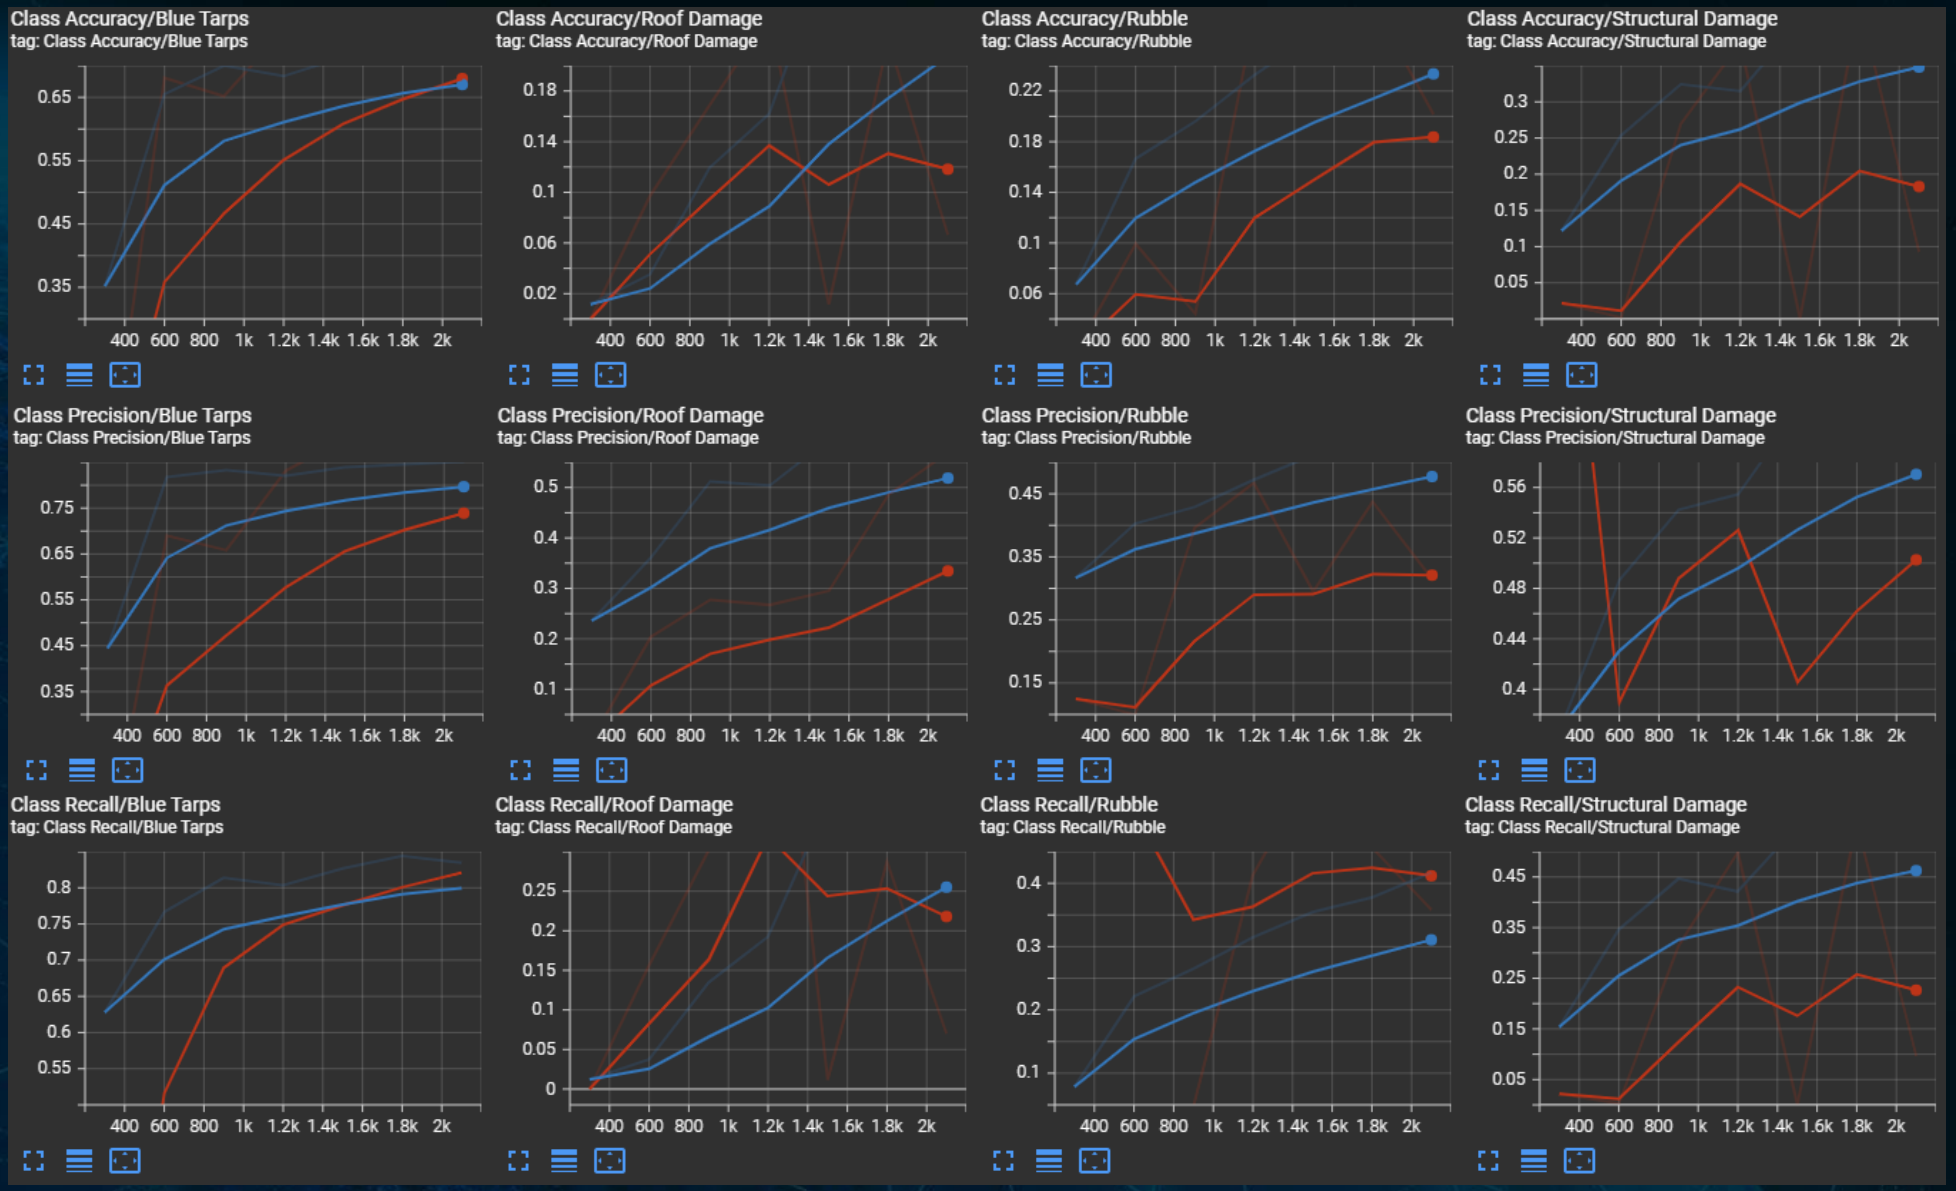

Confusion Matrices

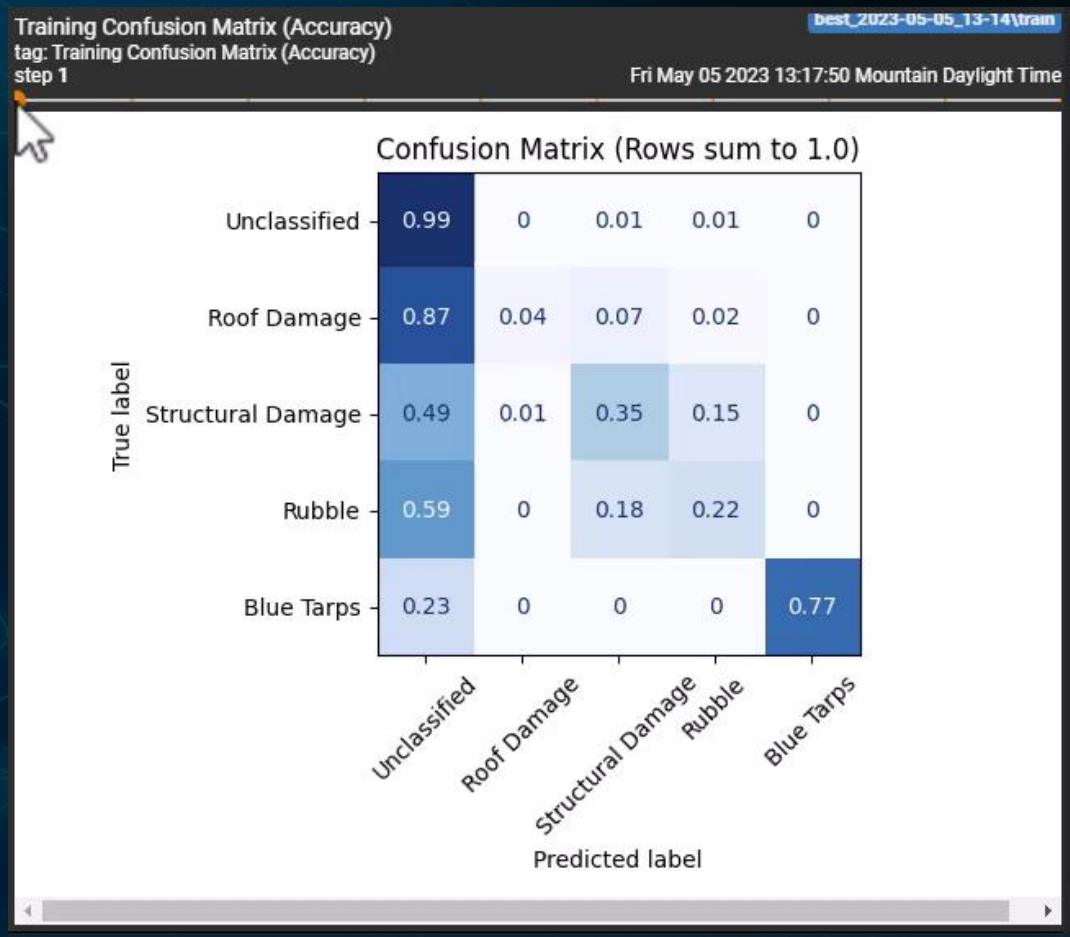

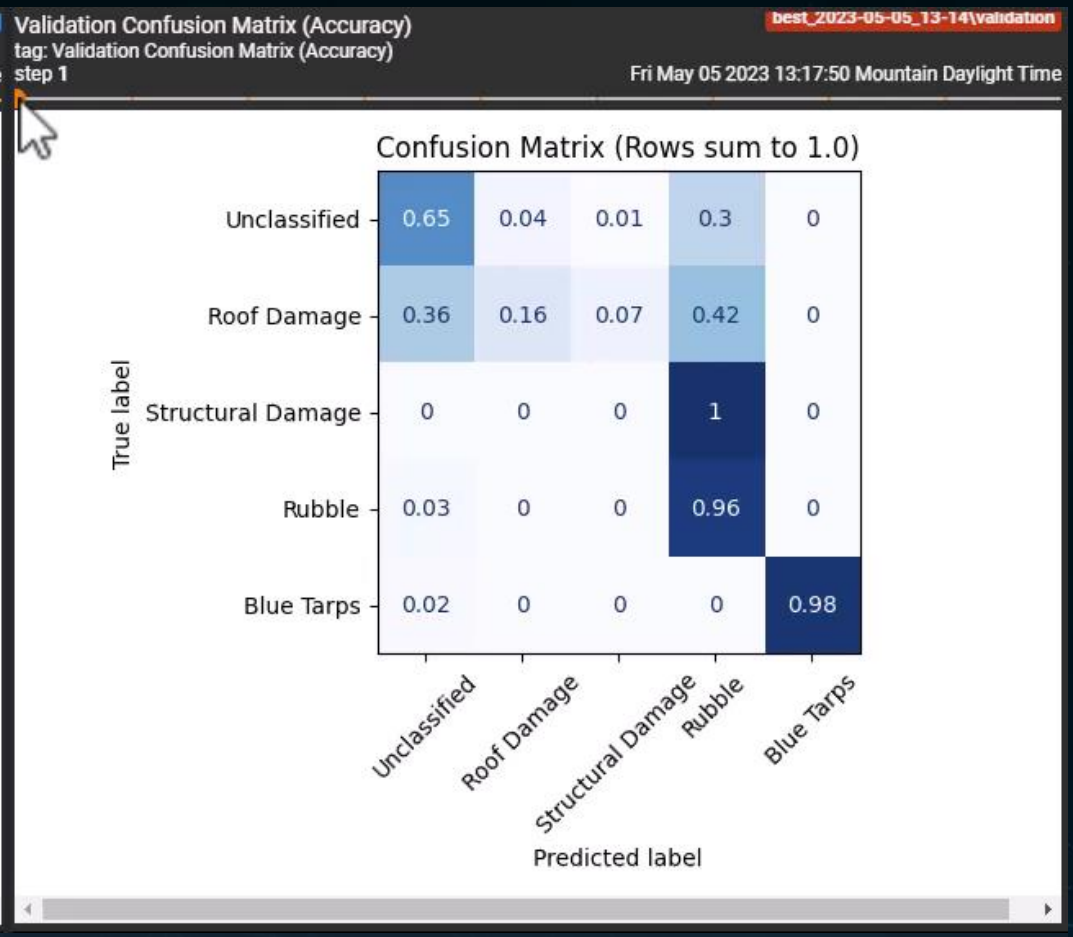

After each epoch, we generate confusion matrices for training and validation. For readability, the rows are normalized from 0 to 1 (you don't get raw class counts)

Distributions and Histograms

Two new styles of metrics: distributions and histograms give insights into the range of values for a metric in a single visual

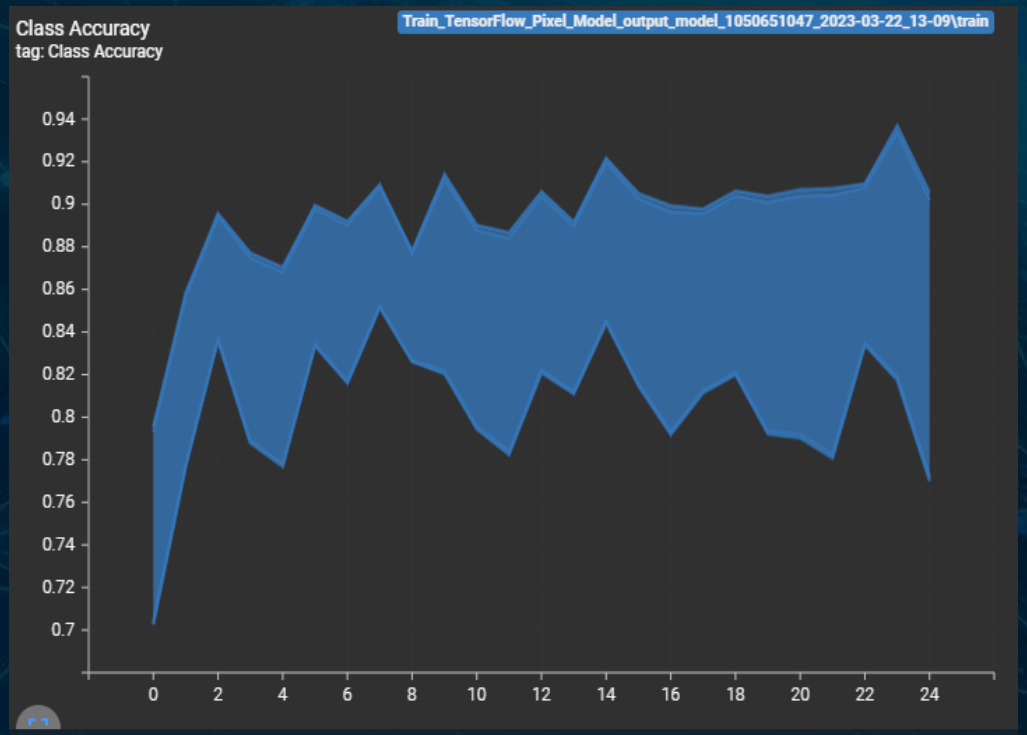

Distributions show the range of values for a given metric

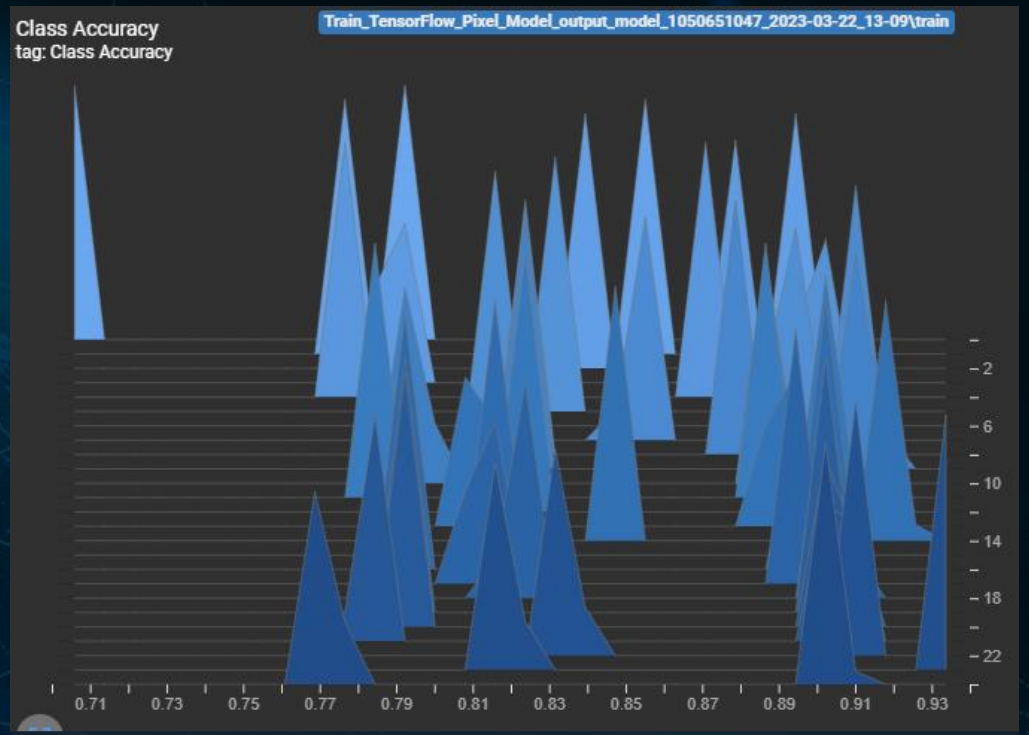

Histograms show the range of values for a metric along with the distribution of values

Distributions and Histograms

Each plot shows the similar information, as shown by the red and light blue lines in the graphs below, corresponding to extents of the values

