Geospatial Deep Learning with ArcGIS

Rohit Singh
Vinay Viswambharan

2020 ESRI DEVELOPER SUMMIT | Palm Springs, CA

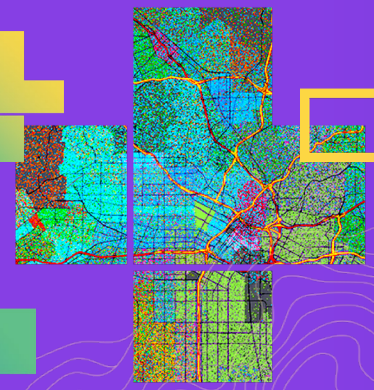

Session Overview

- AI, Machine Learning & Deep Learning
- Deep Learning Workflow
- Training Models
  - ArcGIS Pro
  - arcgis.learn
- Types of models and their applications
- Scalable deep learning with Image Server
- Learning Resources

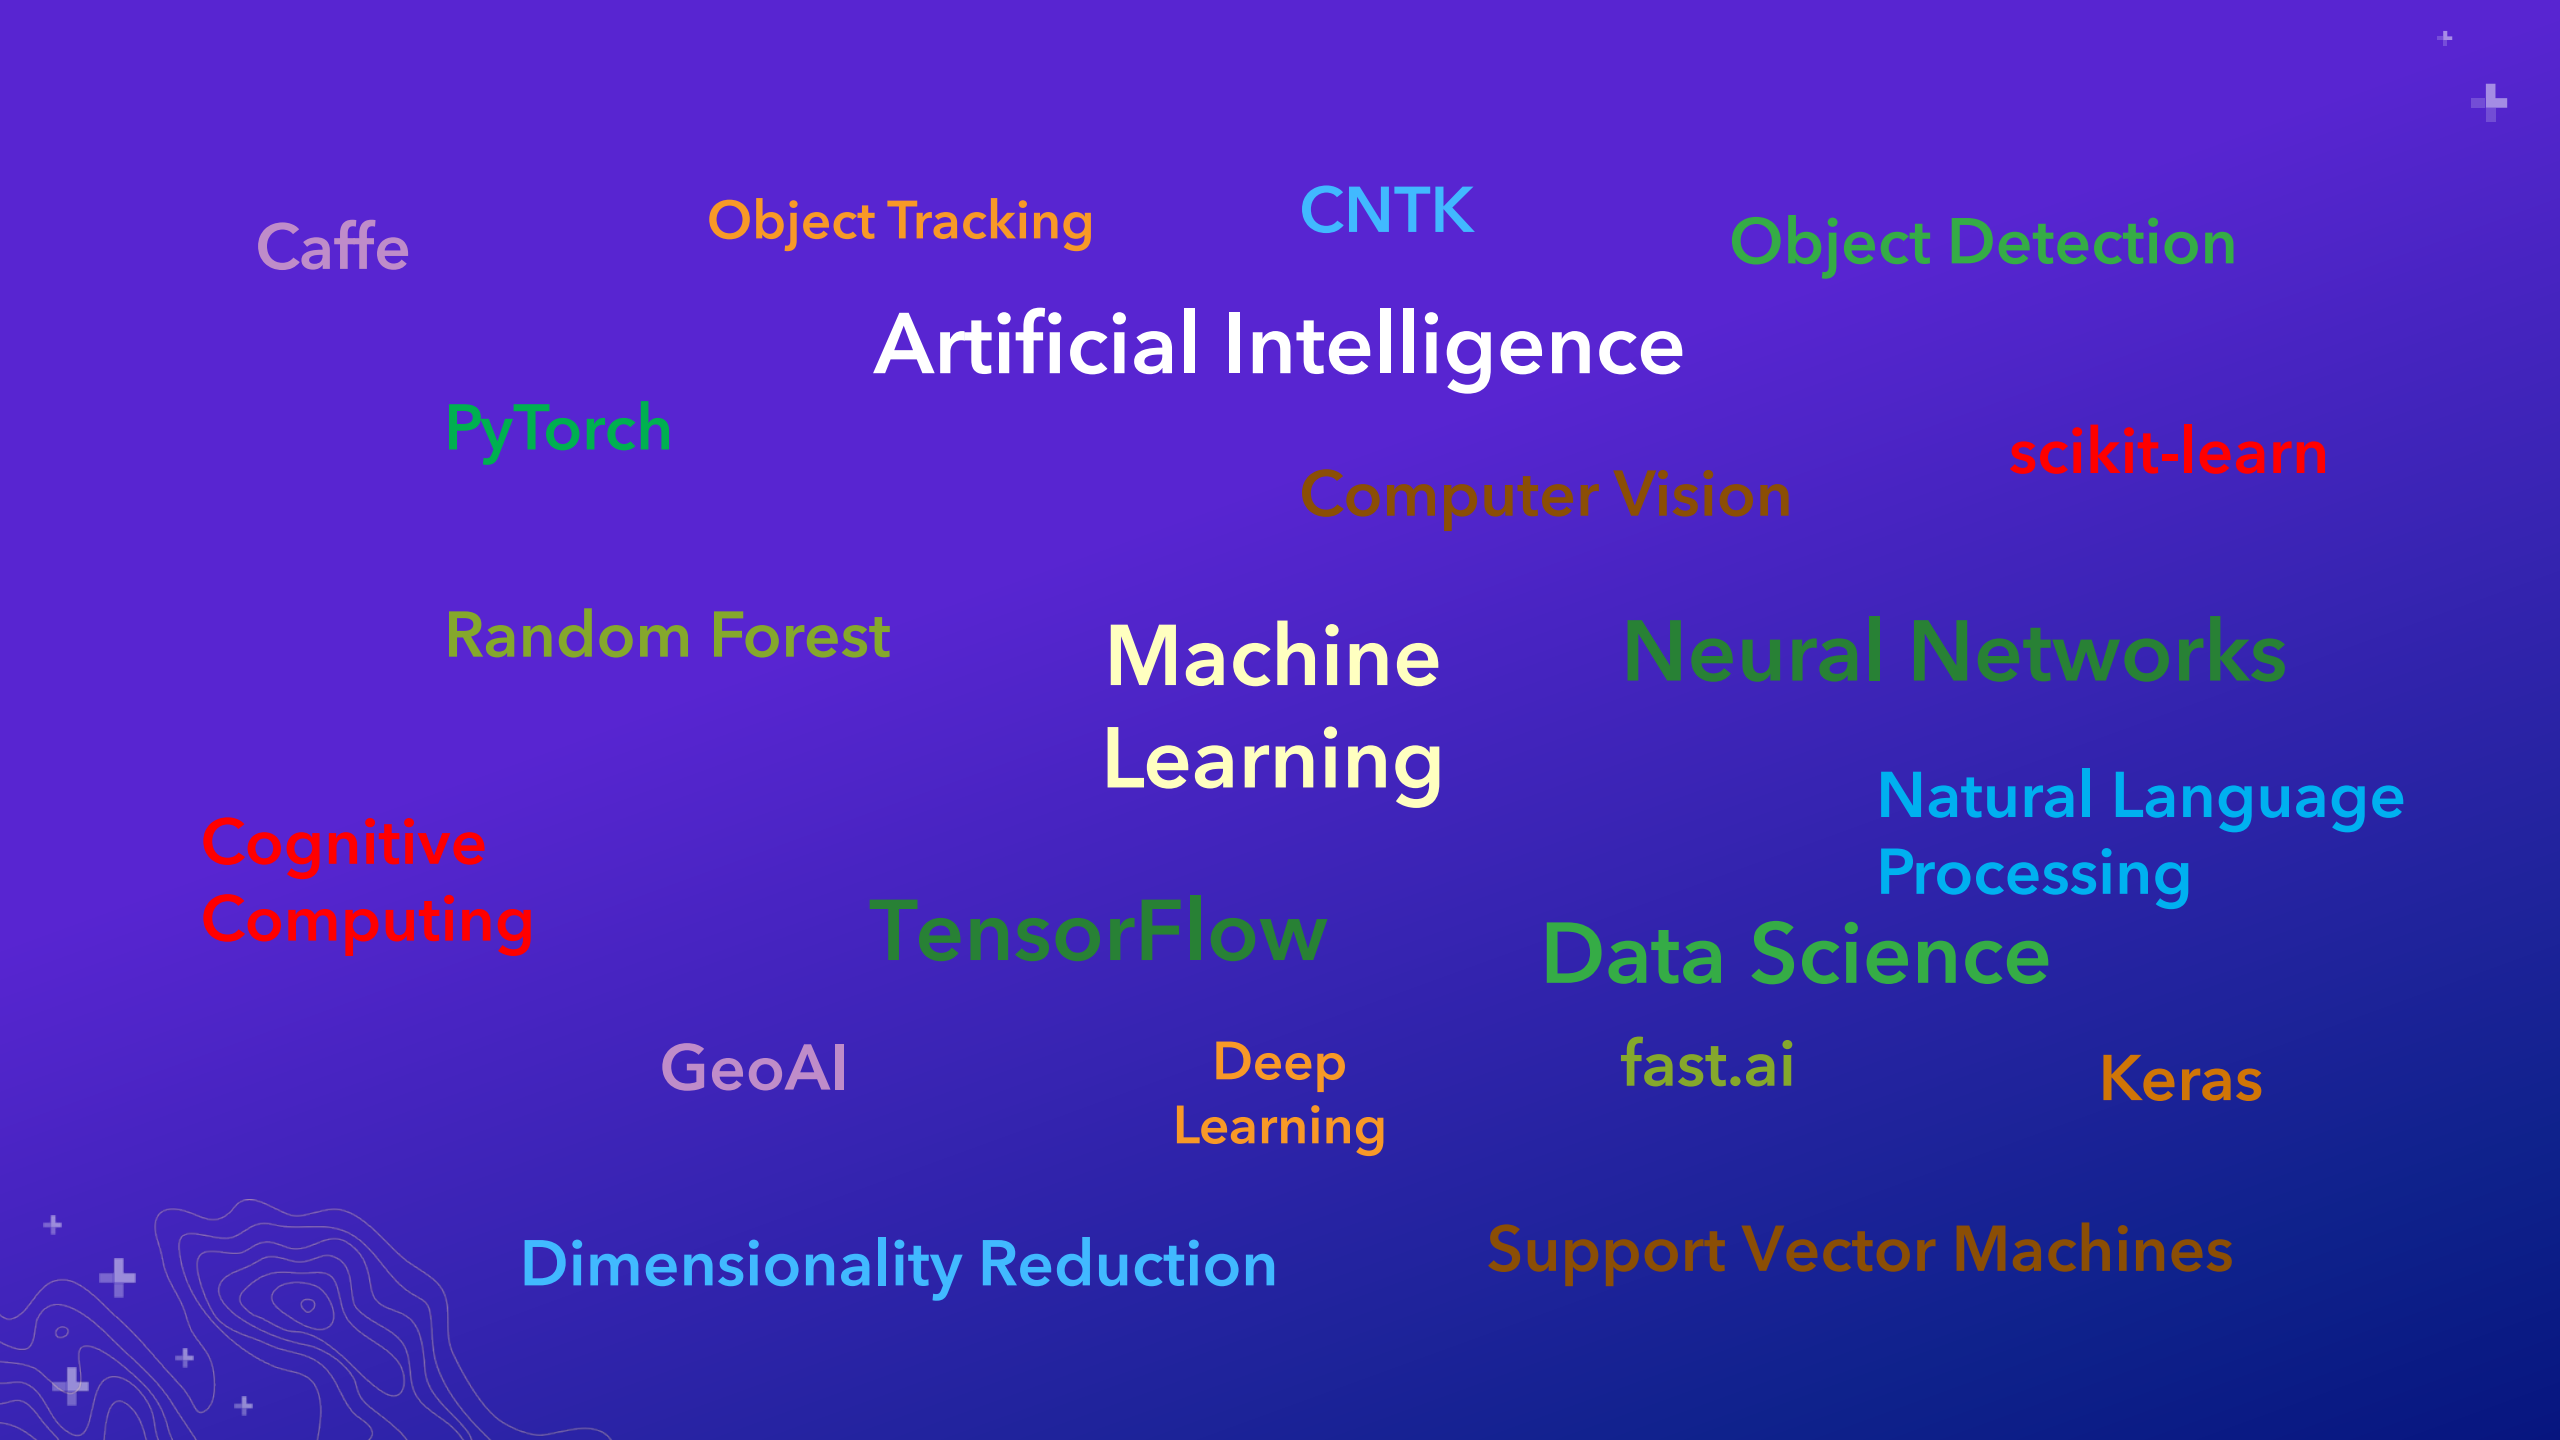
Caffe
Object Tracking
CNTK
Object Detection
Artificial Intelligence
PyTorch
scikit-learn
Computer Vision
Random Forest
Machine Learning
Neural Networks
Natural Language Processing
Cognitive Computing
TensorFlow
Data Science
GeoAI
Deep Learning
fast.ai
Keras
Dimensionality Reduction
Support Vector Machines

