

WHAT'S NEW IN ENVI DEEP LEARNING 1.2

January 2022

Zach Norman JP Metcalf



Contact Information and Introductions





JP Metcalf

Solutions Engineer jp.metcalf@l3harris.com

Zachary Norman

Product Manager zachary.norman@l3harris.com



Agenda



ENVI Deep Learning Overview What is Object Detection? ENVI Deep Learning: Object Detection Workflow Examples and Use Cases Object Detection vs. Pixel Classification Questions and Discussion

ENVI Deep Learning Recap



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



Keras

<u>Without</u> needing to program, the capabilities include:

- Segmentation (i.e. cloud masking)
- 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

ENVI 5.6.1 is needed for DL 1.2

NVIDIA GPU card with CUDA® Compute Capability 3.5 or higher https://developer.nvidia.com/cuda-gpus

CUDA 11 is required for the latest version of ENVI Deep Learning

Minimum 8GB of GPU RAM recommended



What is Object Detection?







Object Detection



Semantic Segmentation



Instance Segmentation



ENVI Deep Learning Pixels are classified, post processing can extract features if they don't touch

Not in ENVI Deep Learning Returns geometry (polygon) for features and detections are allowed to overlap

Image source: https://medium.com/onepanel/instance-segmentation-with-mask-r-cnn-and-tensorflow-on-onepanel-6a072a4273dd

ENVI Deep Learning 1.2

Returns geometry

(bounding box) for

features and detections

are allowed to overlap

ENVI Deep Learning 1.2 Toolbox



Existing capabilities have been moved under "Pixel Segmentation"



Labeling Tool



💽 Create New	Labeling Project		×
Project Type	Object Detection \sim		
Project Name	Awesome Object Detection Project		
Project Folder			
		ОК	Cancel

When you create a new project using the Deep Learning Labeling Tool, there is a new setting to indicate whether you are labeling data for "Pixel Classification" or "Object Detection" A short video showing the Labeling Tool being used to identify ships for an object detection model



Training: Pixel Classification vs Object Detection

Х

OK Cancel



Pixel Classification (UNet)

Train TensorFlow F	Pixel Model	- 0		×	Train TensorFlow	Pixel Model	_	
🔏 Import Parameters f	rom Model				// Import Parameters	from Model		
	-			^	Augment Scale	●Yes ○No		
Input Model	New Model		•		Augment Rotation	● Yes ○ No		
					Number of Epochs	25		
Training Rasters	[<u>^</u>	1.		(optional)			
Indining Nasiers		<u> </u>	2		Patches per Epoch (optional)	* (1	to 500,000)	
	<u>_</u>	/	Î		Patches per Batch			
			1.		(optional)	۷		
alidation Rasters		<i>•</i>	1 ส		Patch Sampling Rate	16 (0.0 d	or higher)	
	<	>	1		(optional)			_
					(optional)			
					nii 2 🖌	Enter one item per	line.	~
					Solid Distance			
					(optional)			
						Enter one number	per line.	
					Blur Distance			-
					(optional)			~
						Enter a 2 by N arra	y like [[1,2],[3,4]	.[5,6]]
					Class Weight (optional)	Min	ax	
					Loss Weight	(0.0 c	or higher)	
					(optional)			
					Output Model			
					Output Last Model			
					0		• (ОК

Object Detection

Model Name (optional)	ENVI Deep Learning OD
Model Description (optional)	<u></u>
Training Rasters	<pre></pre>
Validation Rasters	<pre></pre>
Pad Small Features	● Yes ○ No
Augment Scale	● Yes ○ No
Augment Rotation	● Yes ○ No
Number of Epochs (optional)	25 ×
Patches per Batch (optional)	1 0
Feature Patch Percentage (optional)	1 (0.01 to 1.0)
Background Patch Ratio (optional)	0.15 (0.0 or higher)
Output Model	
Output Last Model	
0	▼ OK Cancel

Object Detection Training Parameters





TensorBoard is a tool for visualizing the performance of your neural networks during and after training

Logs for object detection are reset each training session

Only loss is reported for object detection

Loss around 0.2 indicates a model is converging well

0



Object Detection: Results and Attributes





Classification results:

×

- Provided as a vector (shapefile) with class information and confidence (probability)
- When displayed in ENVI, features will have their class color applied
- Shapefiles are used for all outputs, even when images are not georeferenced

Examining Detections



Fun fact: when you are viewing the outputs of object detection in ENVI, you can colorize the vector data by the confidence of detects.

This is a great tool to quickly gauge the performance of models and see where lowprobability detects are found.







Before and after results using the new "Postprocess Classification Vector" task Once you have classified an image, there are two options to cleanup results:

- 1. Use the "Postprocess Classification Vector" to filter detections based on confidence, IoU, and optionally intersection
- 2. Within ENVI, you can use the vector editing tool to add or remove features



Easy-to-use ENVI API



Custom Data Preprocessing

29 ;process and generate labels

30 foreach file, process, fIdx do begin

31 print, fIdx 32 ;build the output

- 3 outRaster = file.replace(' metadata.xml', ' labels.dat')
- if file_test(outRaster) then begin
- 5 labelRasters.add, e.openRaster(outRaster)

6 continu

7 endi

;get the ROI file roiFile = labels[fIdx]

. :open evervthina

roi = e.openROI(roiFile)
rasters = e.openRaster(file)

;normalize

	norm =	ENVILinearRang	geStretchRaster	(rasters[1	ι],	MIN=0,	MAX=10000)
--	--------	----------------	-----------------	------------	-----	--------	------------

;build the label raster

task = ENVITask('BuildLabelRasterFromROI')
task.INPUT_RASTER = norm
task.INPUT_ROI = roi
task.OUTPUT_RASTER_URI = outRaster
task.execute
content = cont