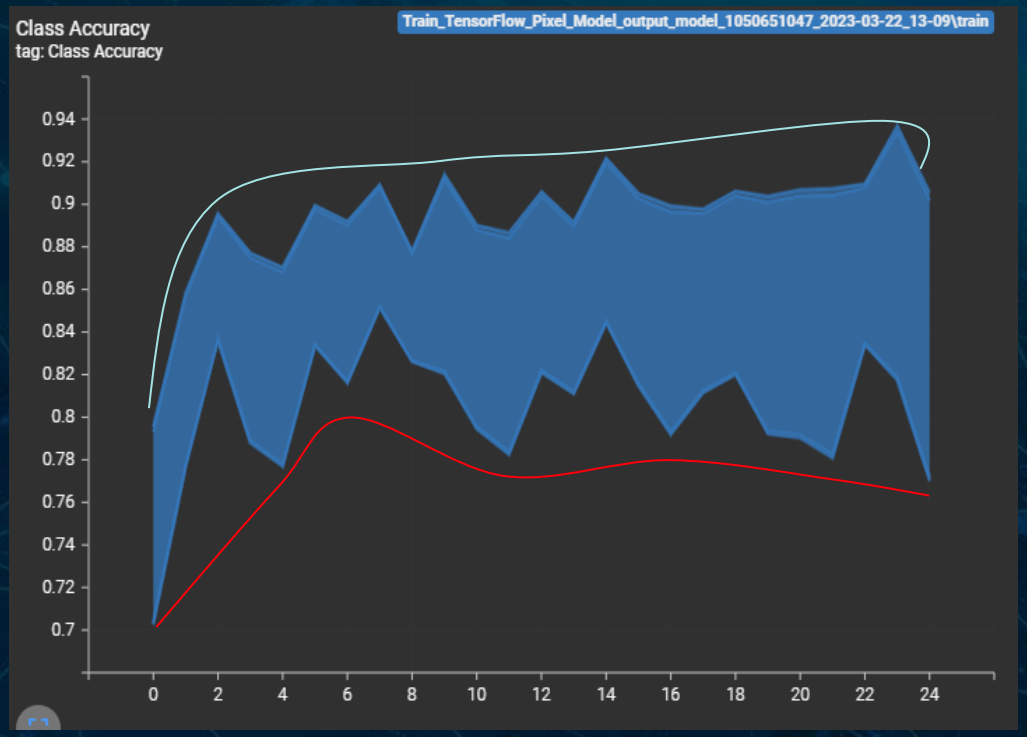

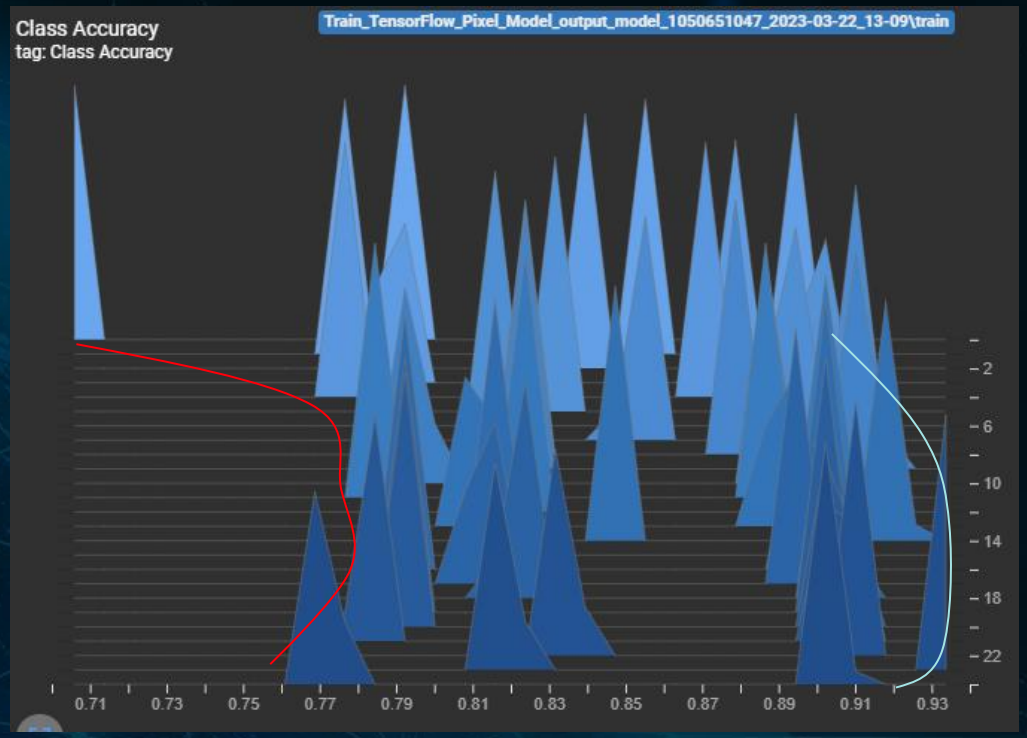

TensorBoard: Metrics vs Real-World Performance

With TensorBoard, you might be surprised by what gets reported for the accuracy, precision, or recall of different classes

Zach’s pro tip is to always process an image and have a look at the results yourself

Also keep in mind that we can use additional post-processing to clean up our results

Reminder: Cleanup Tools

RASTER

- **Classification Aggregation**
  - This task aggregates smaller class regions to a larger, adjacent region as part of the classification cleanup.
- **Classification Clumping**
  - This task performs a clumping method on a classification image. This operation clumps adjacent similar classified areas using morphological operators.
- **Classification Sieving**
  - This task removes isolated classified pixels using blob grouping.
- **Classification Smoothing**
  - This task removes speckling noise from a classification image. It uses majority analysis to change spurious pixels within a large single class to that class.
- **Convert Classification to Vector Shapefile**
  - Converts a classification image to a shapefile
- **Class Activation to Polyline Shapefile**
  - Converts a linear mask of features to a polyline, where the polyline represents the centerline (i.e polygon of roads converted to the centerline of roads)

VECTOR

- **Filter Vector**
  - Remove geometry from a vector based on area (pixels and normal units such as meters
- **Smooth Vector**
  - Reduce the number of vertices in a vector
- **Vector Record to Bounding Box**
  - Convert polygons to a bounding box (extent-based or oriented with the shape of the features)
- **Vector Records to Centroid**
  - Convert polygons to a centroid
- **Postprocess Classification Vector**
  - Takes bounding boxes from object detection and allows you to do additional filtering to cleanup results

New Architecture: Pixel Segmentation

- Creating a model for segmentation now includes two architectures:
  - SegUNet++
  - SegUNet
- Differences with the new model:
  - Better ability to detect features with smaller training sessions
  - Generates less noise and fewer spurious (i.e. standalone) pixels
  - About 10% faster
  - Slightly smaller network size