Artificial Intelligence

Machine Learning

Deep Learning

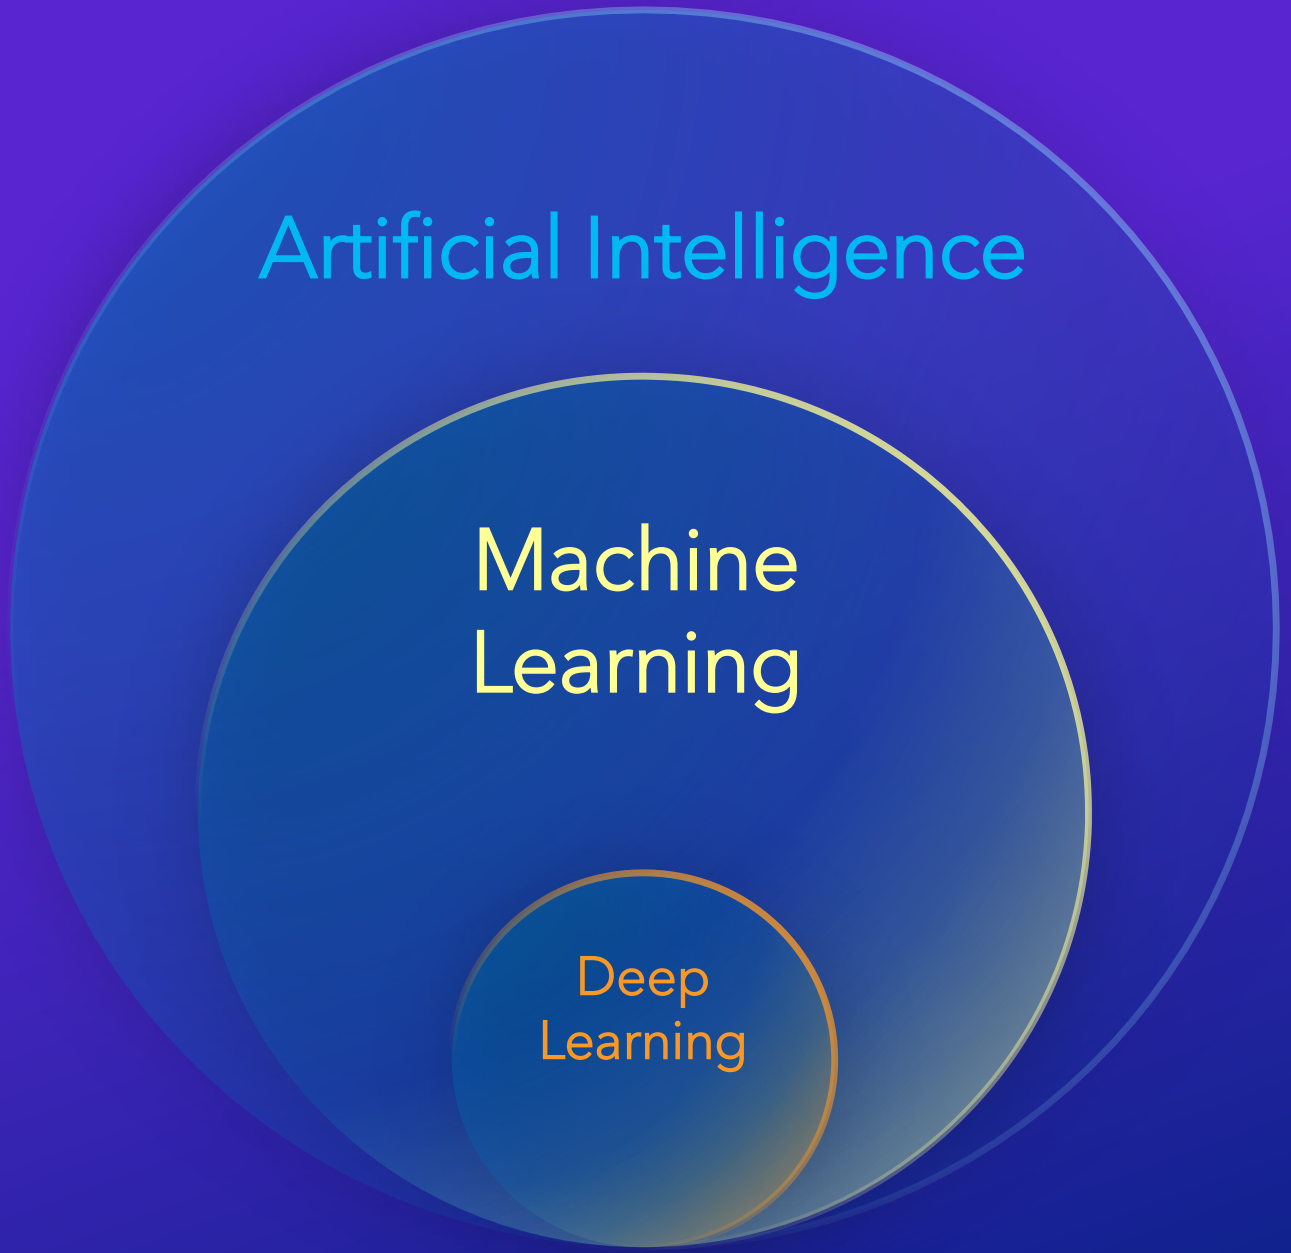
Artificial Intelligence
Machine Learning
Deep Learning