7 norm.close

- i8 foreach r, rasters do r.close
- 59 foreach r, roi do r.close

50 labelRasters.add, e.openRaster(outRaster)

i1 endforeach

Classifier Generation

; create our task
task = ENVITask('TrainTensorflowObjectModel')

; specify training time task.EPOCHS = epochs task.FEATURE_PATCH_PERCENTAGE = dataPerEpoch task.BACKGROUND_PATCH_RATIO = bpr

; set some properties for our project task.MODEL_NAME = 'Crosswalks' task.MODEL_DESCRIPTION = 'Crosswalks from aerial data'

; set our rasters
task.TRAINING_RASTERS = training
task.VALIDATION_RASTERS = validation
task.PATCHES_PER_BATCH = patchesPerBatch

; set our outputs
task.OUTPUT_MODEL_URI = bestUri
task.OUTPUT_LAST_MODEL_URI = lastUri

; execute task.execute

Creating Custom Training Data



; create our ROI, let ENVI handle colors roi = ENVIROI(NAME = "Airplane")

; buffer around points, pixels buffer = 20

; add all of our features to the ROI foreach point, points do begin ; make a polygon bounding box from our points roi.AddGeometry, [\$ [point[0]-buffer, point[1]-buffer],\$ [point[0]-buffer, point[1]+buffer],\$ [point[0]+buffer, point[1]+buffer],\$ [point[0]+buffer, point[1]-buffer],\$ [point[0]-buffer, point[1]-buffer]],\$ /POLYGON endforeach

; get our task to make an object detection raster Task = ENVITask('BuildObjectDetectionRasterFromROI')

; specify our inputs
Task.INPUT_RASTER = e.openRaster("C:\some\file.dat")
Task.INPUT_ROI = [roi]

; run Task.<mark>Execute</mark>

Recommended

When the out-of-the-box tools for data preparation are not enough, you have several options via the ENVI API to choose from to get your data in the right format.

Out-of-the-box tools allow you to create training data from ROIs, annotations, and vectors (after conversion to ROIs)

Advanced usage

specify class information

names = list("Airplane", "Helicopter", "Fighter Jet")
colors = list([255, 0, 0], [0, 255, 0], [0, 0, 255])

; specify some training data boxes = ...; [nBoxes, 4] array of [xmin, ymin, xmax, ymax] pixel coordinates classes = ...; [nBoxes] array of classes, zero-based

; make bounding box set objBox = ENVIBoundingBoxSet()

; add our classes to the bounding box set

foreach name, names, idx do \$
 objBox.AddClass, CLASS=idx, LABEL=name, COLOR=colors[idx]

process each bounding box

Foreach class, classes, idx do begin
 ; unpack the bbox
 bbox = boxes[*,idx] ; [xmin, ymin, xmax, ymax]

; convert extent to [2,5] bbox: ; [[Upper Left XY], ; [Upper Right XY], ; [Lower Right XY], ; [Lower Left XY], ; [Upper Left XY]] bbox = reform([bbox[[0,1]], bbox[[0,3]], \$ bbox[[2,3]], bbox[[2,1]], bbox[[0,1]]], [2,5])

; add to our training data

objBox.AddBoundingBox, CLASS=0, BOUNDING_BOX=bbox endforeach

; make GeoJSON
geoJSON = objBox.GetGeoJSON()

update and add metadata fields

meta = myRaster.METADATA
meta['dlobjectbounds'] = ENVIDeepLearningObjectDetectionRaster.Encode(geoJSON)
meta['dlclassnames'] = names.toArray() ; 1D array
meta['dlclasscolors'] = colors.toArray() ; 1D*3 array
myRaster.WriteMetadata

Verifying Custom Training Data



Use the "<u>View Object Detection</u> <u>Raster Labels</u>" tool to verify your bounding boxes are correct before training

Important note!

There is an ENVI bug you may encounter while using this tool. It will be fixed in the next release of ENVI (likely Q1 2022)



Object Detection Examples



1

Data Source: WorldView panchromatic

Ship Detection: Optical

Data Source: PlanetScope VNIR, surface reflectance

N

Ship Detection: SAR





Data Source: Sentinel 1 (SAR)

23

Turn Arrows





Vehicle Detection





Data Source: RGB aerial imagery



Feature	Object Detection	Pixel Classification
Label types	Bounding box	 Points, polylines, or polygons
What is used for training?	 Source data, bounding boxes, and box classes Each bounding box is a training sample, so there are few to learn from 	 Source data and classification image (i.e. mask) Each pixel is a training sample, so there are many to learn from
Epoch for training	 Recommended 60-75 minimum Target loss should be <=0.2; loss=0.5 will likely produce bad results 	 Depends on feature and training data, can get good models at 20 epochs Loss can vary, lower the better
Size of features	 Minimum: 25 pixels across (architecture limitation) Maximum: 400 pixels across (patch size limitation) 	 Minimum: At least one pixel (no sub-pixel detections) Maximum: Can learn features larger than your patch size
Relative size of classes	 Don't train a model with classes that vary dramatically in size For example: Not the best idea to detect parking lots and cars 	 Size is not as important Multi-class size differences can impact performance because of class balance
Model output	 Bounding box, class, and confidence/score 	 A class or probability for each pixel



Release details: <u>https://www.l3harrisgeospatial.com/Support/Maintenance</u>

JP Metcalf

Solutions Engineer jp.metcalf@l3harris.com

Zachary Norman

Product Manager zachary.norman@l3harris.com

L3Harris Geospatial

www.L3HarrisGeospatial.com geospatialinfo@I3harris.com 303-786-9900