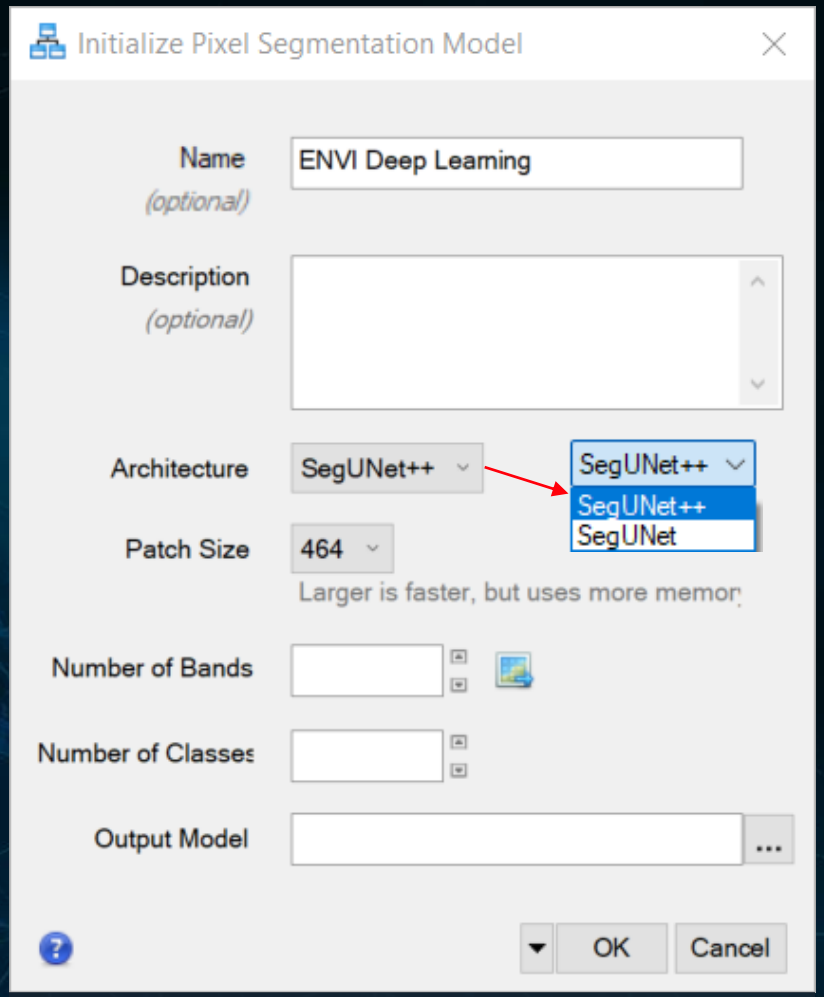

Architecture selection in the model initialization dialog

Model Comparison: Limited Training

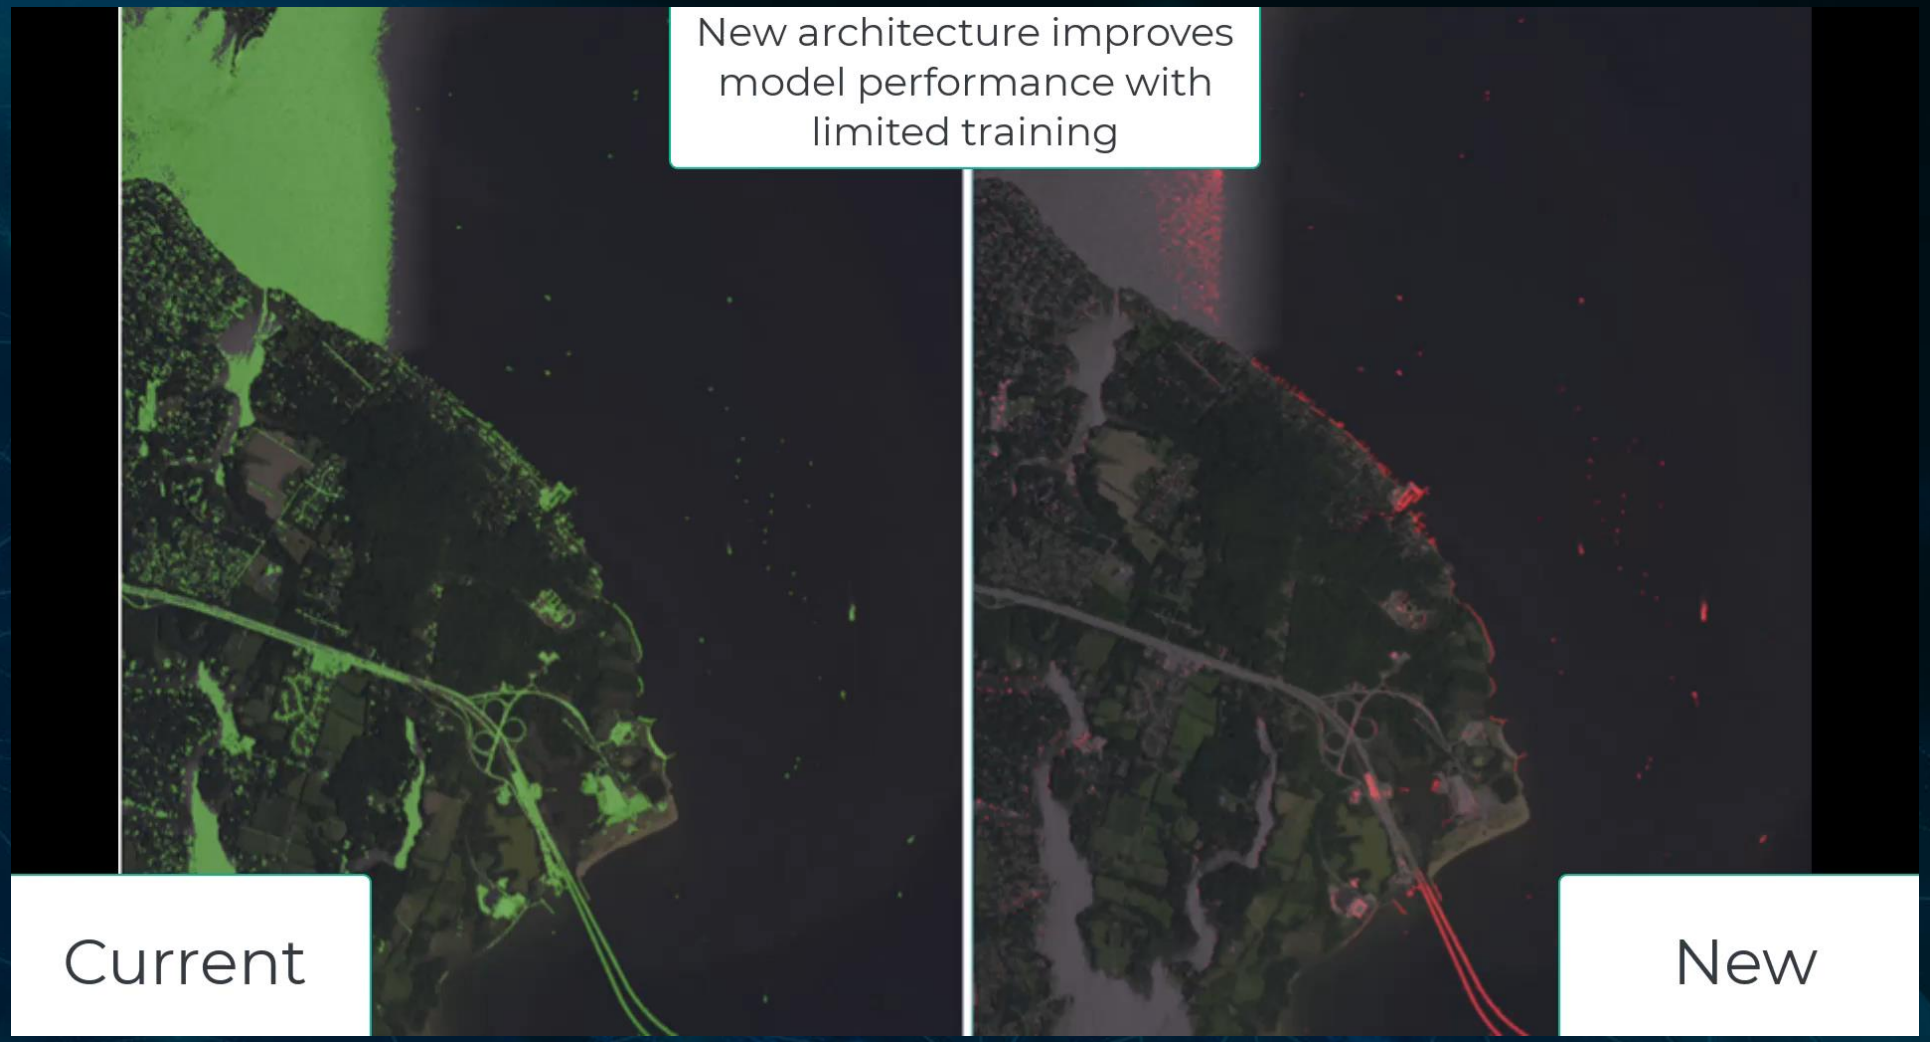

Tornado Damage 1