Contrasting Machine Learning with Deep Learning

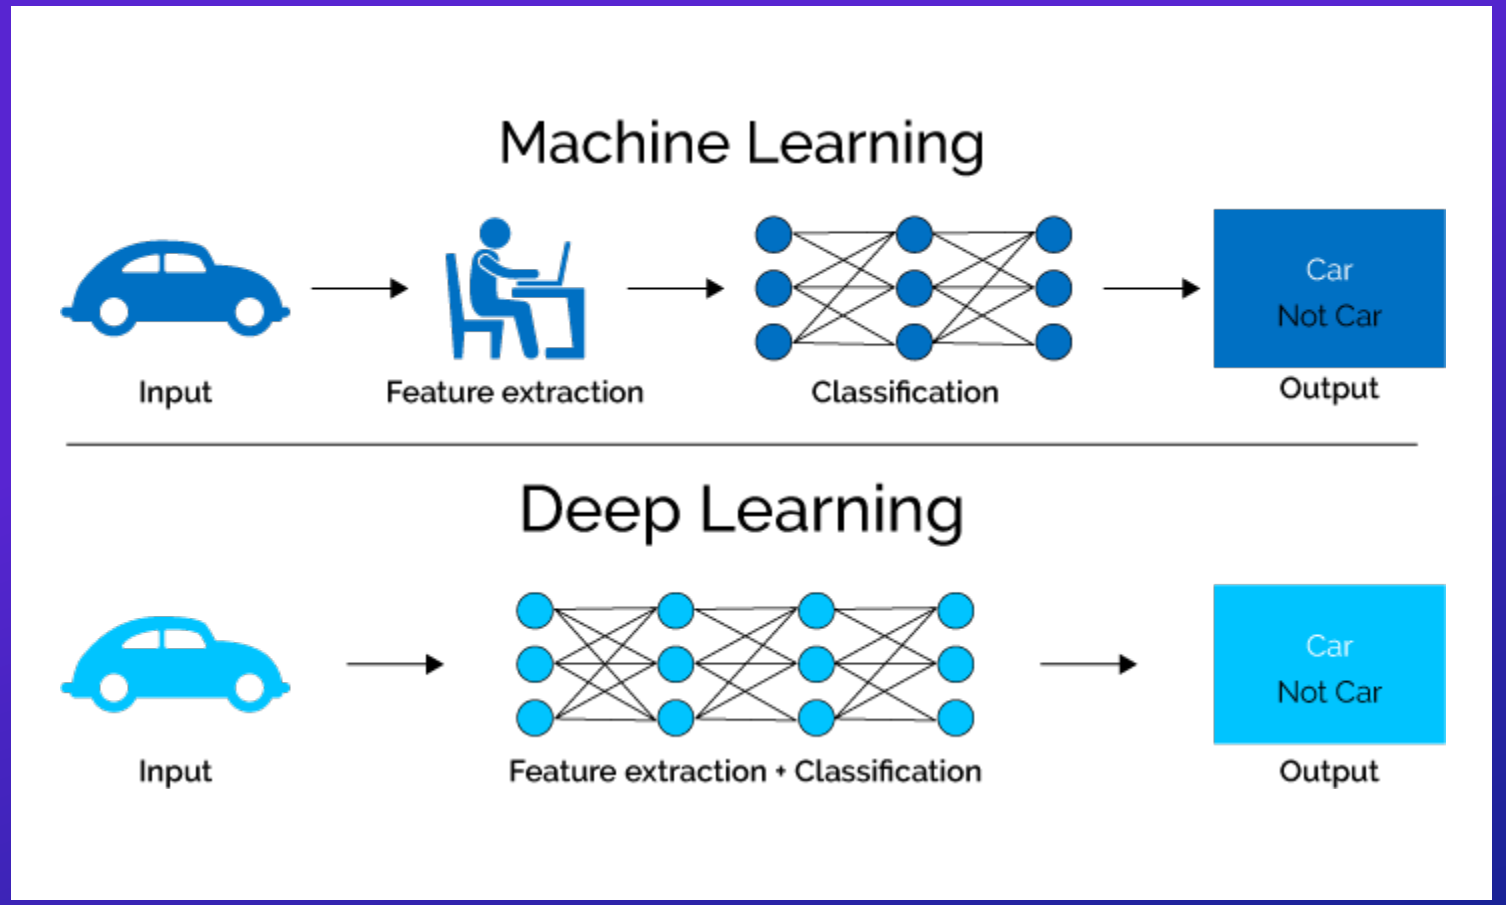

Machine Learning in ArcGIS

Classification

- Pixel & Object Based
- Image Segmentation
- Maximum Likelihood
- Random Trees
- Support Vector Machine

Clustering

- Spatially Constrained Multivariate Clustering
- Multivariate Clustering
- Density-based Clustering
- Hot Spot Analysis
- Cluster and Outlier Analysis
- Space Time Pattern Mining

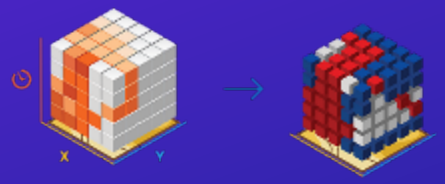

Prediction

- Empirical Bayesian Kriging
- Areal Interpolation
- EBK Regression Prediction
- Ordinary Least Squares Regression and Exploratory Regression
- Geographically Weighted Regression