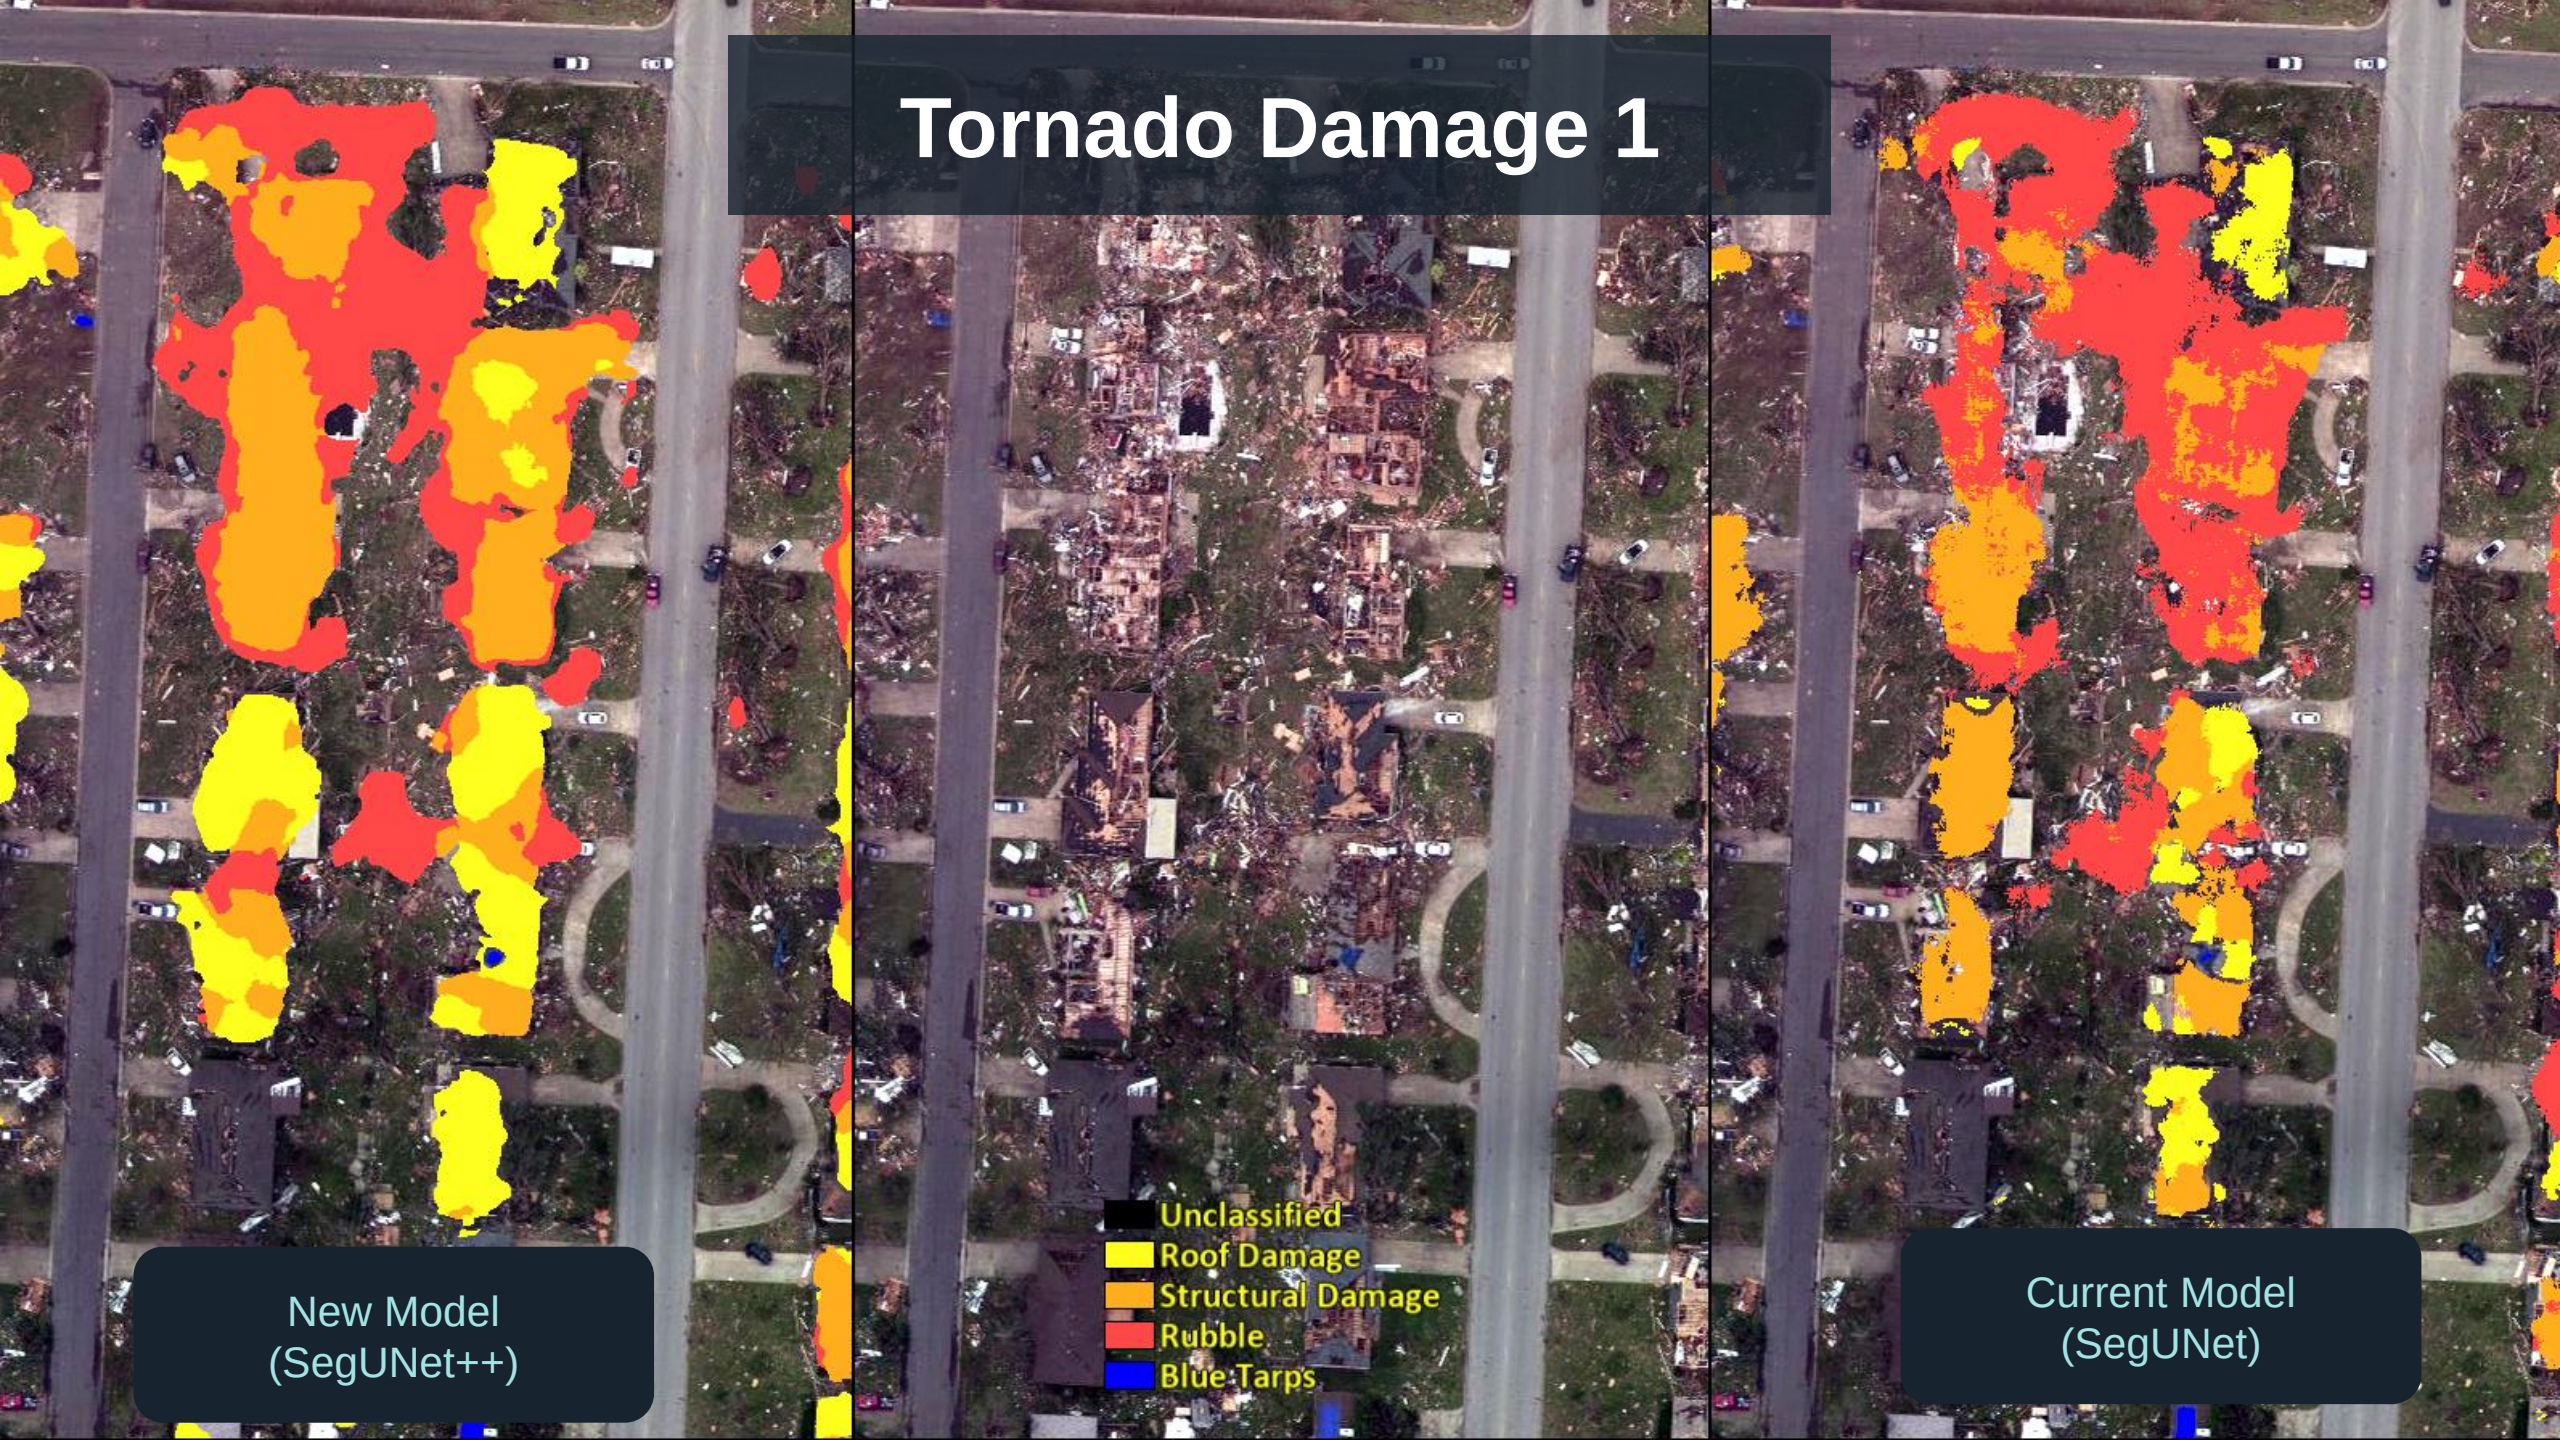

Tornado Damage 2

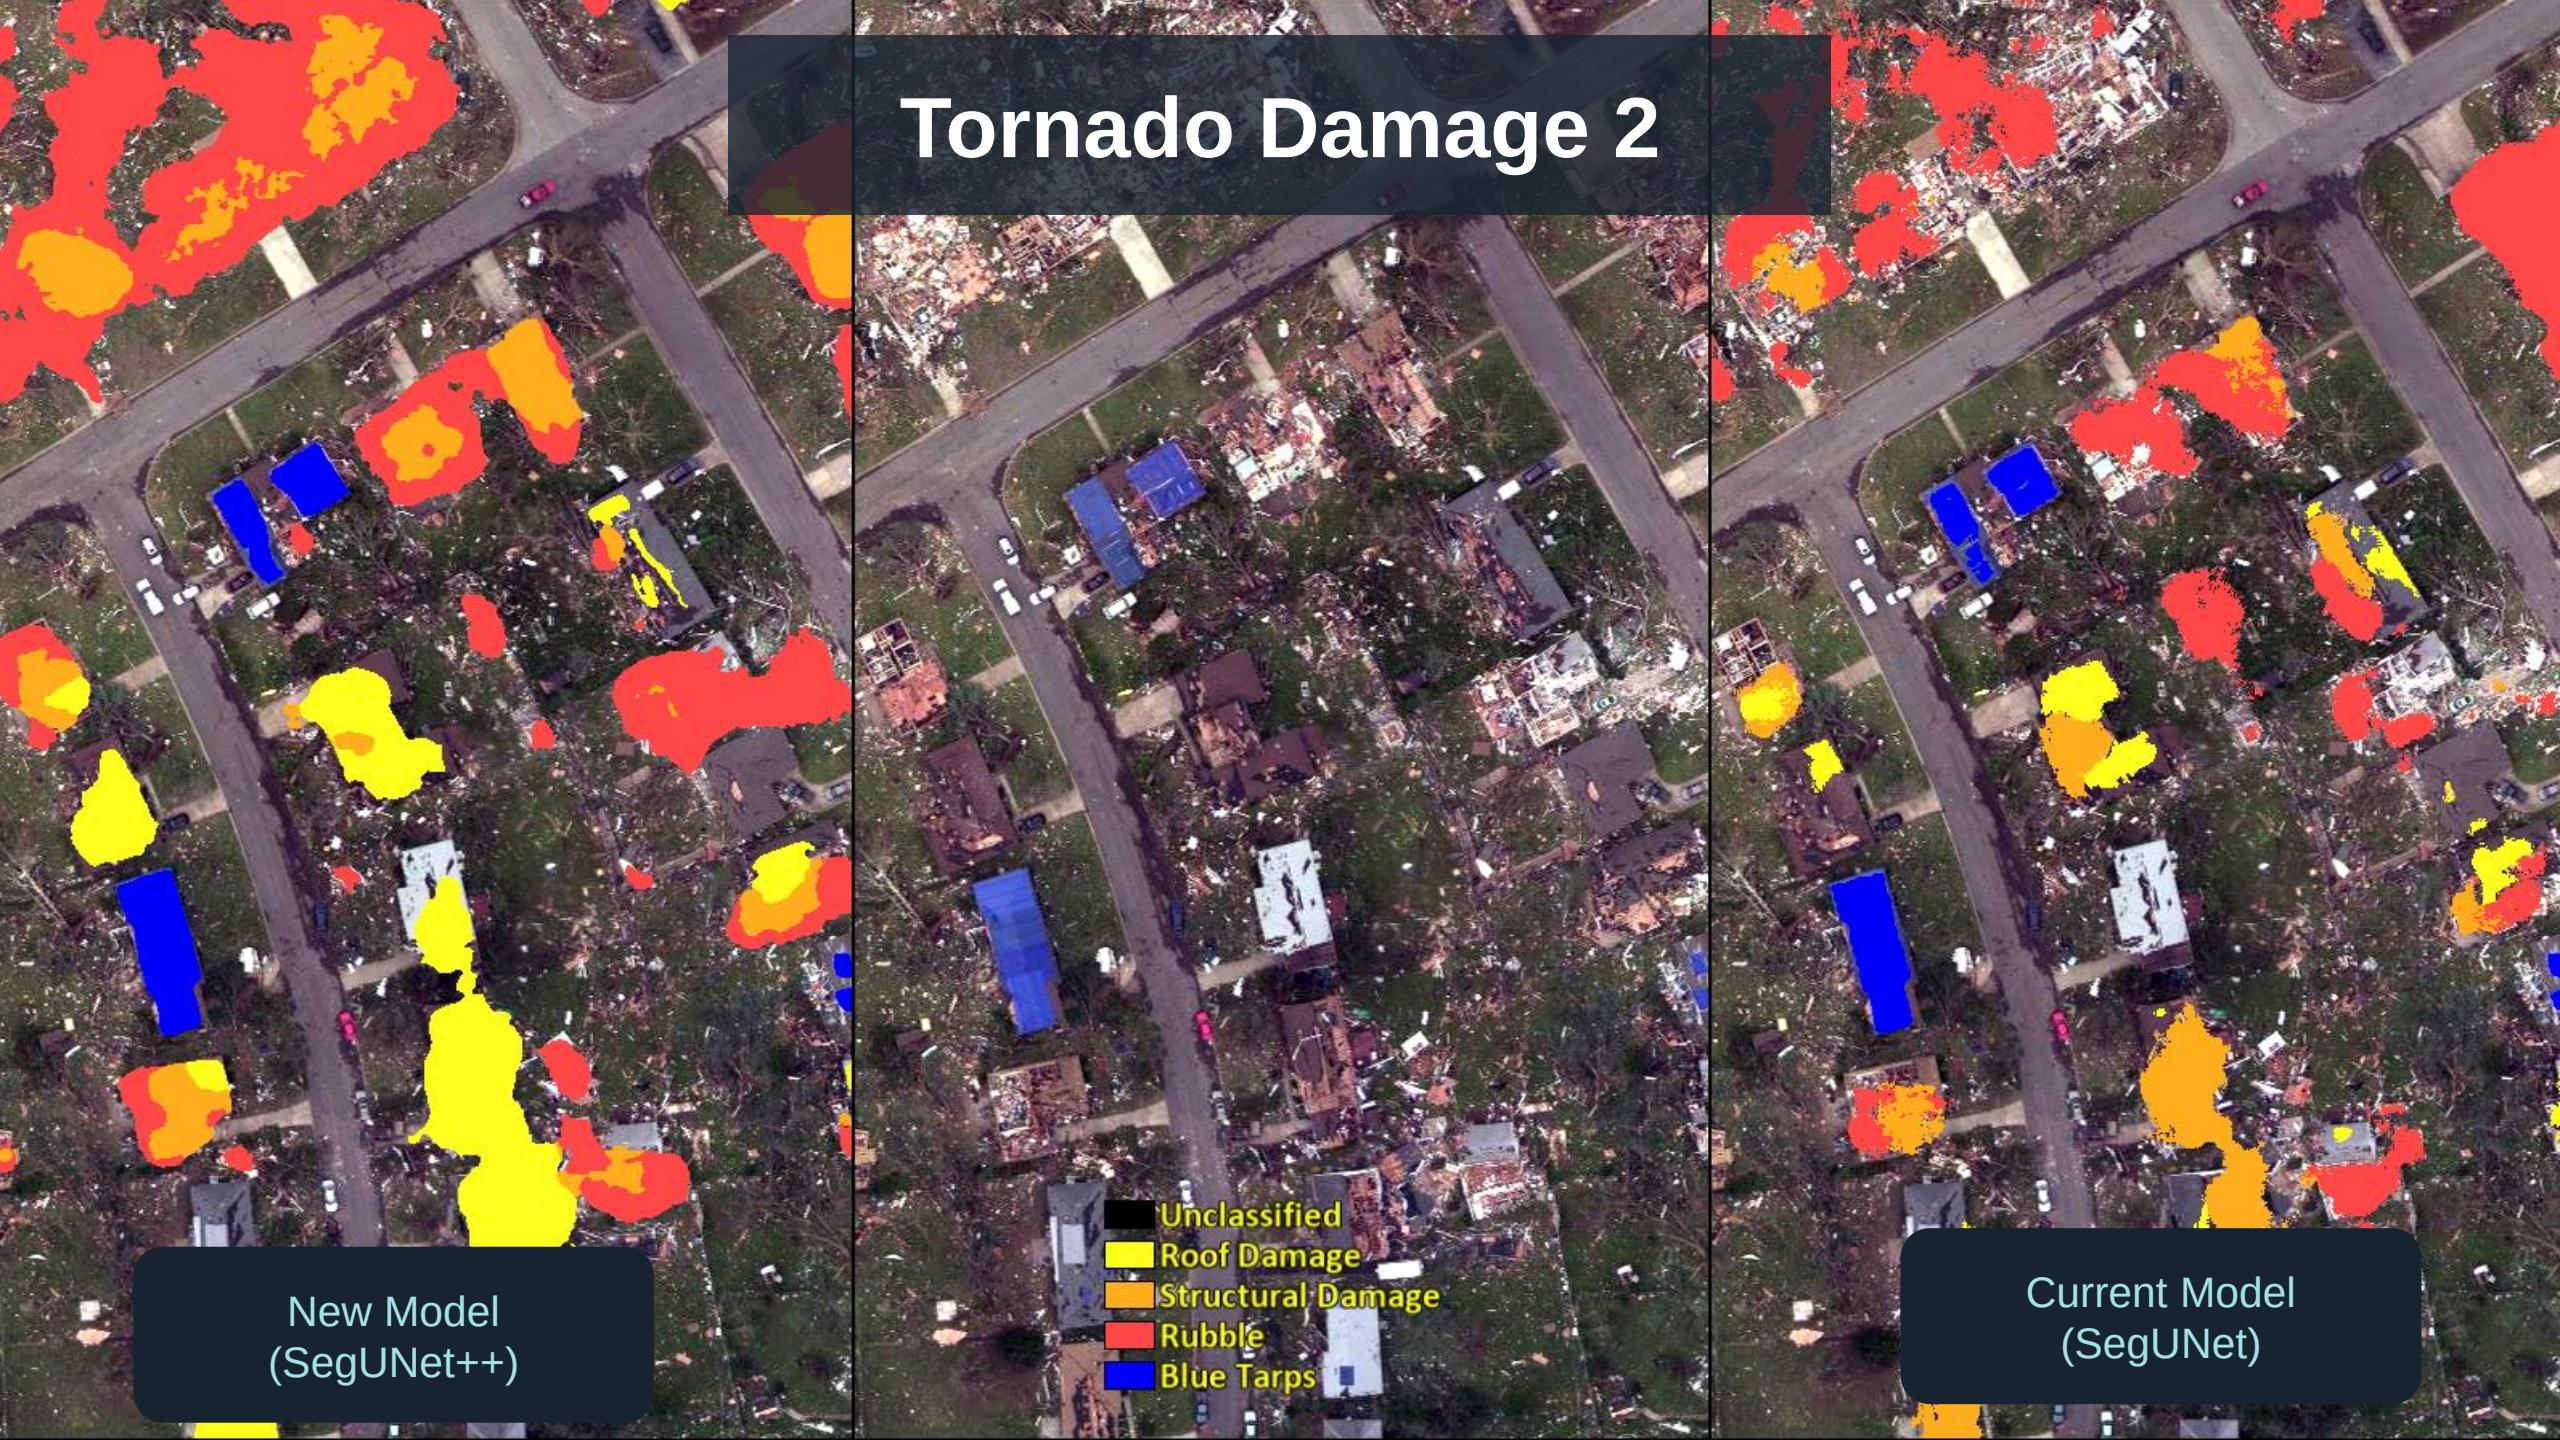

Crosswalks 1

Crosswalks 2

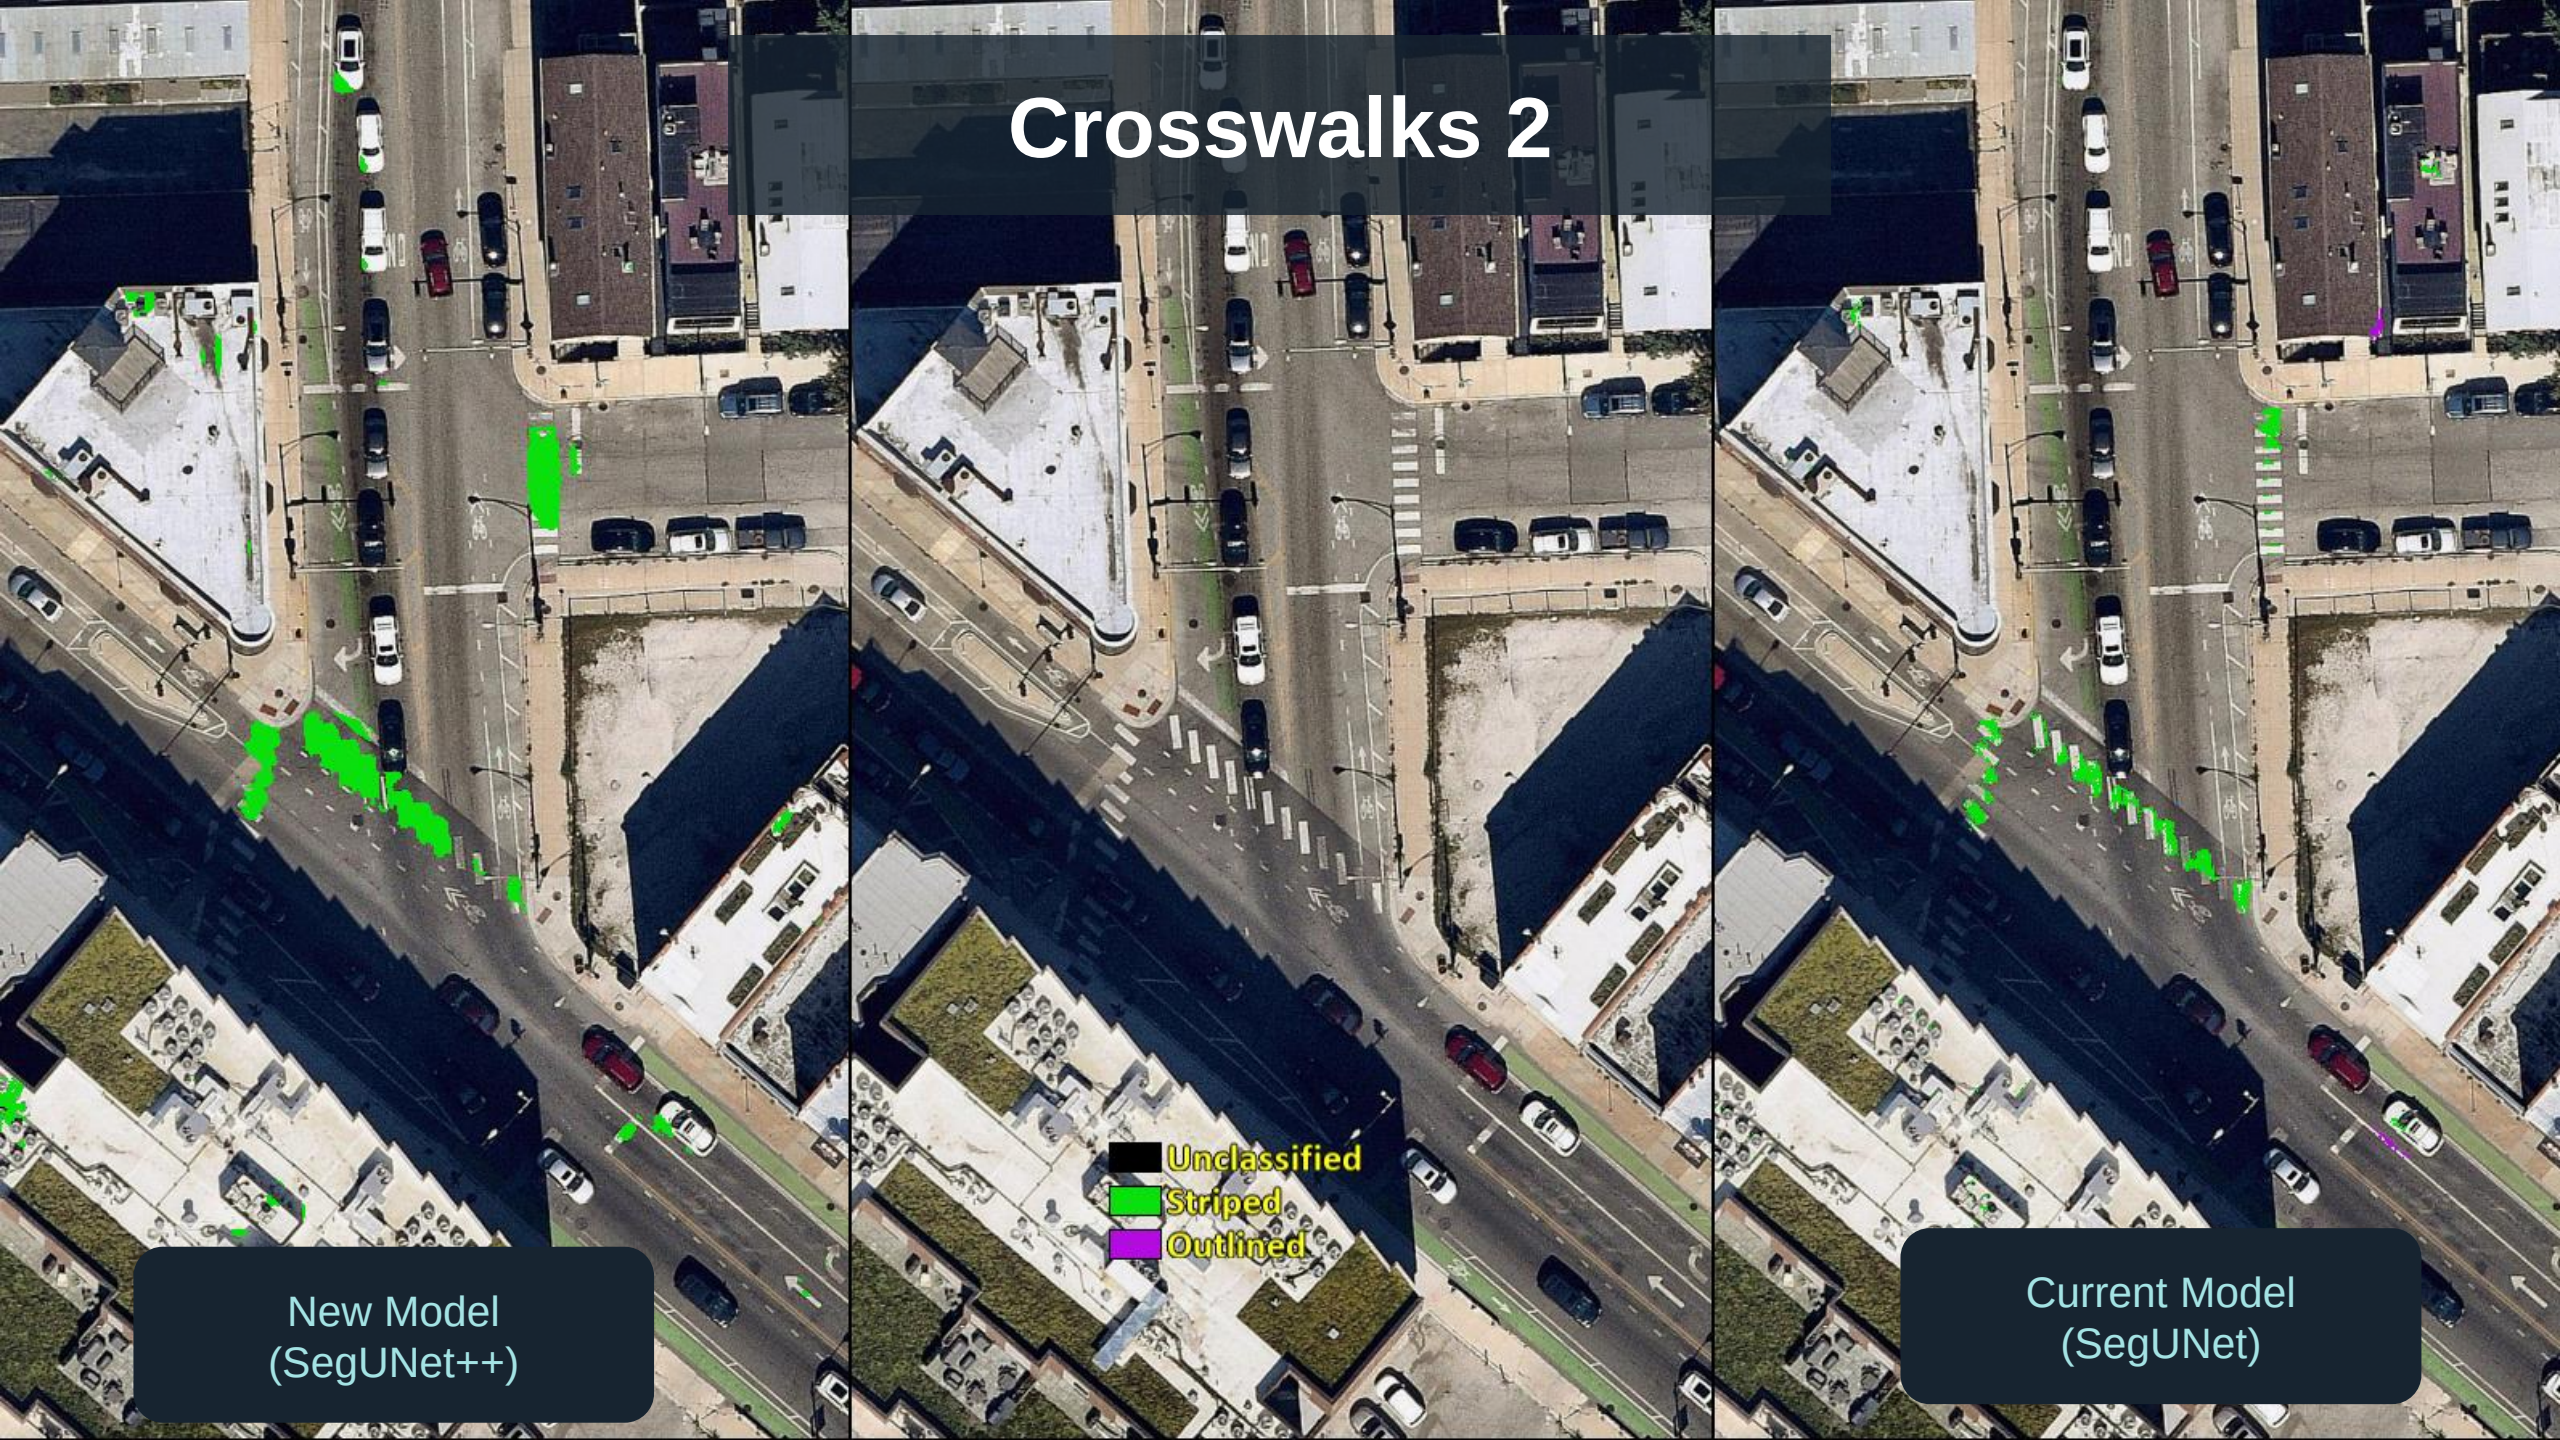

Crosswalks 3

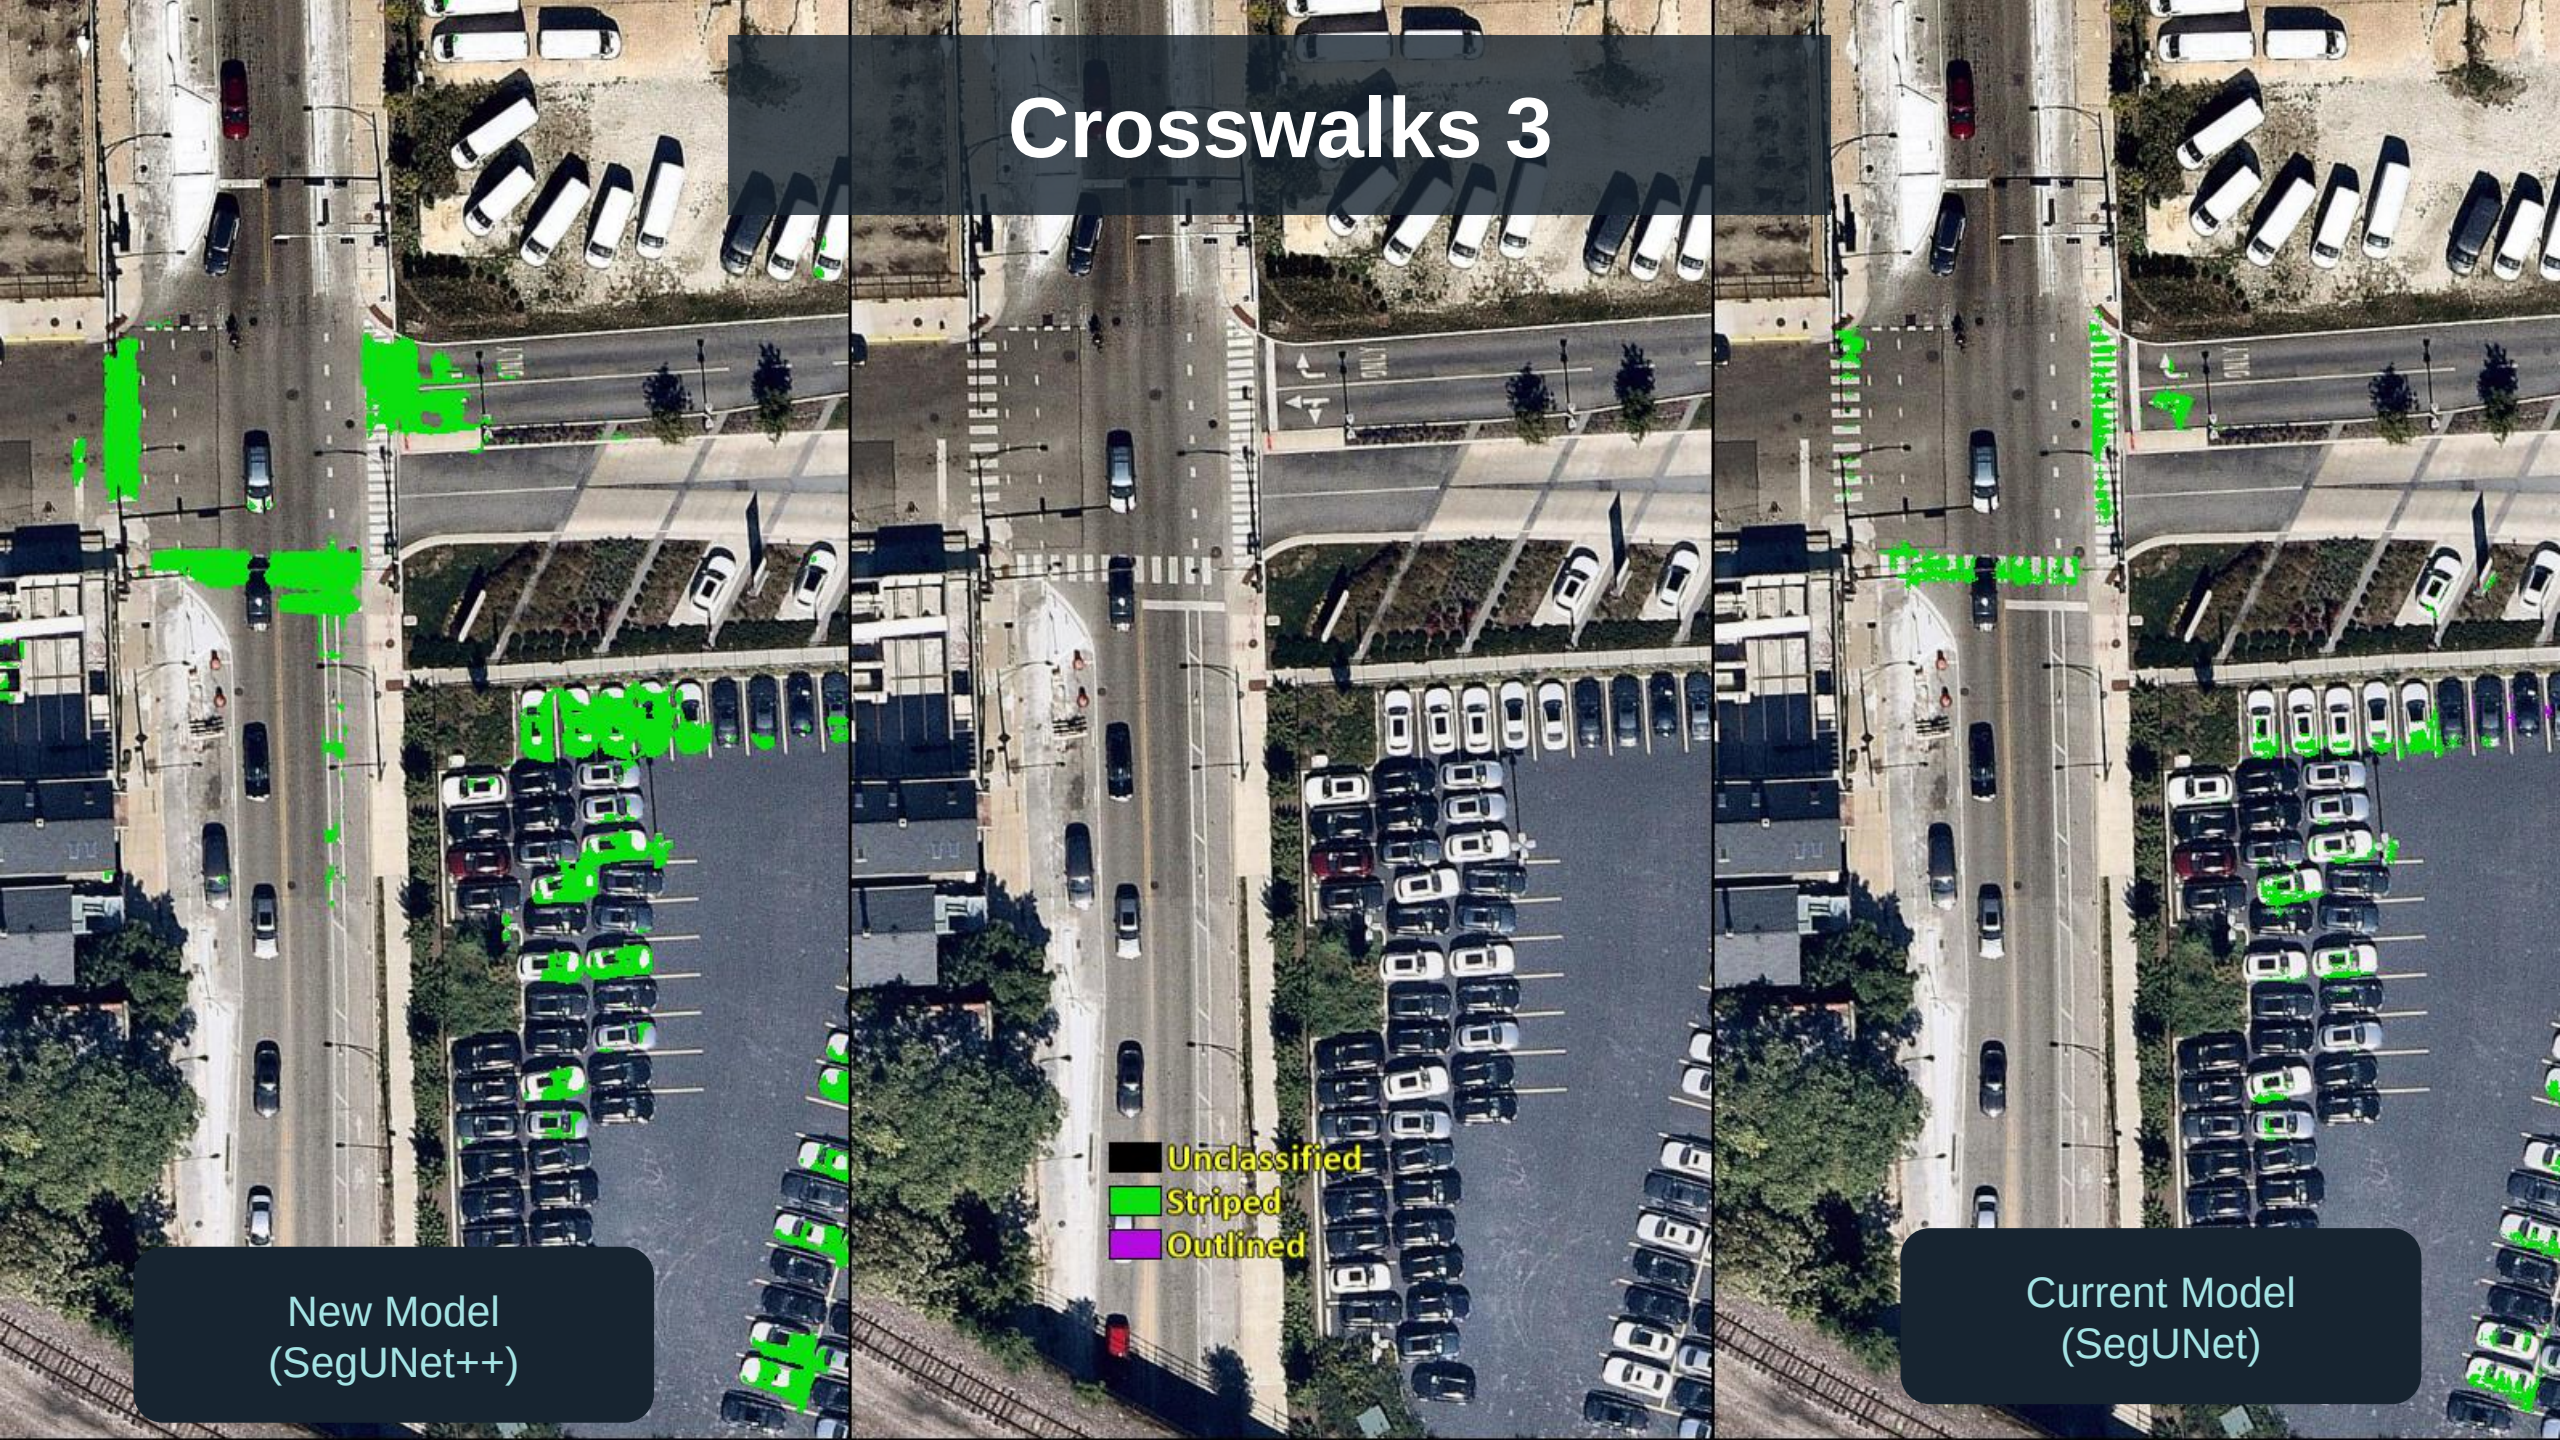