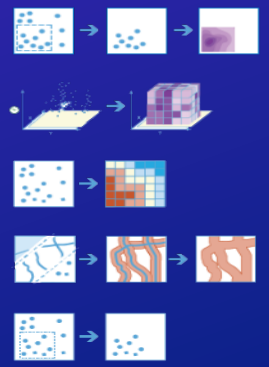

Deep Learning in ArcGIS

Data Preparation

- Label Objects
- Training Samples Manager
- Export Training Samples

Training

- Train Deep Learning Model
  - Object Detection
  - Object Classification
  - Pixel Classification
  - Instance Segmentation
- arcgis.learn module (ArcGIS API for Python)

Inferencing

- Detect Objects
- Classify Pixels
- Classify Objects
- Non Maximum Suppression

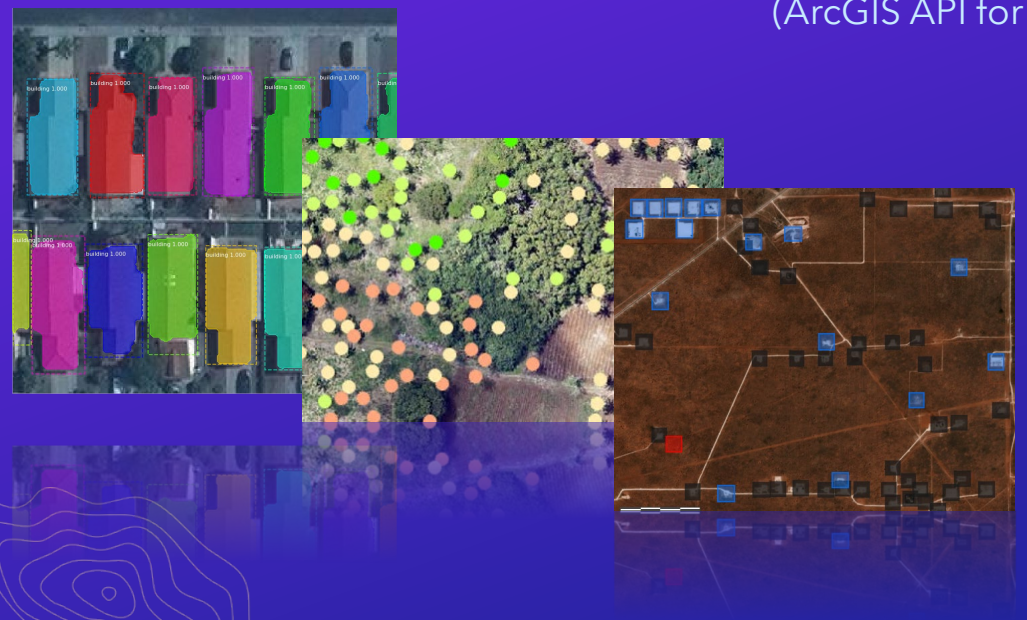

Why use Deep Learning for Imagery

Computer vision is now almost as good, if not better, than human vision

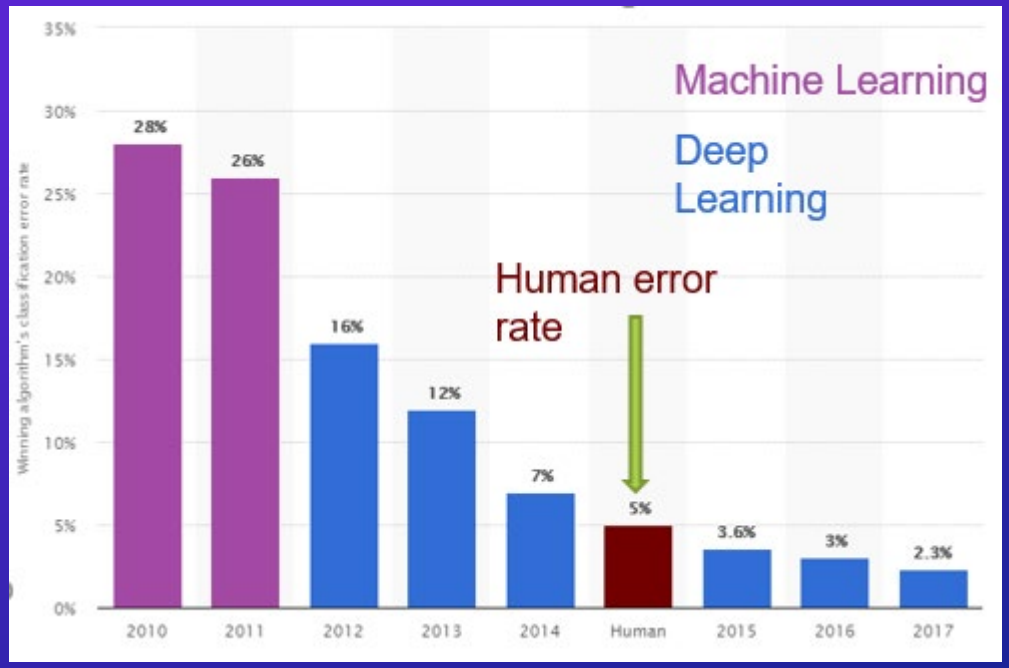

ImageNet Visual Recognition Challenge error rate

Applications of Deep Learning to GIS

Impervious Surface Classification

Coconut Tree Detection

Building Footprint Extraction

Damaged House Classification

End to End Deep Learning – Wide spectrum of deep learning models

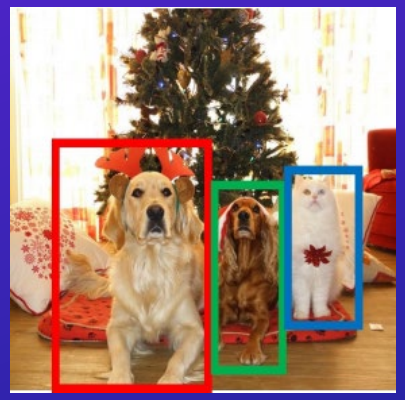

Pixel Classification

Object Detection

Instance Segmentation

Image Classification

Deep Learning with Imagery in ArcGIS

ArcGIS supports end-to-end deep learning workflows

- Tools for:
  - Labeling training samples
  - Preparing data to train models
  - Training Models
  - Running Inferencing
- Supports the key imagery deep learning categories
- Supported environments
  - ArcGIS Pro
  - Map Viewer
  - ArcGIS Notebooks/Jupyter Notebook

Part of ArcGIS Image Analyst

Run distributed on ArcGIS Image Server

Deep Learning Workflow in ArcGIS

Label Training Samples

Prepare data for training

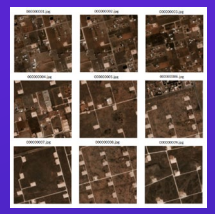

Train Deep Learning Model

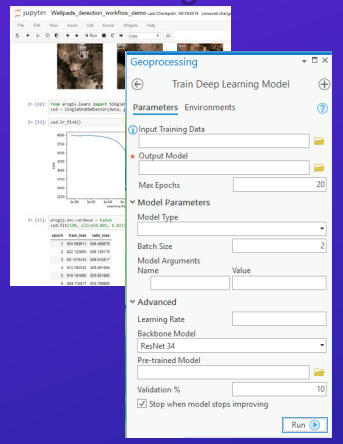

Inference

Detect/Classify Objects

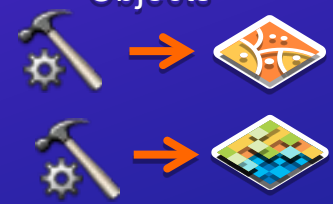

Classify Pixels

Collect Samples → Export Training Samples → Train → Perform Inference

Collect Training Samples / Label data

- Different methods
  - **Label Objects for Deep Learning – ArcGIS Pro (2.5)**
  - Training sample manager – ArcGIS Pro
  - Feature editing
    - ArcGIS Pro
    - Map Viewer
    - JS Web Apps
- Different data models
  - Feature class (local single user)
  - Feature services (collaborative experience)
  - Classified thematic rasters