Crosswalks 4

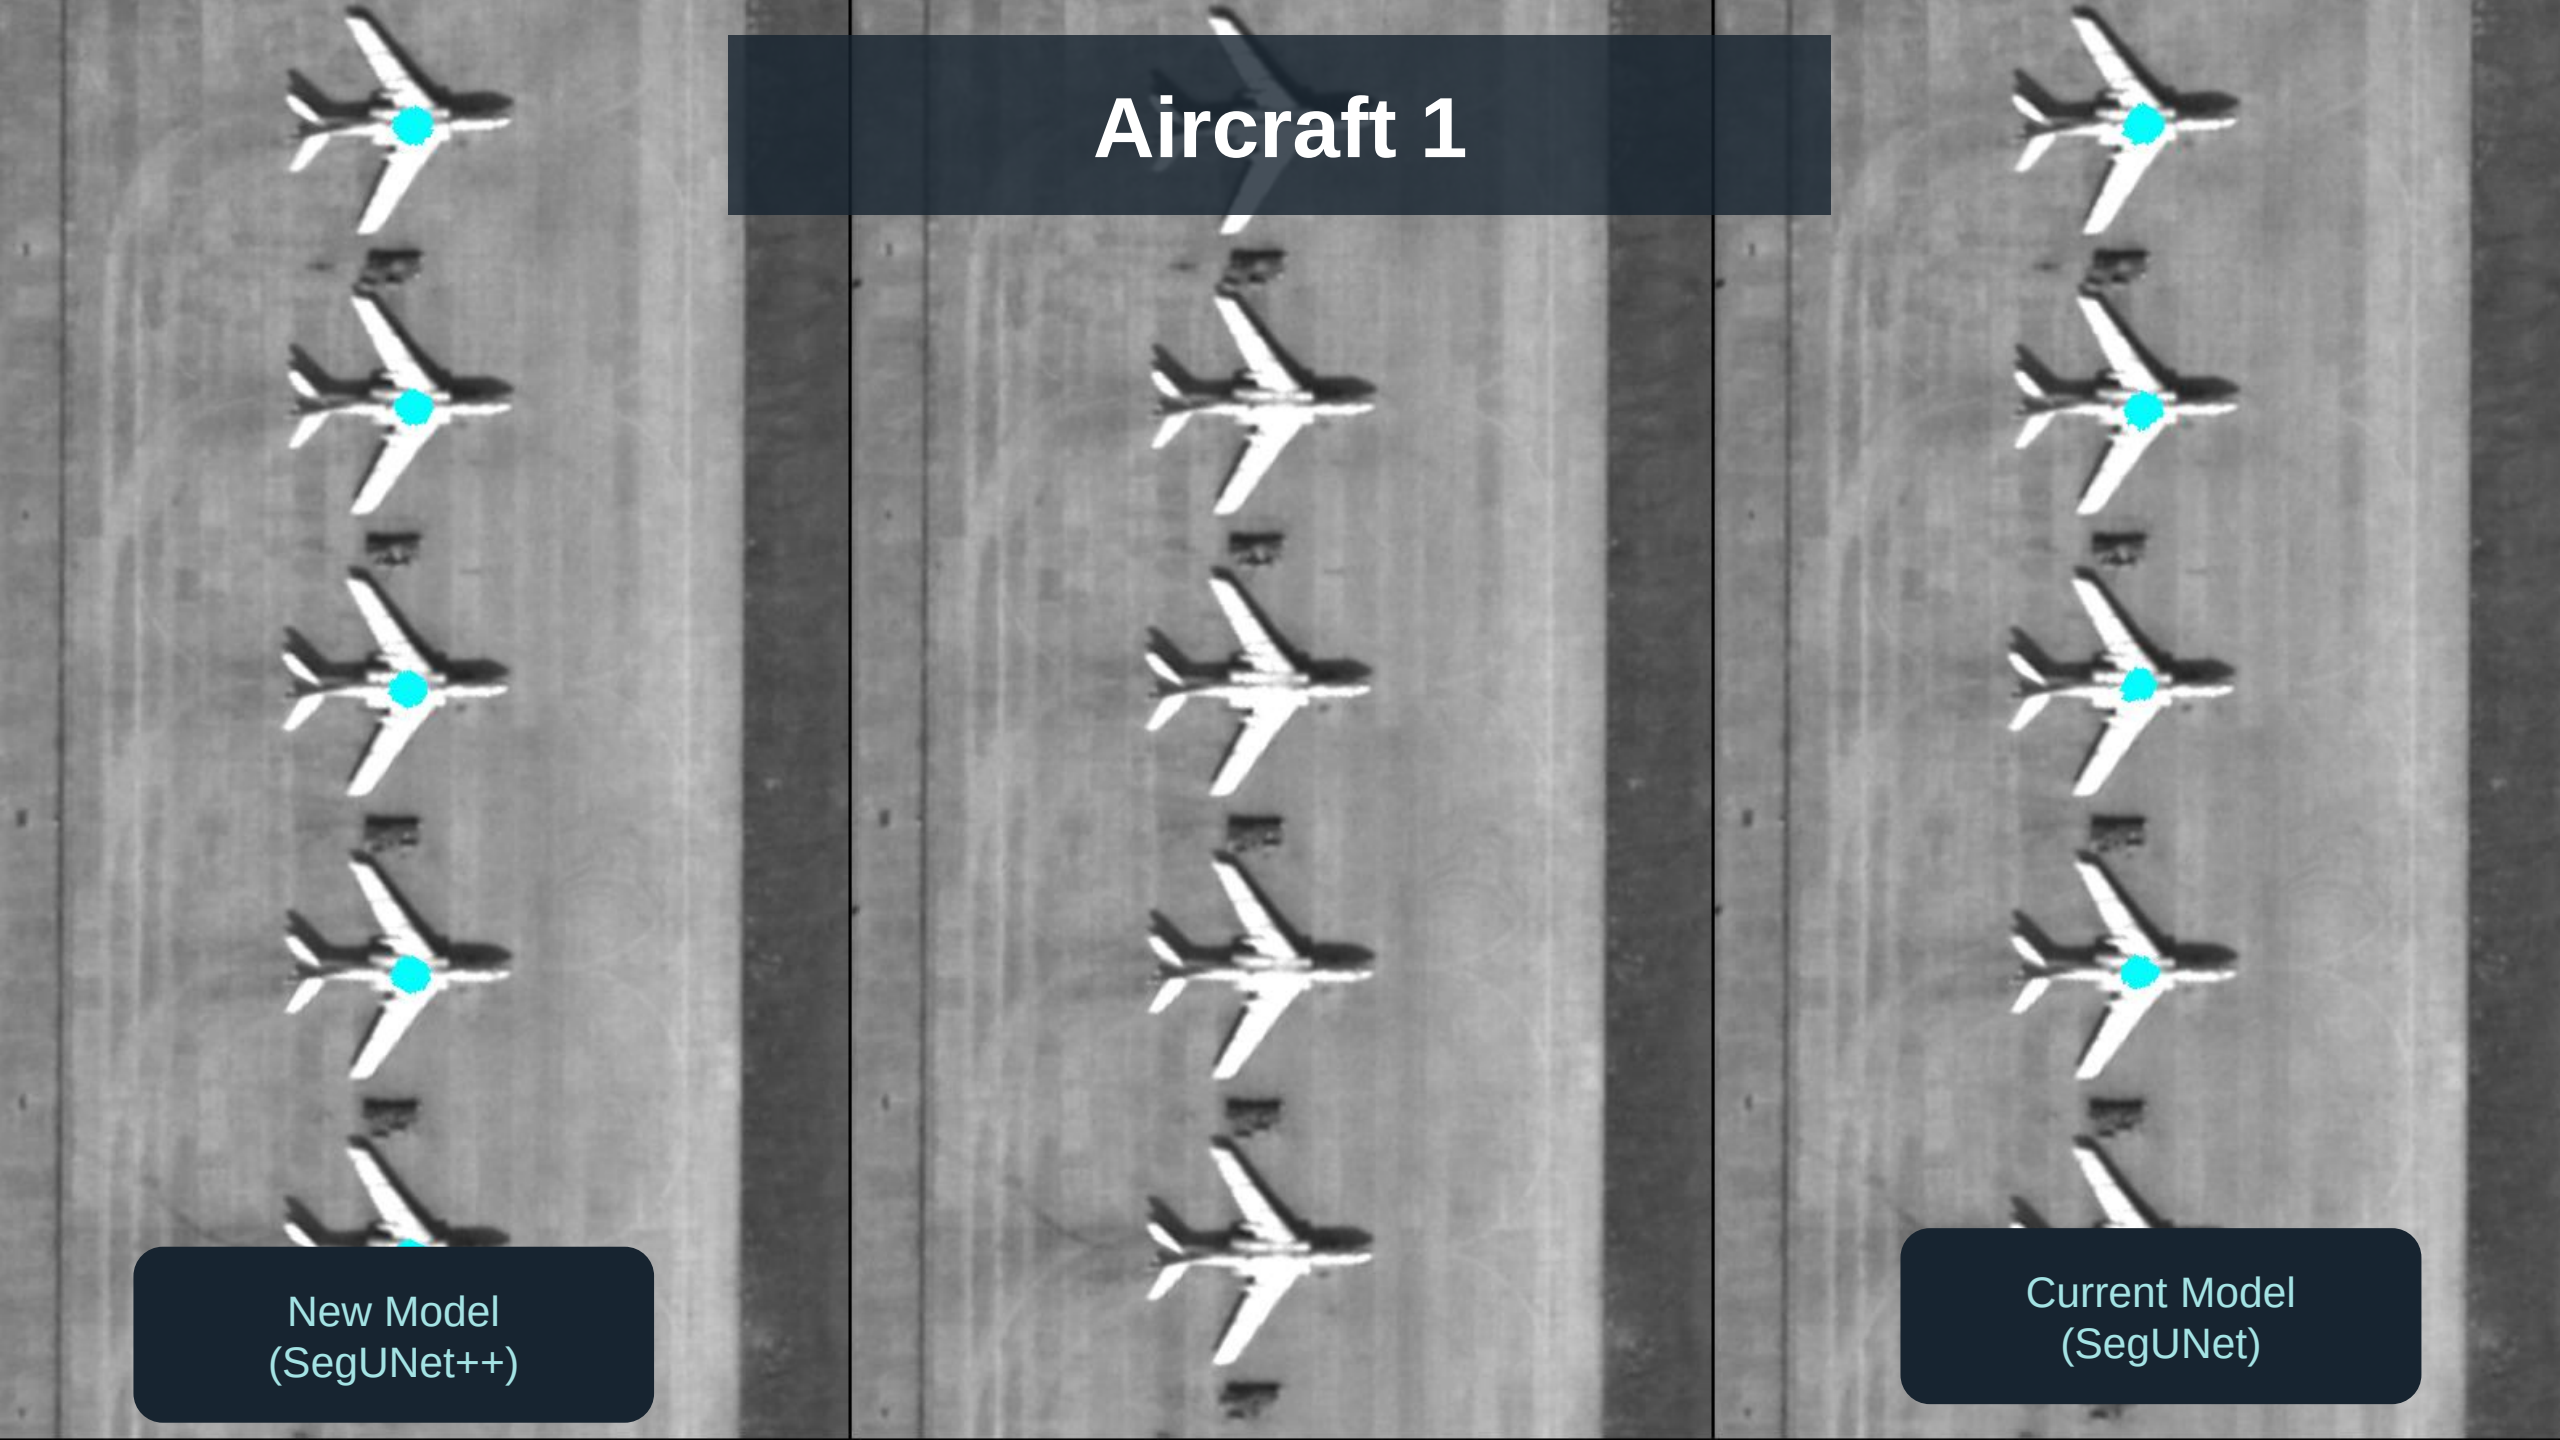
Aircraft 1
New Model
(SegUNet++)
Current Model
(SegUNet)

Aircraft 2

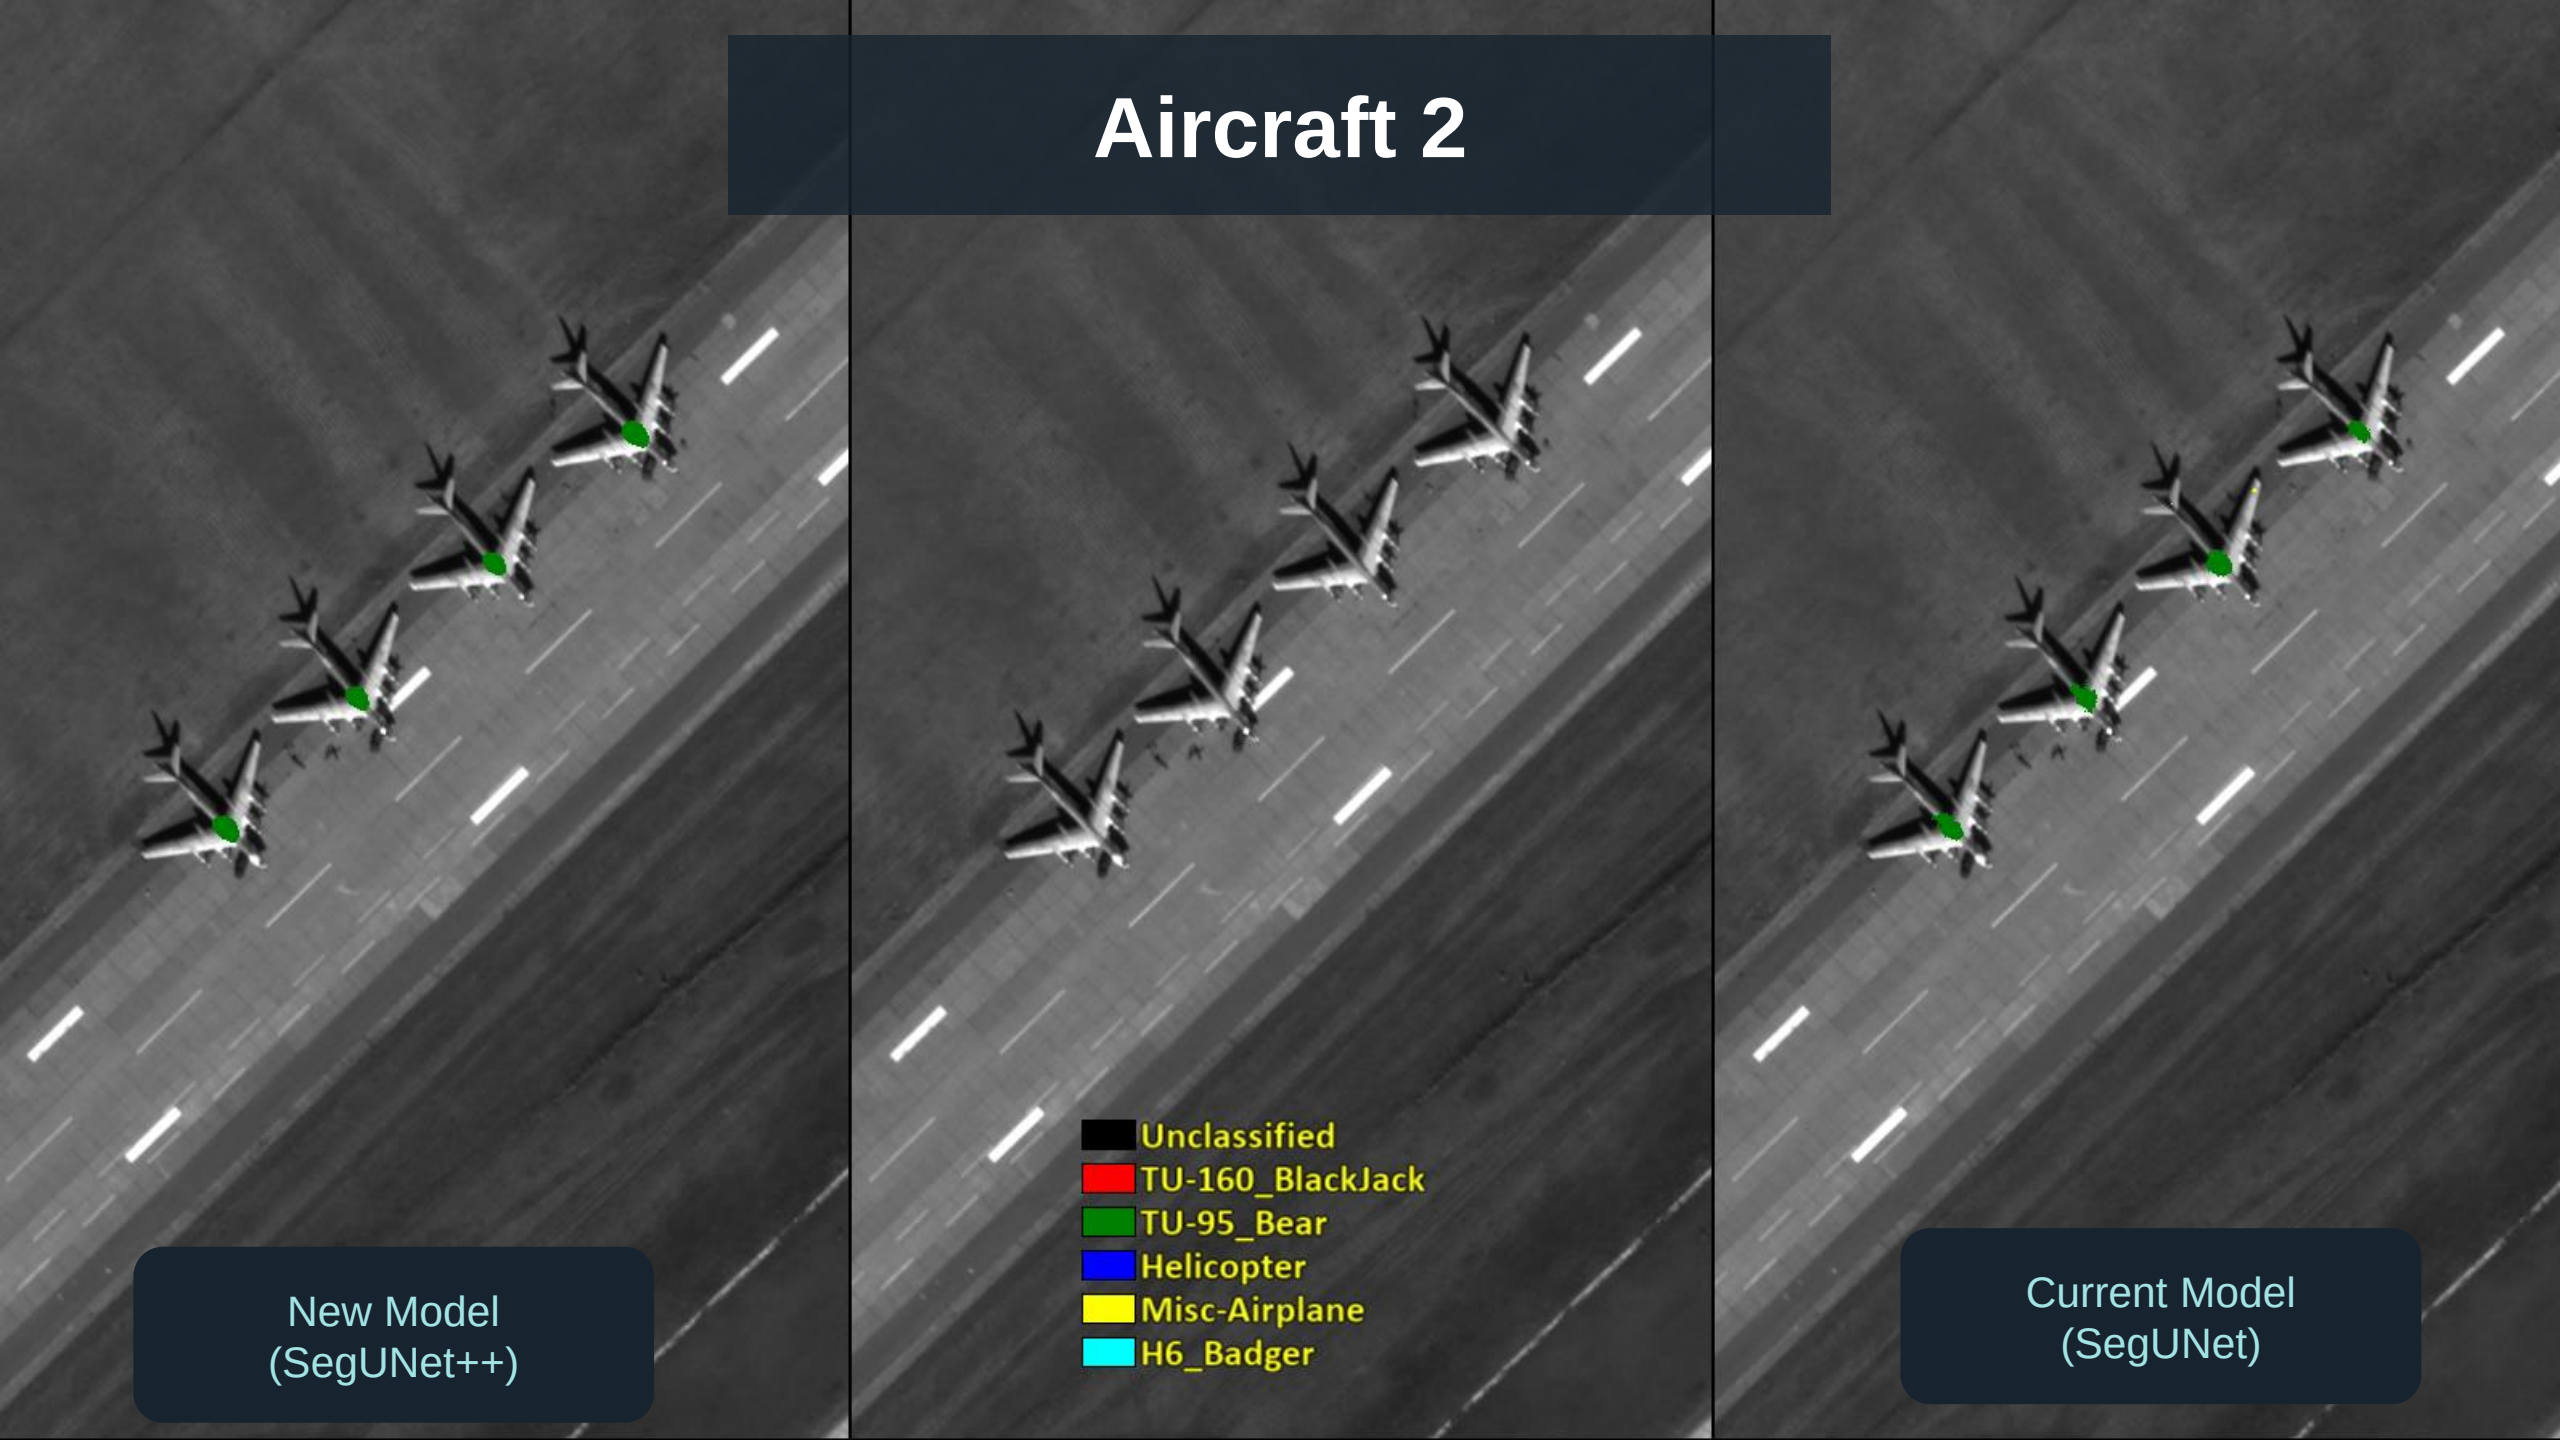

New Model
(SegUNet++)

Current Model
(SegUNet)

Aircraft 3

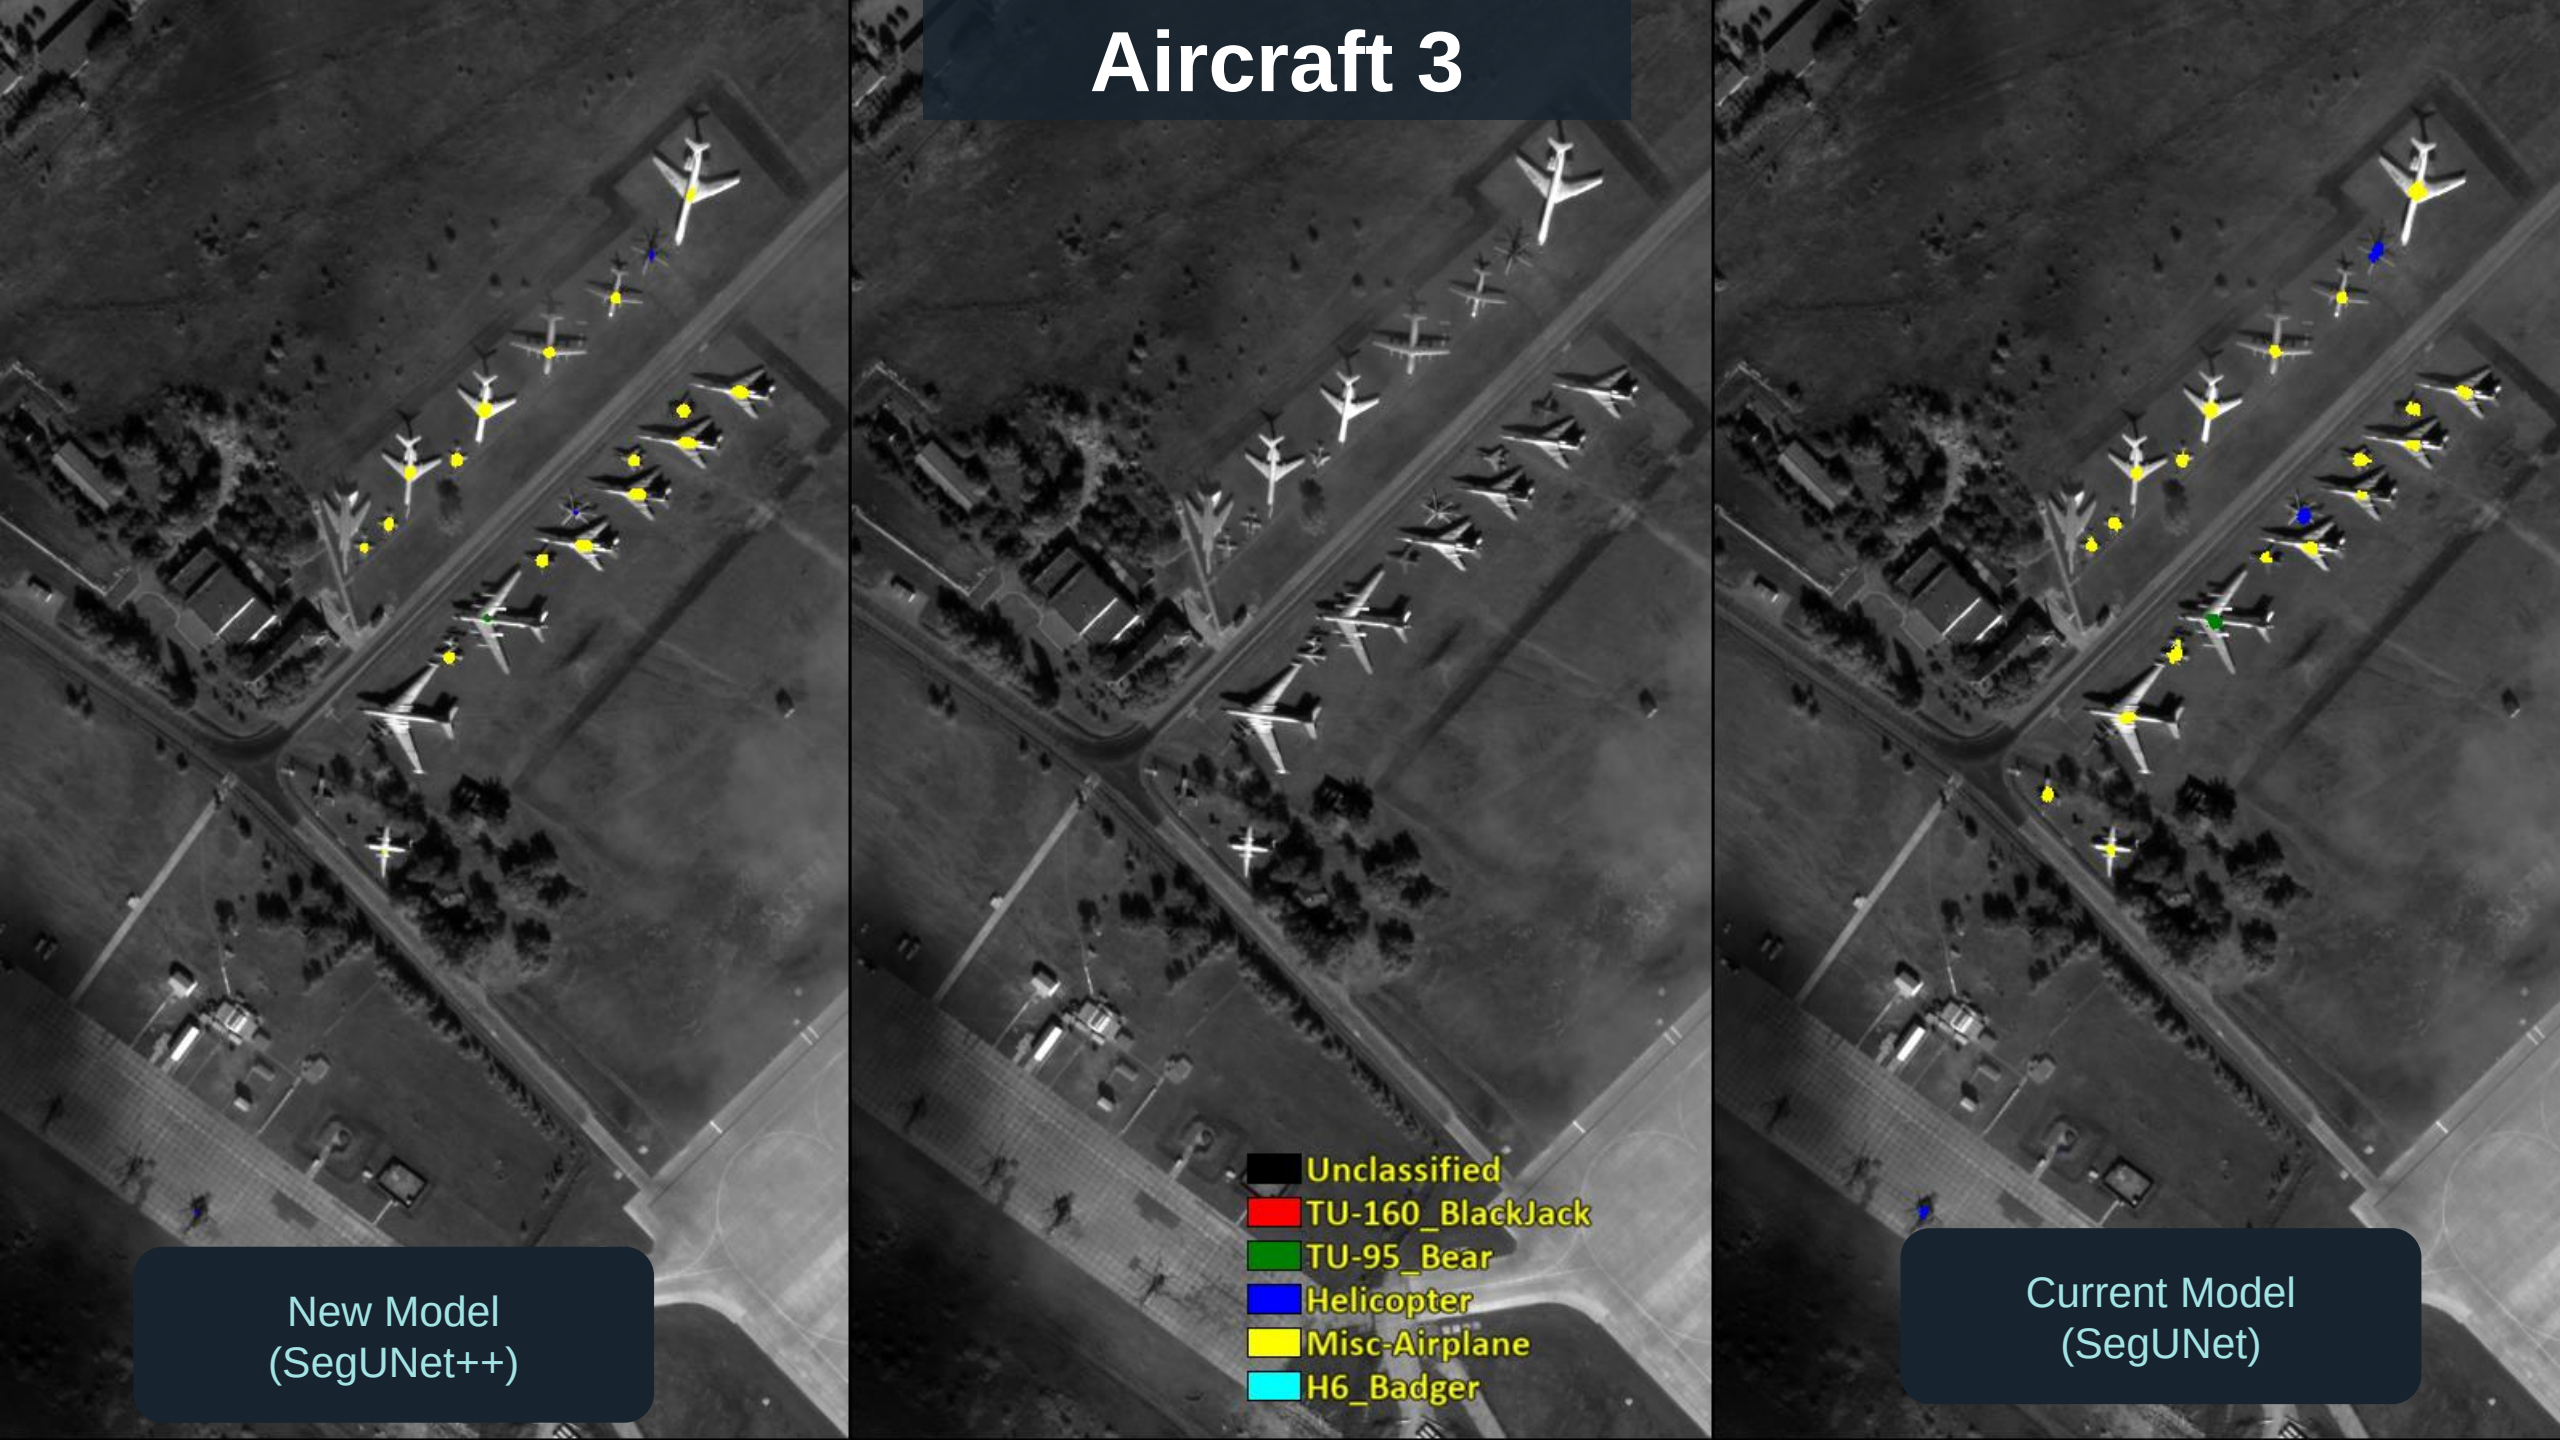
Unclassified
TU-160_BlackJack
TU-95_Bear
Helicopter
Misc-Airplane
H6_Badger
New Model
(SegUNet++)
Current Model
(SegUNet)

BREAKING CHANGE

- The legacy “Single Class Model” can no longer be created in ENVI Deep Learning
- This is an old concept and no longer applicable to the current state of the application
- You can still process imagery using legacy single-class models

Progress Reporting

- Progress dialogs in ENVI when using ENVI Deep Learning were not always helpful or accurate
- They would also show “Not Responding” while we were waiting for message updates
- TensorBoard was more useful to track progress because you would get more frequent updates

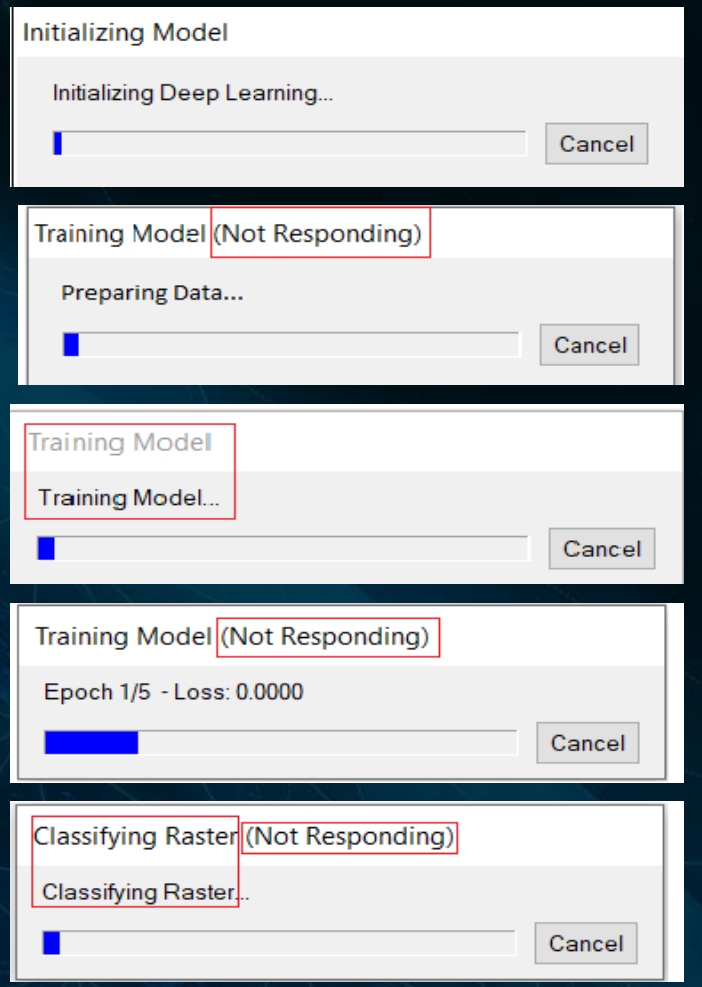

Progress Reporting: New User Experience

- Now, all dialogs provide much more insight into what is going on
- Many more progress updates provided for training to give you progress through each epoch
- Different types of deep learning now present the same progress dialogs

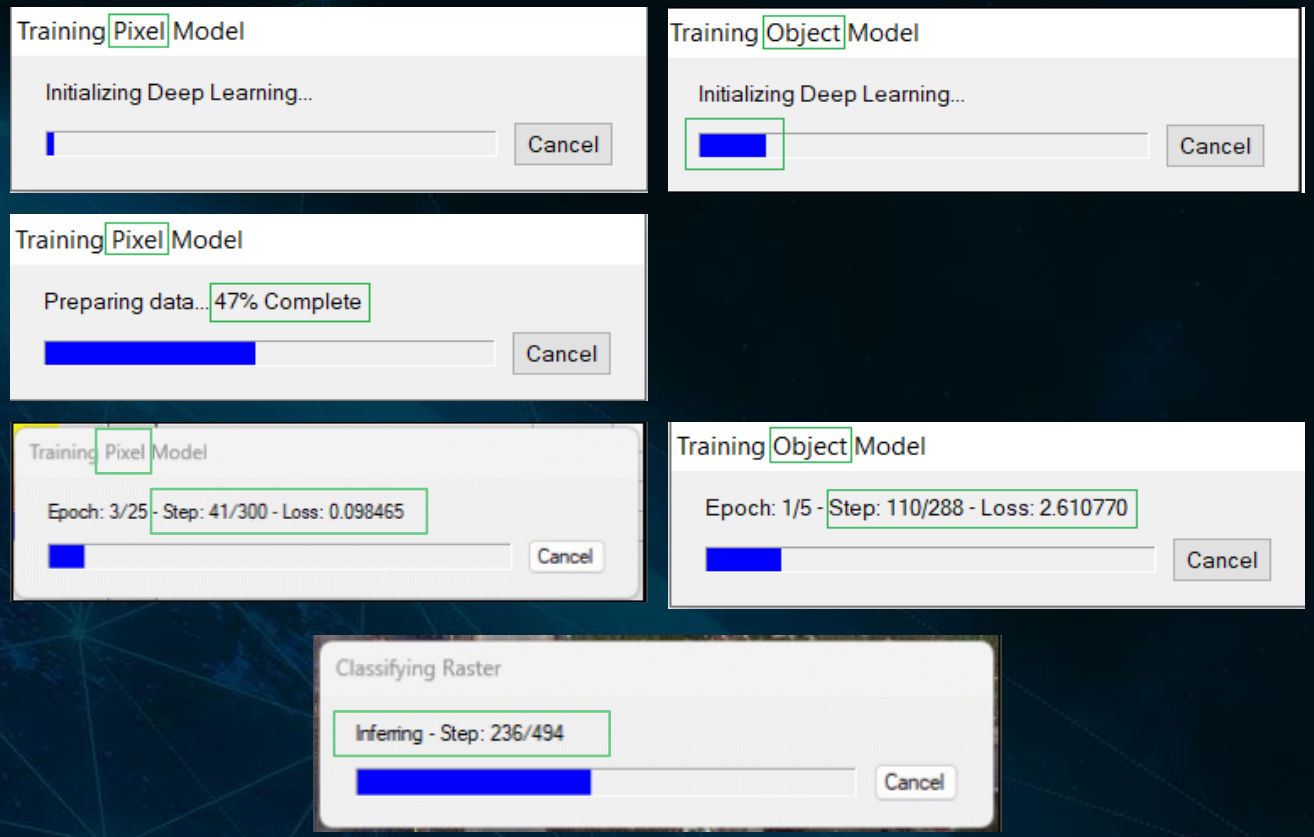

Screenshots showing new dialogs and differences from previous versions

MACHINE LEARNING

Machine Learning: Normalization

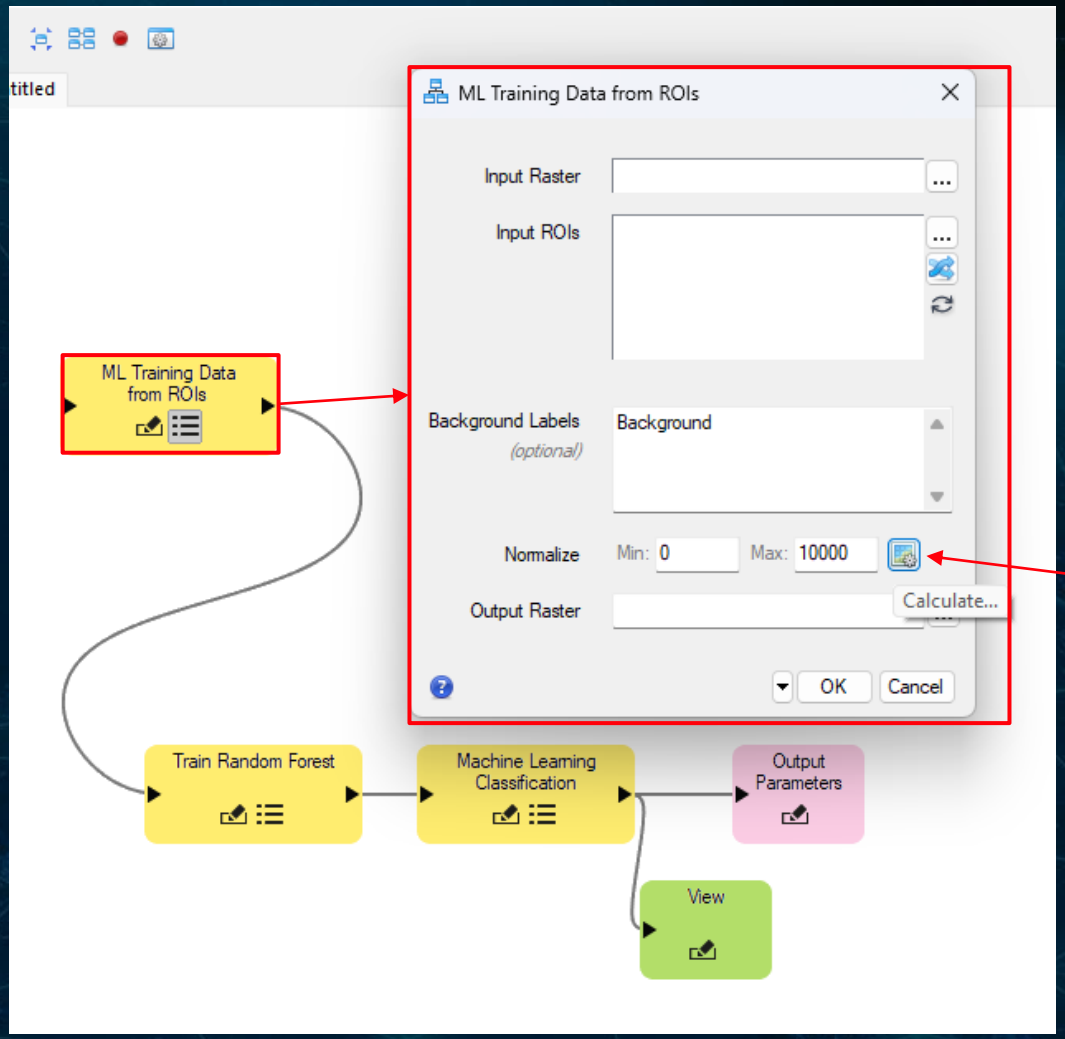

We have improved the automatically generated normalization statistics for machine learning to work better with hyperspectral datasets.