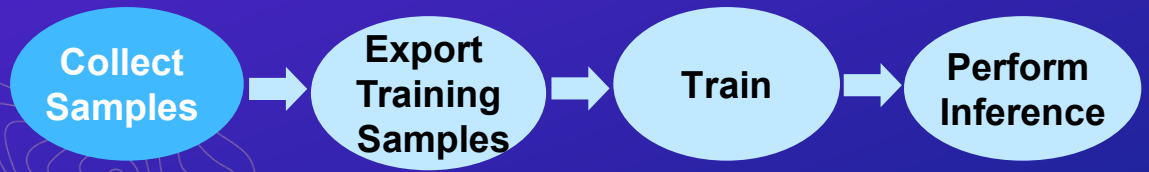

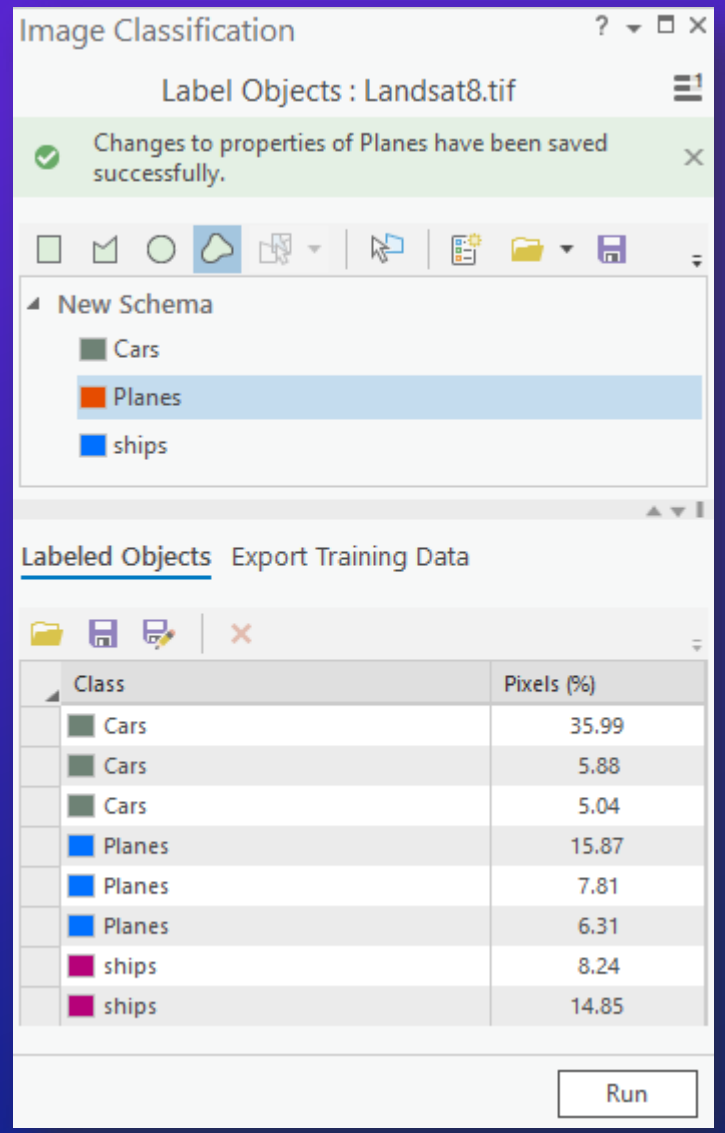

Export Training Data for Deep Learning

- Exports samples to training images
- Images have associated labels/metadata
- Writes out and ECD
- Used as inputs for model training
- Supports various formats

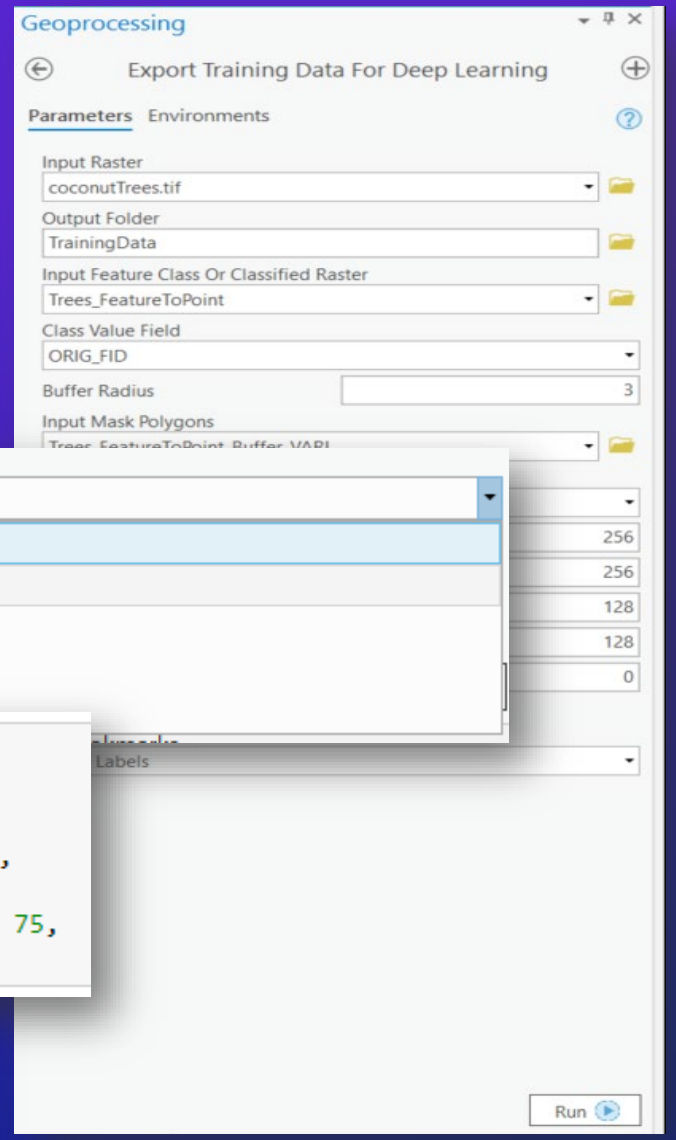


```
#Export training data
chips = export_training_data(sentinel_data,
                             well_pads,
                             "PNG", {"x":448,"y":448},
                             {"x":224,"y":224},
                             "PASCAL_VOC_rectangles", 75,
                             "well_pads")
```

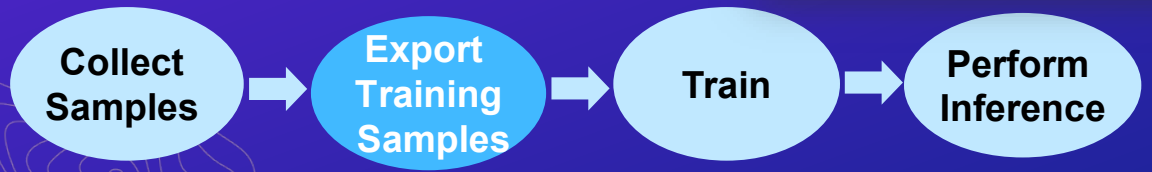

Train Deep Learning Model

- ArcGIS Pro and ArcGIS API for Python supports training
- ArcGIS Pro “Train Deep Learning Model” tool
- arcgis.learn module in ArcGIS API for Python
- Supported Models:
  - Object Detection - SSD, RetinaNet, MaskRCNN
  - Object Classification – Feature classifier
  - Pixel Classification – UNET, PSPNet
- External Deep Learning Frameworks
  - TensorFlow
  - CNTK…

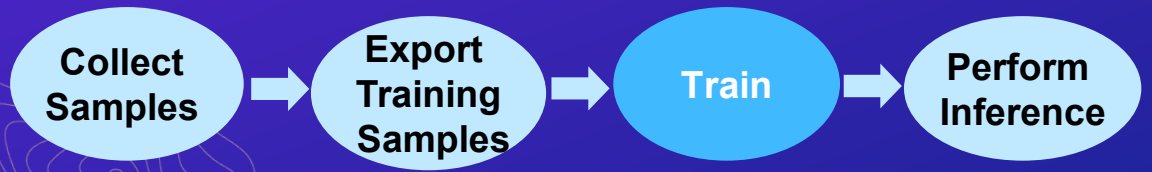

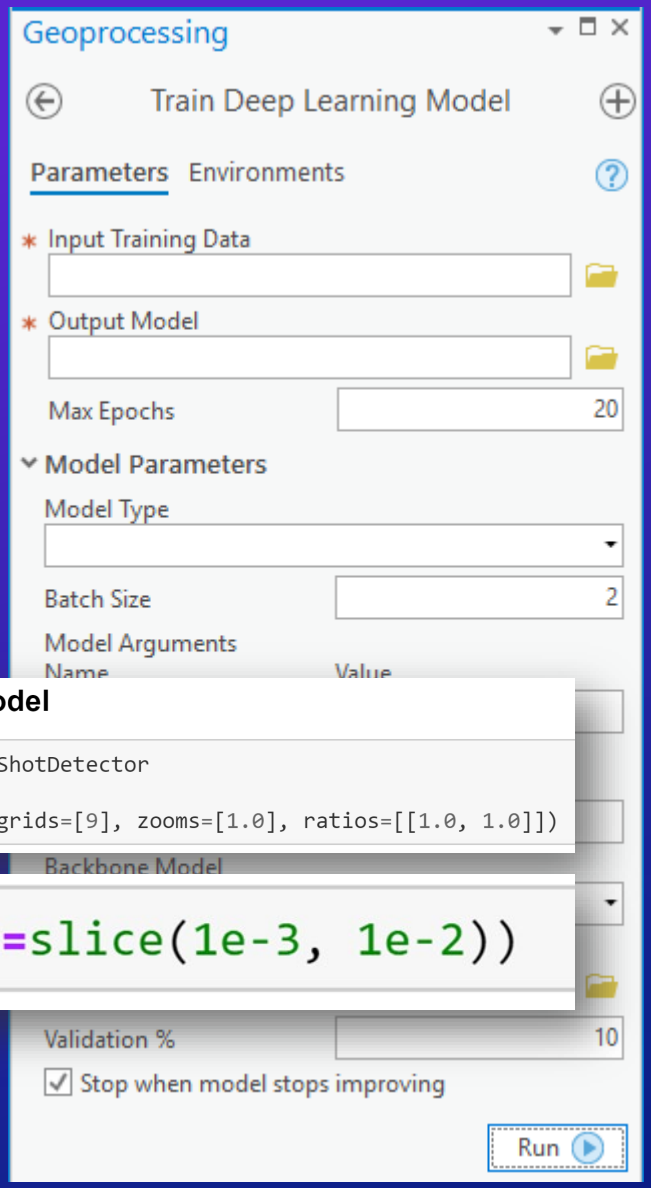

Train SingleShotDetector Model

```
from arcgis.learn import SingleShotDetector

ssd = SingleShotDetector(data, grids=[9], zooms=[1.0], ratios=[[1.0, 1.0]])
```

```
ssd.fit(10, lr=slice(1e-3, 1e-2))
```

Use Deep Learning Models

- Run on desktop and enterprise
- Parallel processing using enterprise
- Types of inferencing
  - Object detection
  - Classify objects
  - Pixel classification

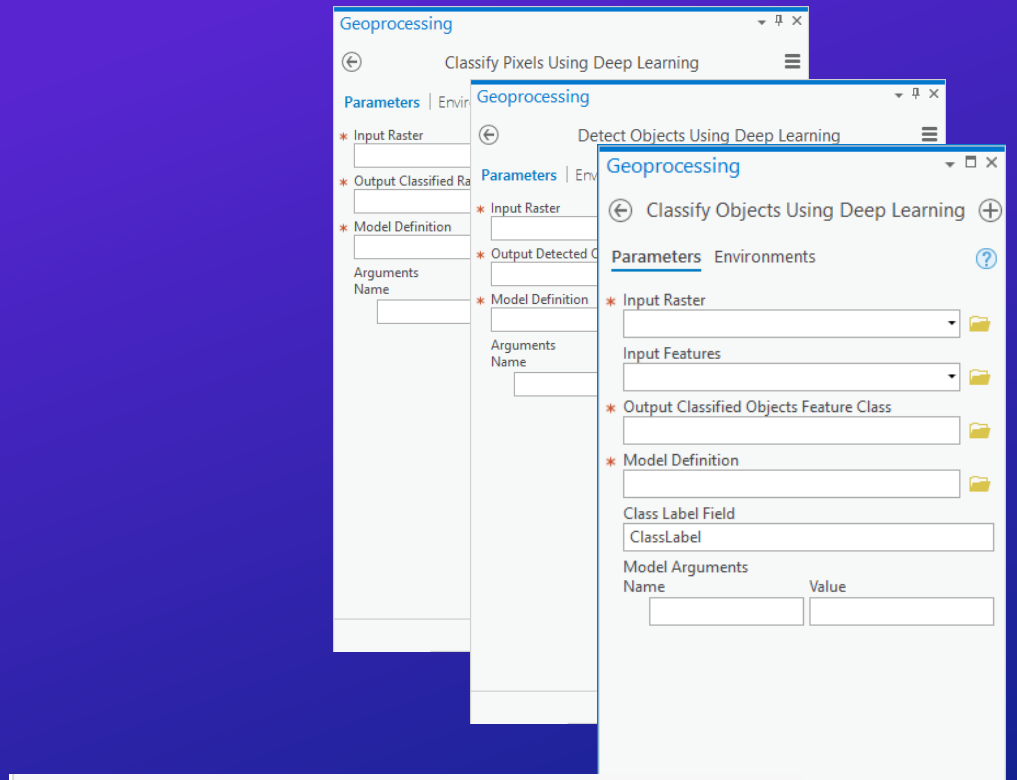


```
detect_objects(input_raster=input_layer,
               model=detect_objects_model,
               output_name="Well_Pad_Detection_Sentinel",
               context=context,
               gis=gis)
```