However, if you know the range of your data (from 0% to 100% surface reflectance) it is recommended to use those values

Updates: Random Forest and Extra Trees

- New parameters for both:
  - **Max features** controls the number of features used to find the best split
  - **Custom max features** lets you manually specify the number of features used for the best fit
- Random forest only:
  - **OOB score** tells ENVI Machine Learning to calculate and return the OOB (Out-of-Bag) score when it generates the model

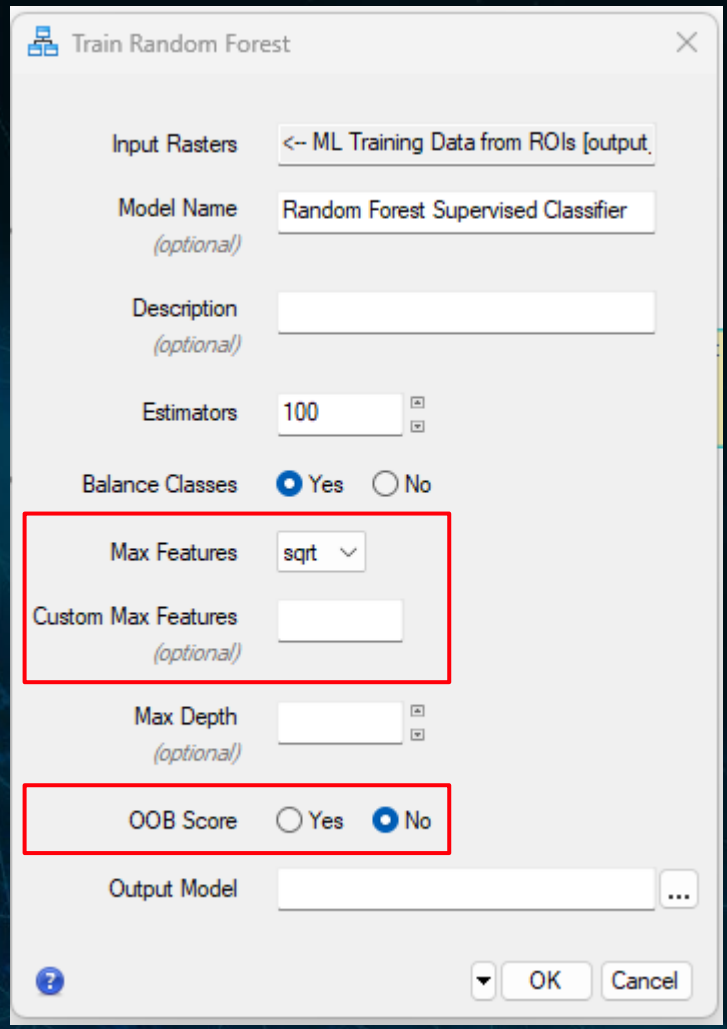

Using OOB Score

```
;+ get training data from ROIs and raster
trainingDataTask = ENVITask('MLTrainingDataFromROIs')
trainingDataTask.input_raster = raster
trainingDataTask.input_roi = rois
trainingDataTask.normalize_min_max = normTask.normalization
trainingDataTask.execute

; track the OOB errors
oobs = list()

; get number of estimators
foreach val, [15 : 200 : 5] do begin
  ; print some progress
  print, `Number of estimators: ${val}`

  ; training task
  rfTask = ENVITask('TrainRandomForest')
  rfTask.input_rasters = trainingDataTask.output_raster
  rfTask.num_estimators = val
  rfTask.oob_score = !true
  rfTask.execute

  ; parse the model file
  parsed = json_parse(rfTask.output_model_uri, /fold_case)

  ; get and save the OOB score
  oobs.add, parsed['statistics', 'report', 'oob_score']
endforeach

; plot our data
p = plot([15 : 200 : 5], 1 - oobs.toArray(), title = 'OOB Score')
ax = p.axes
ax[0].title = 'Num. Estimators'
ax[1].title = '1 - OOB'
```

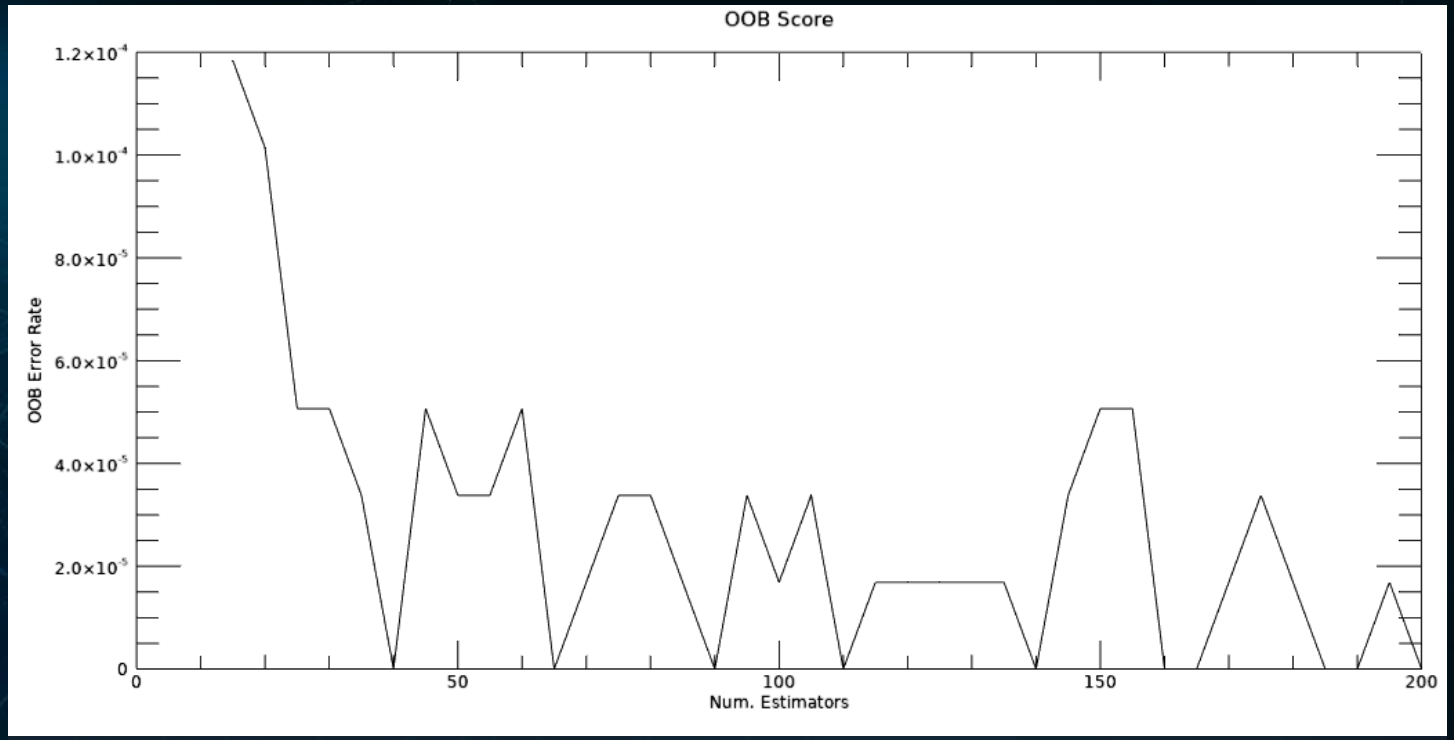

Out-of-Bag (OOB) error rate plot from random forest classifiers made with ENVI Machine Learning and the ENVI API in IDL

NOTABLES

ENVI Deep Learning: Updated Tutorials

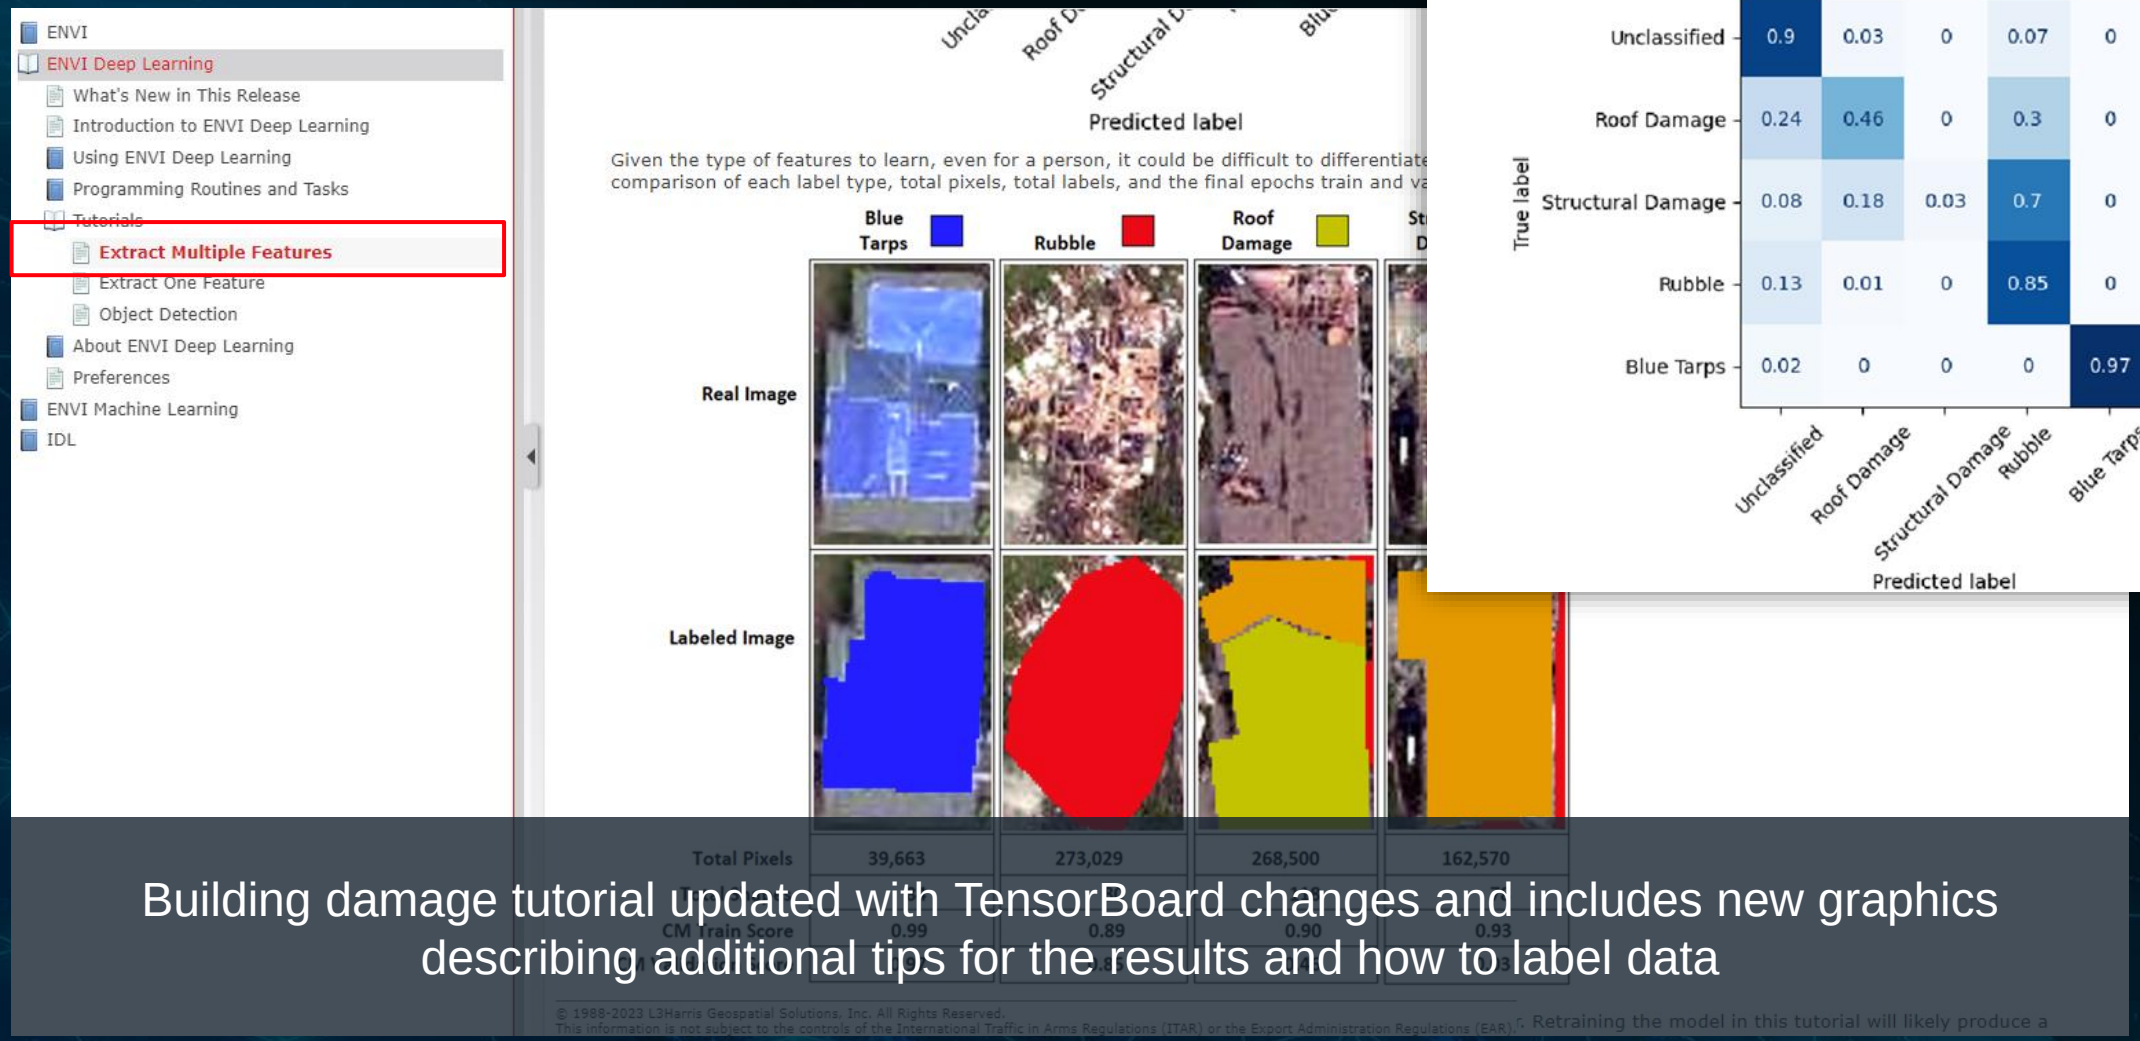

Building damage tutorial updated with TensorBoard changes and includes new graphics describing additional tips for the results and how to label data

ENVI Deep Learning: Updated Tutorials

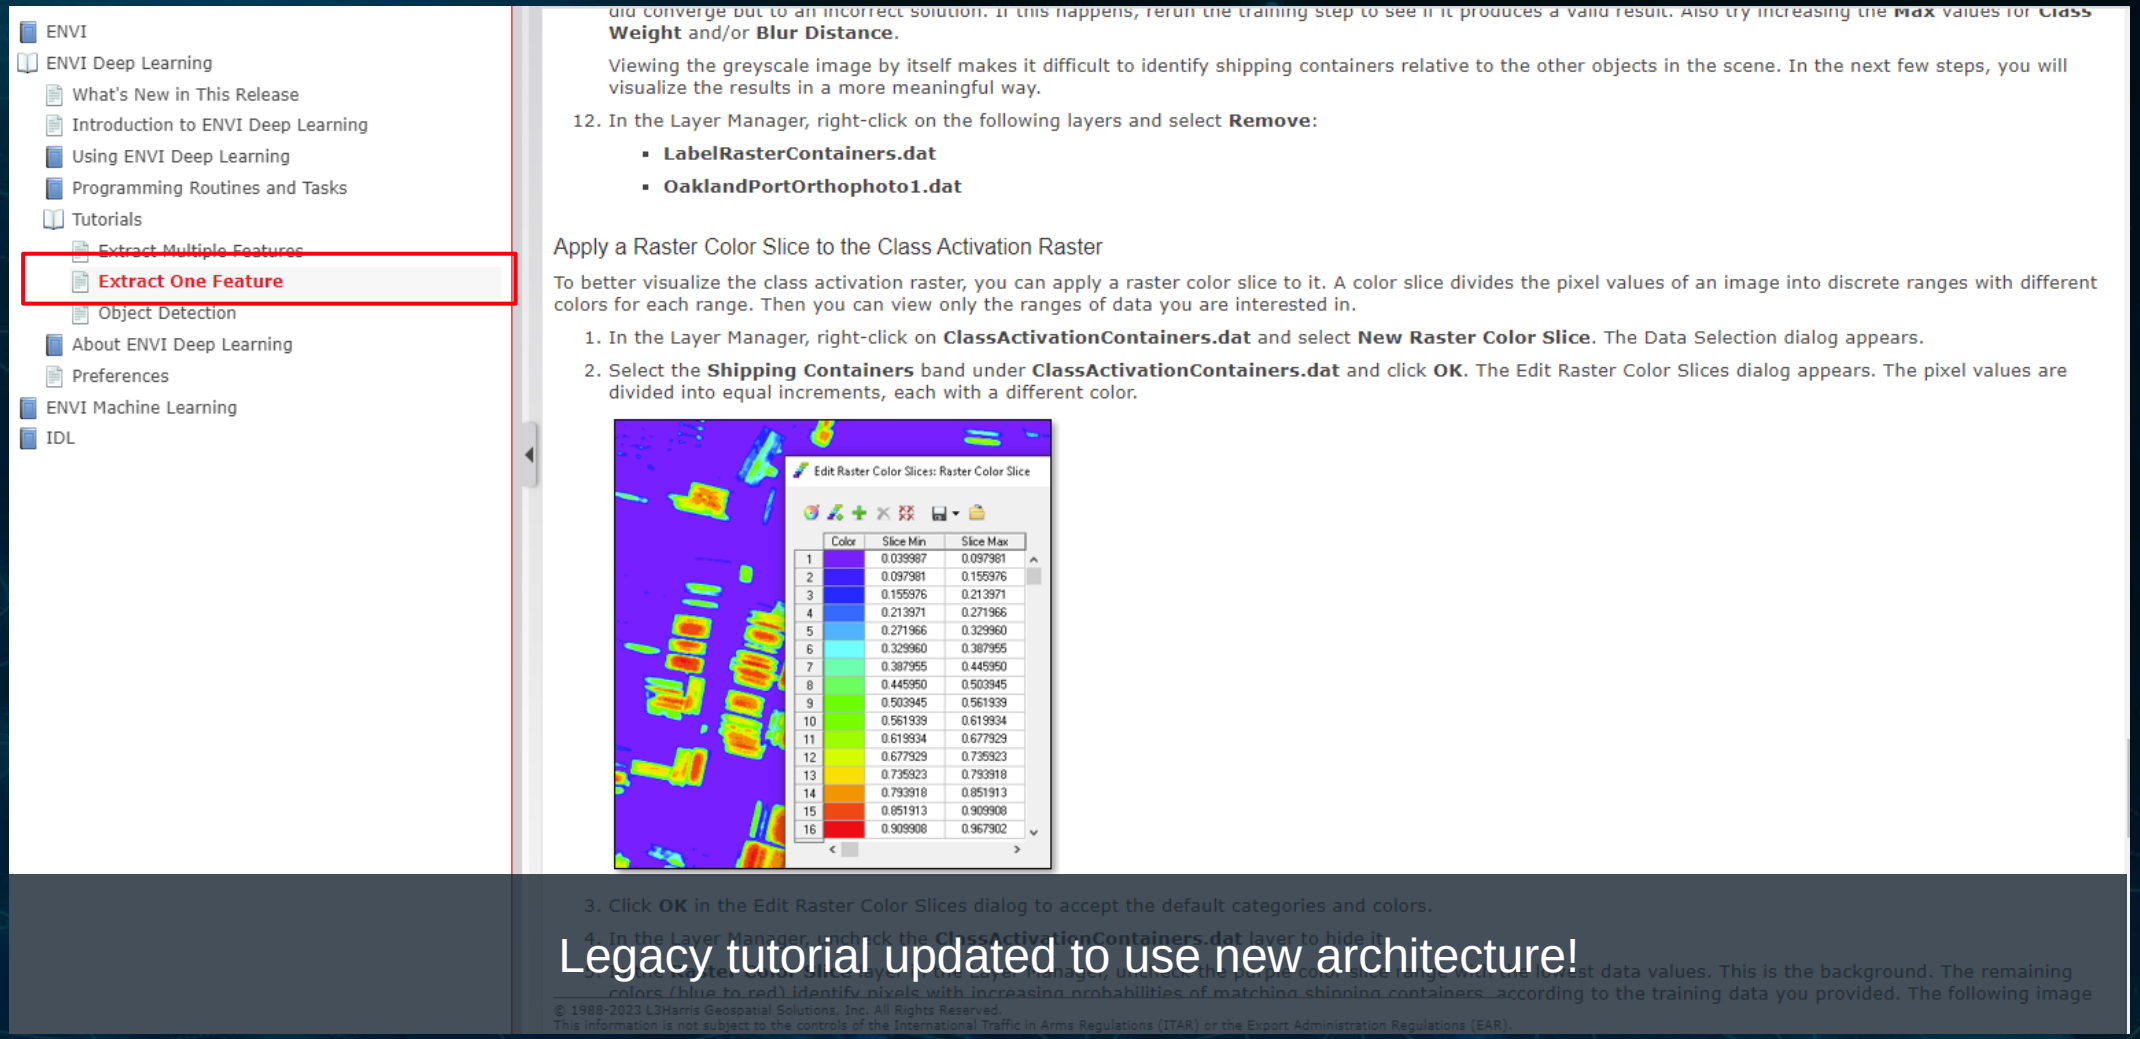

Legacy tutorial updated to use new architecture!

Machine Learning: Algorithm Background

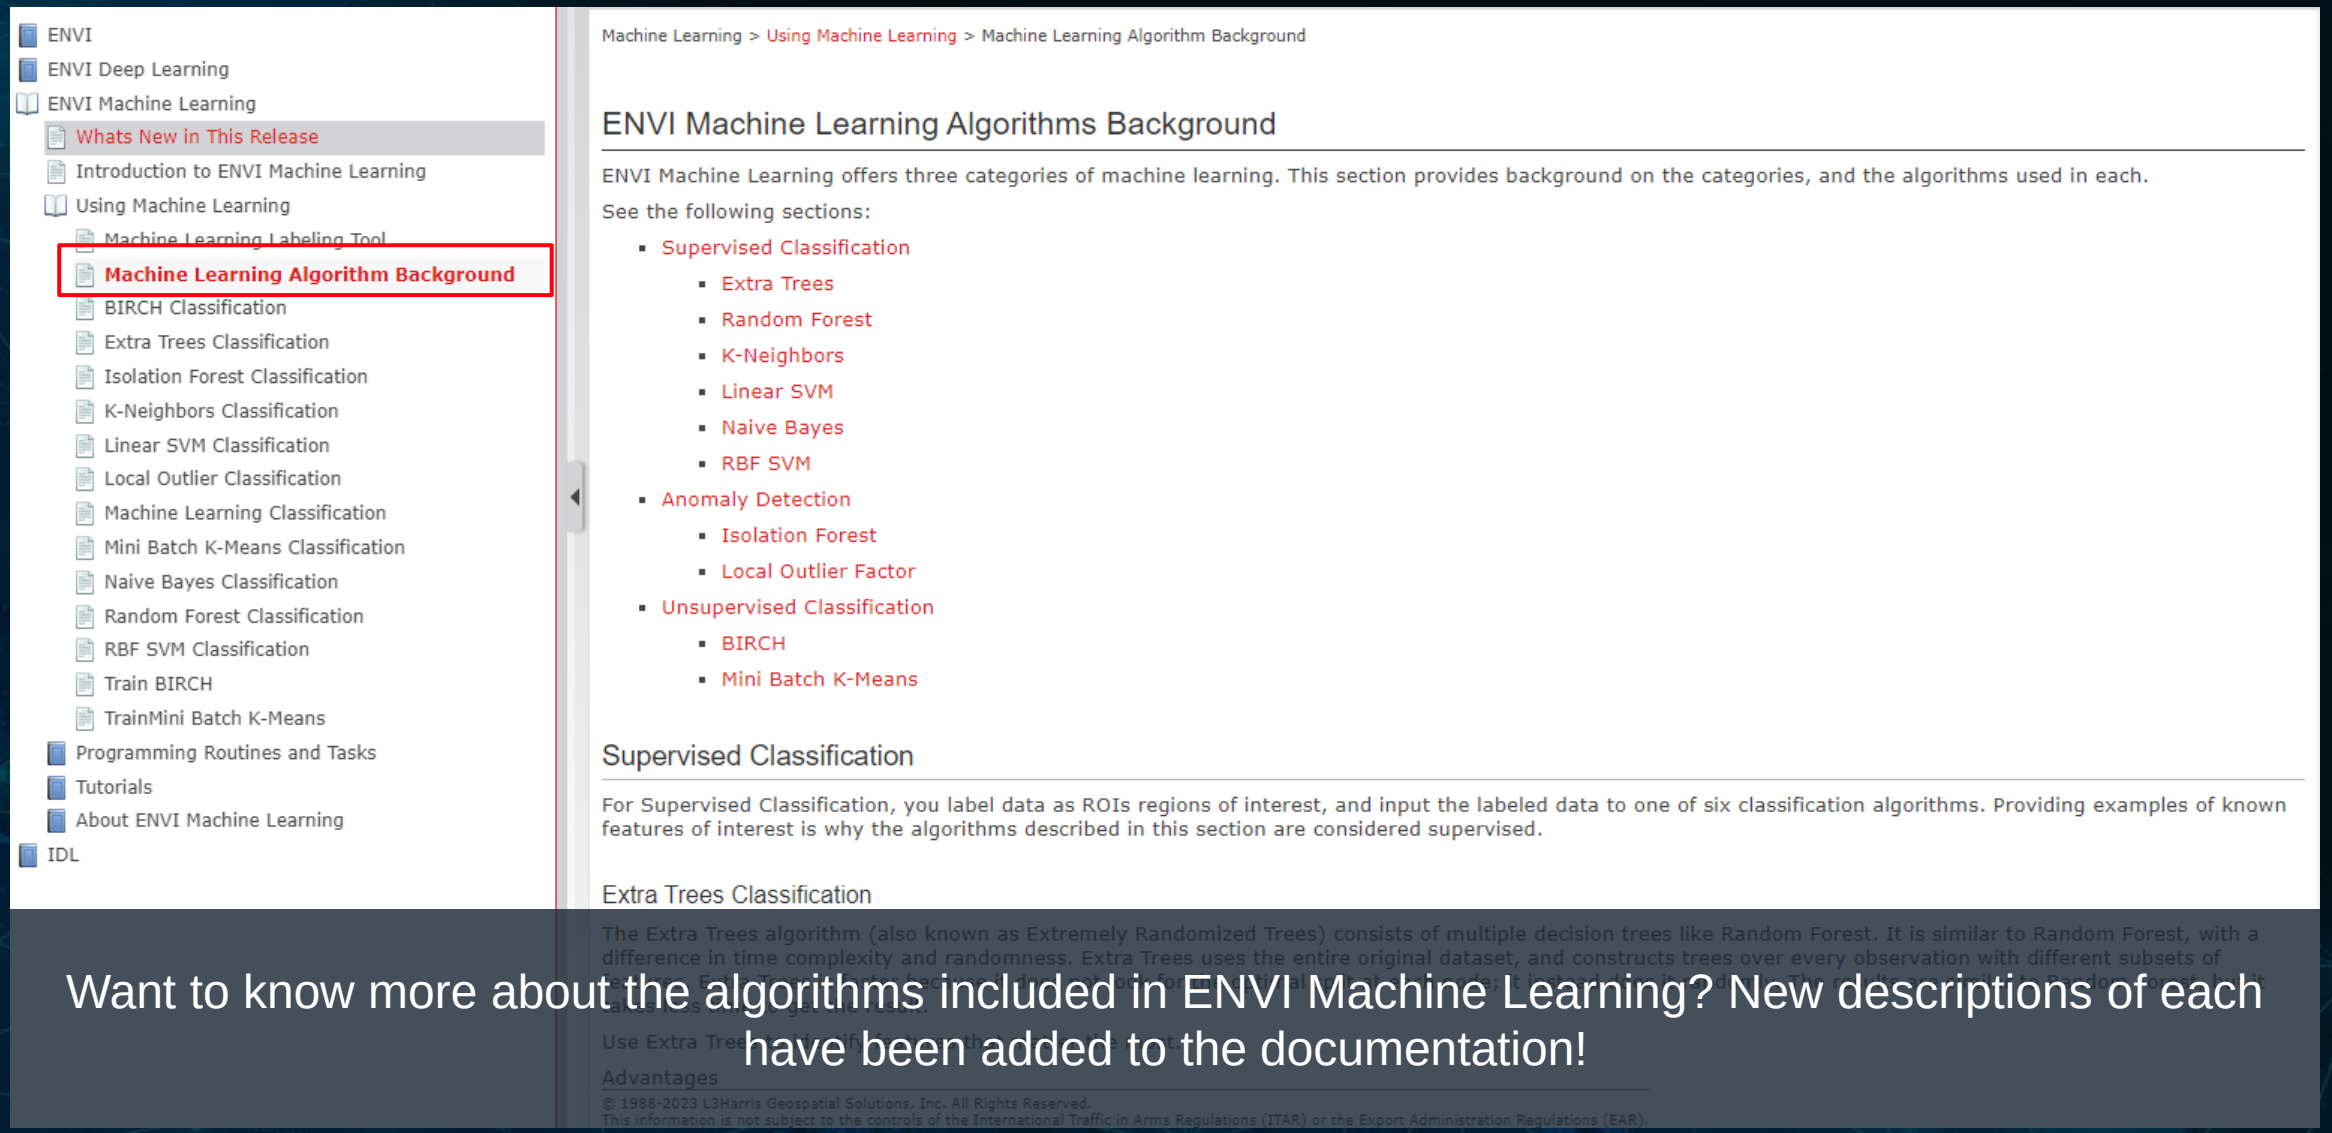

Want to know more about the algorithms included in ENVI Machine Learning? New descriptions of each have been added to the documentation!

User Experience: Band Names

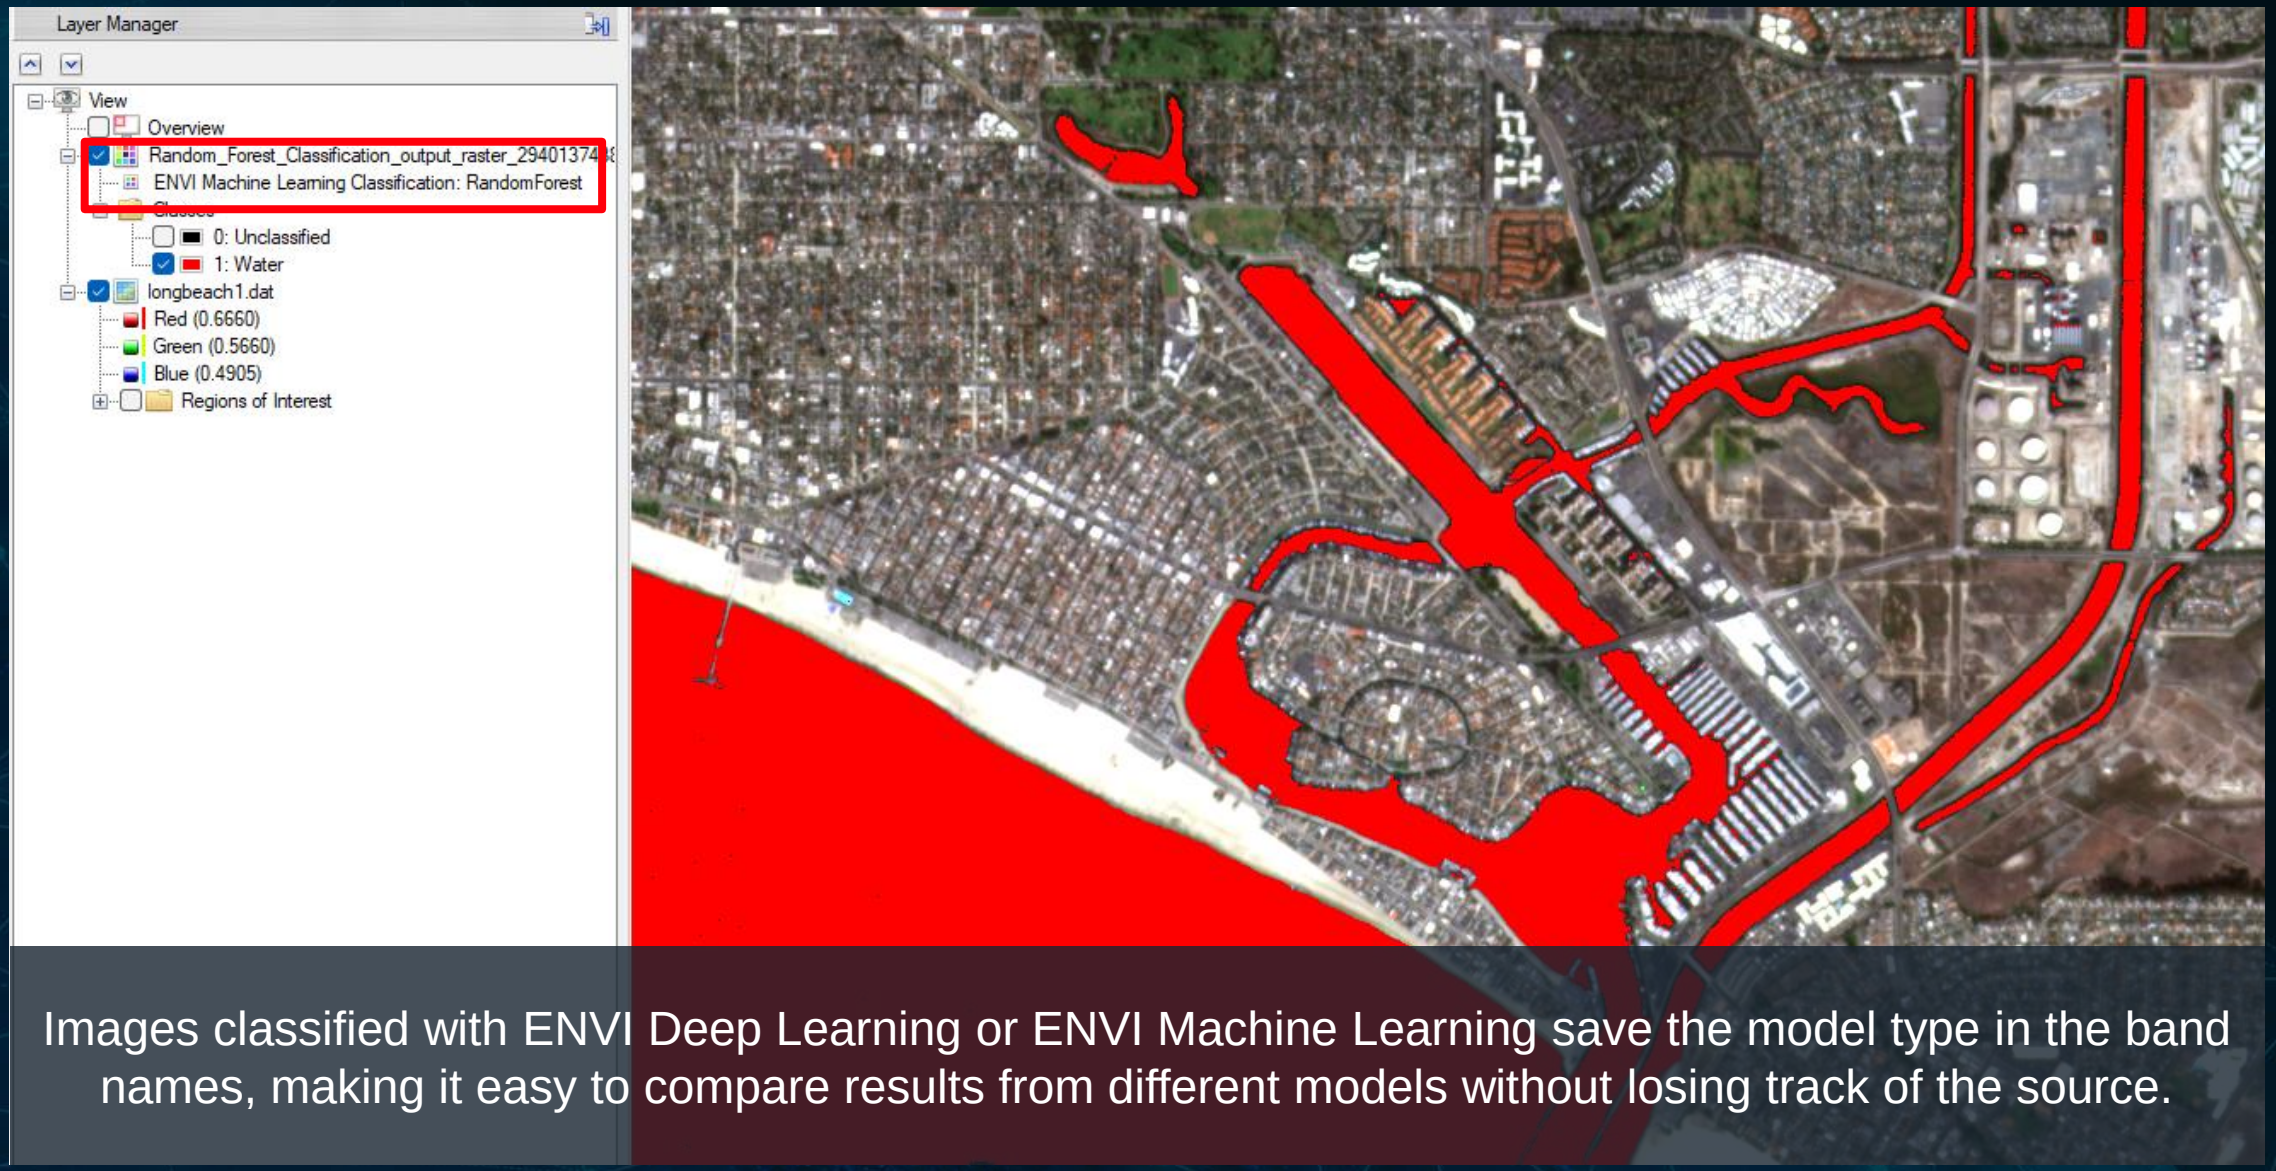

Images classified with ENVI Deep Learning or ENVI Machine Learning save the model type in the band names, making it easy to compare results from different models without losing track of the source.

Usability: New Tooltips

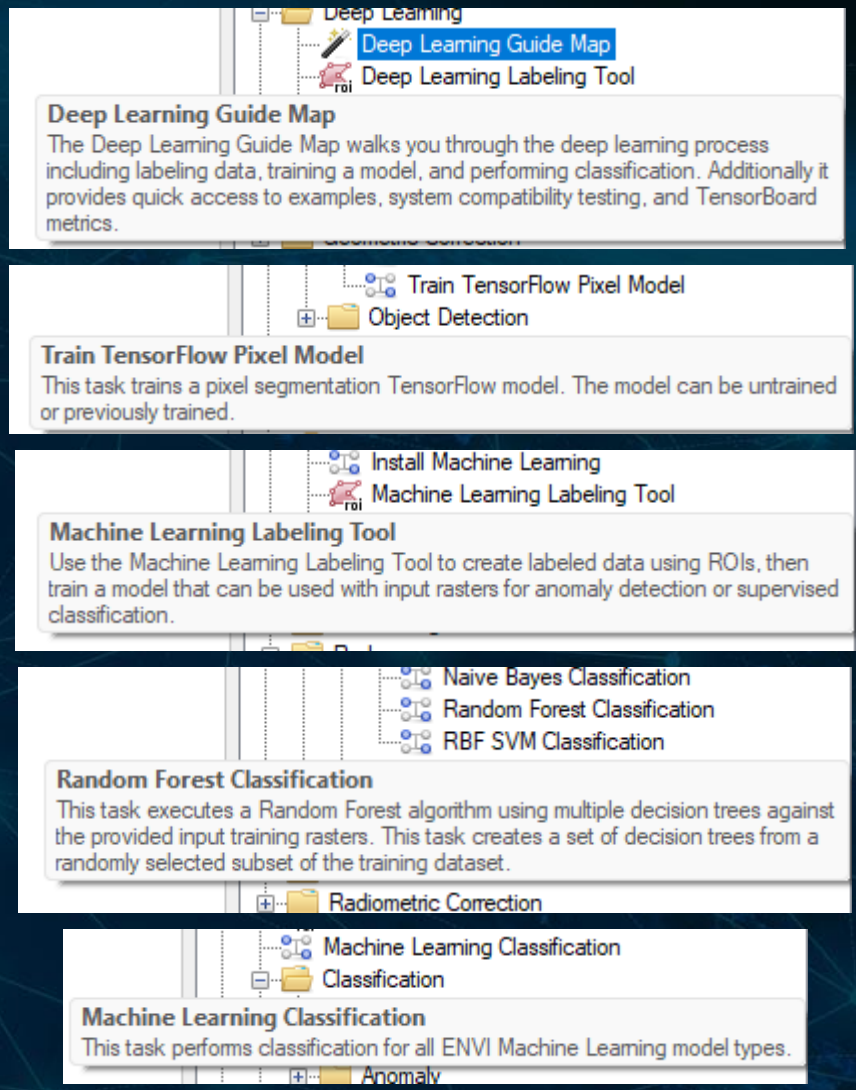

Like ENVI, the machine learning and deep learning toolbox entries now have tooltips with basic descriptions of each tool

Questions and Discussion

Bill Okubo

Product Manager

bill.okubo@nv5.com

Zach Norman

Product Manager

zachary.norman@nv5.com

NV5 Geospatial

www.NV5GeospatialSoftware.com

geospatialinfo@nv5.com

303-786-9900