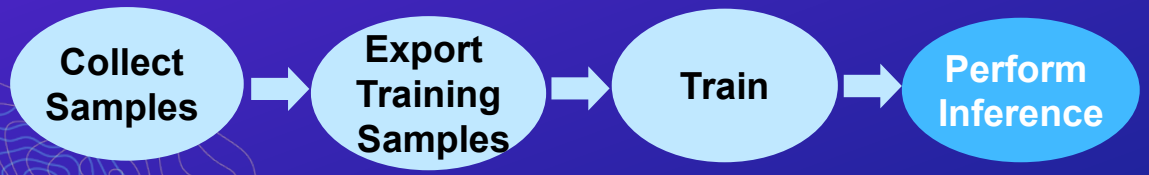

ArcGIS – Deep Learning Workflow

End-to-end deep learning workflow

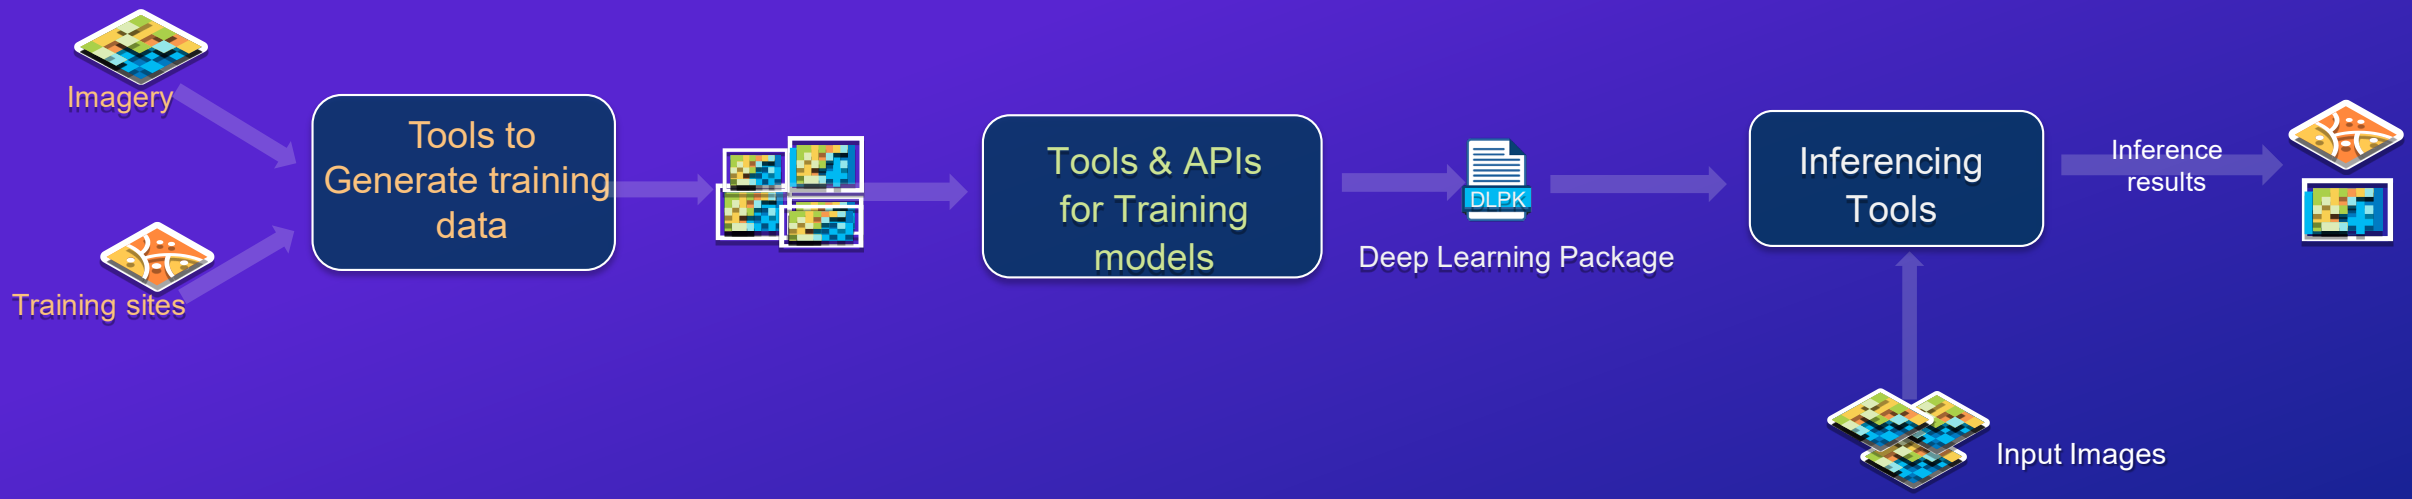

Tools to generate training samples

- Image Analyst in ArcGIS Pro
- Image Server on ArcGIS Enterprise

Model Training

- ArcGIS Pro
- Notebooks

Inferencing

- Image Analyst in ArcGIS Pro
- Image Server on ArcGIS Enterprise

Deep Learning Package

- Zip with a .dlpk file extension
  - Created by Train Deep Learning Model tool and arcgis.learn (ArcGIS API for Python)
- Contents of the dlpk
  - Model definition file (.emd)
  - Deep learning model file (framework specific)
  - Python Raster Function (.py, optional if using an out-of-the-box model)
- Can be shared across your organization

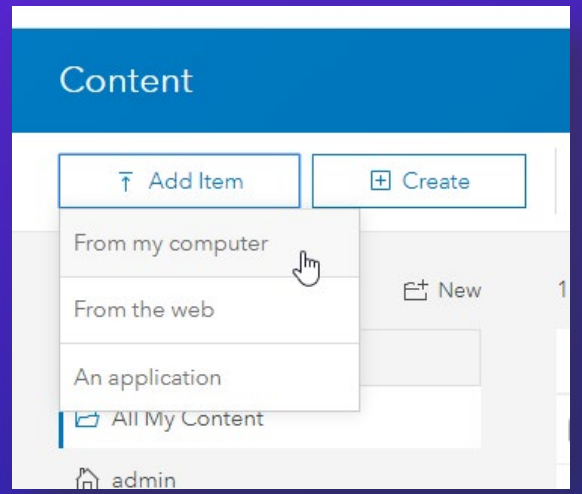

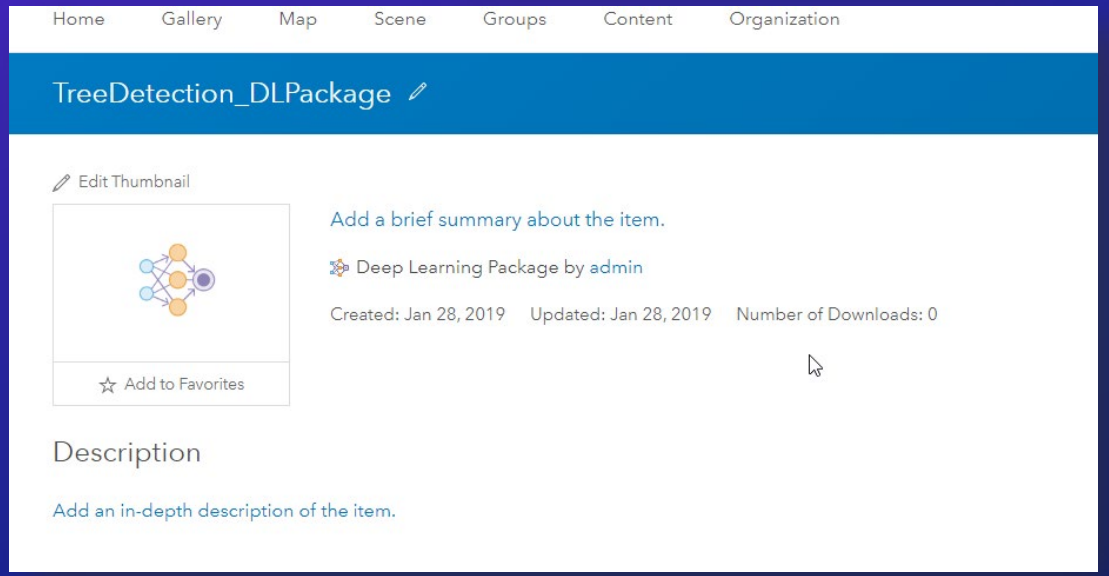

Supported Capabilities

	Labelling	Exporting Training Samples	Training Deep Learning Models	Inferencing
ArcGIS Pro	✓	✓	✓	✓
ArcGIS API for Python		✓	✓	✓
Map Viewer	✓	✗	✗	✓
Image Server		✓	✓	✓

Deep Learning Workflow in ArcGIS

End-to-end from raw imagery to structured information products

Image Management

Labelling

Data Prep

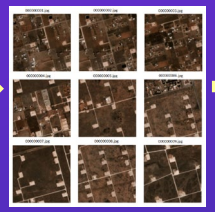

Train Model

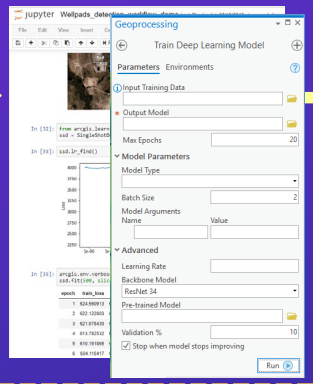

Inferencing

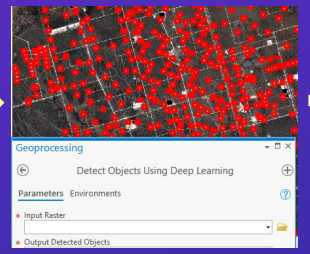

Analysis

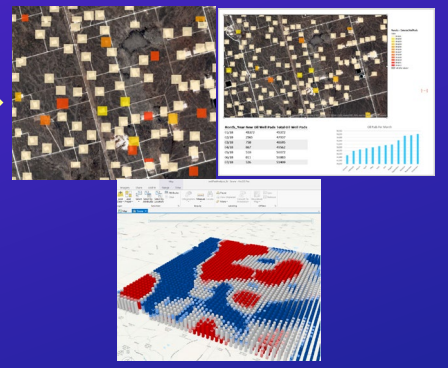

Field Mobility, Monitoring

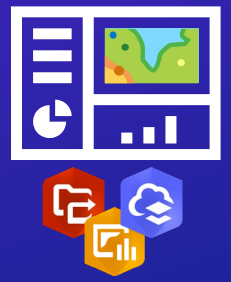

ArcGIS being used for each step of the deep learning workflow

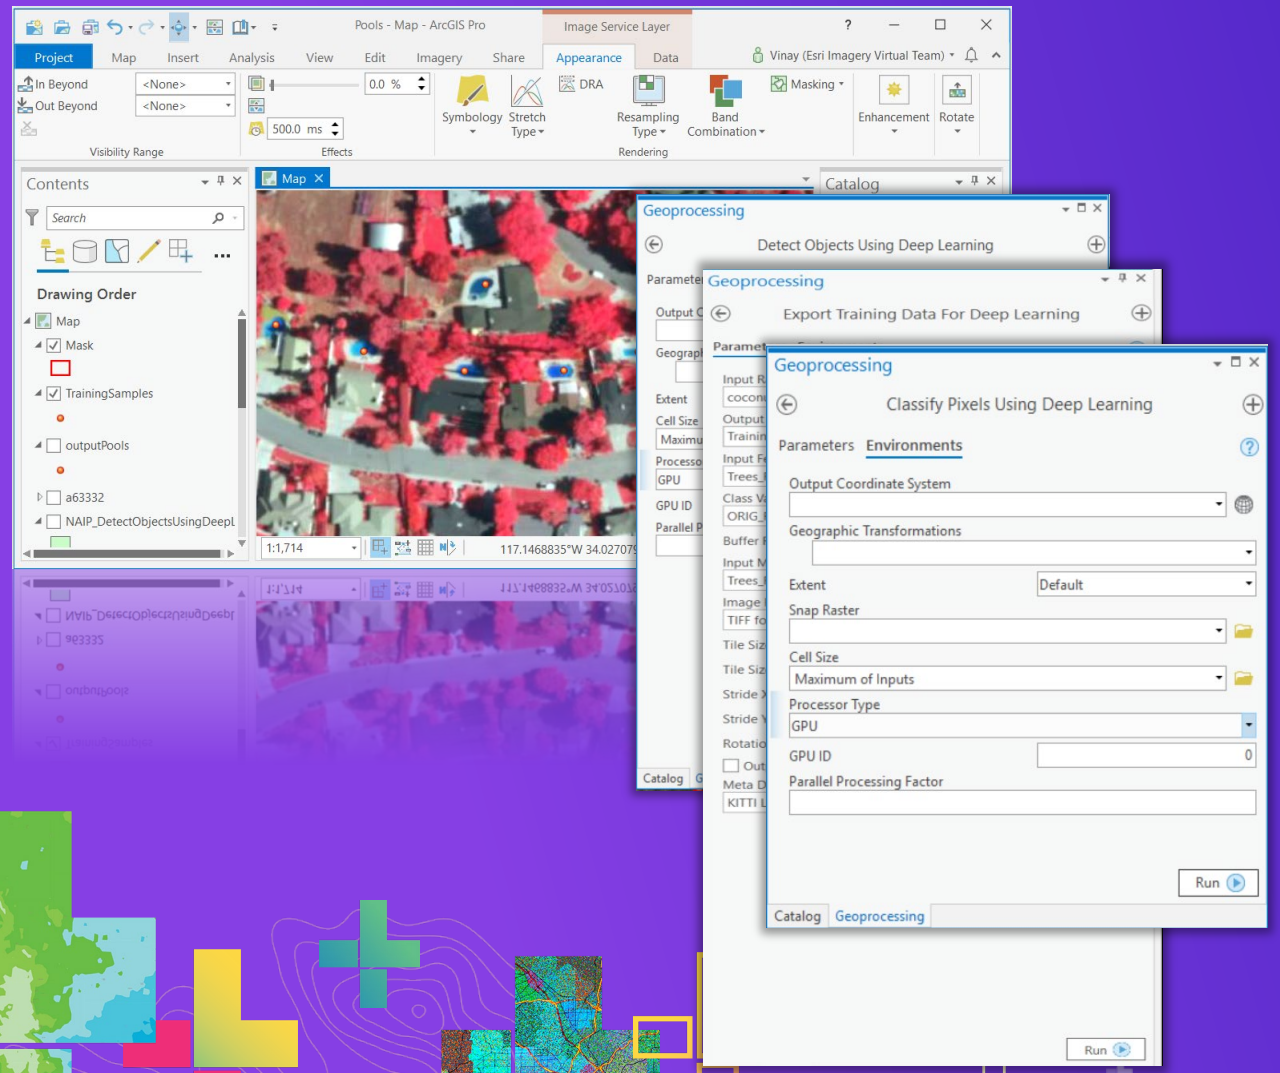

Demo

Using Deep Learning Tools in ArcGIS Pro – Well pad detection

ArcGIS + Notebooks = ♥

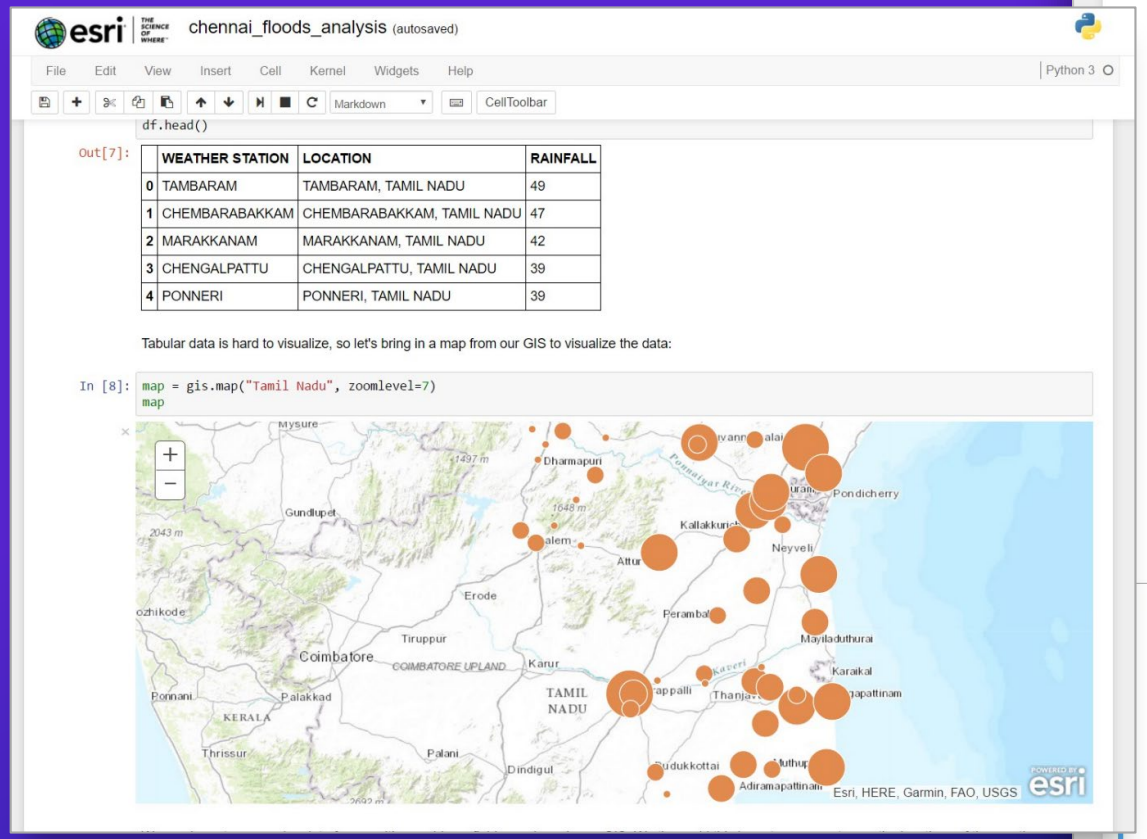

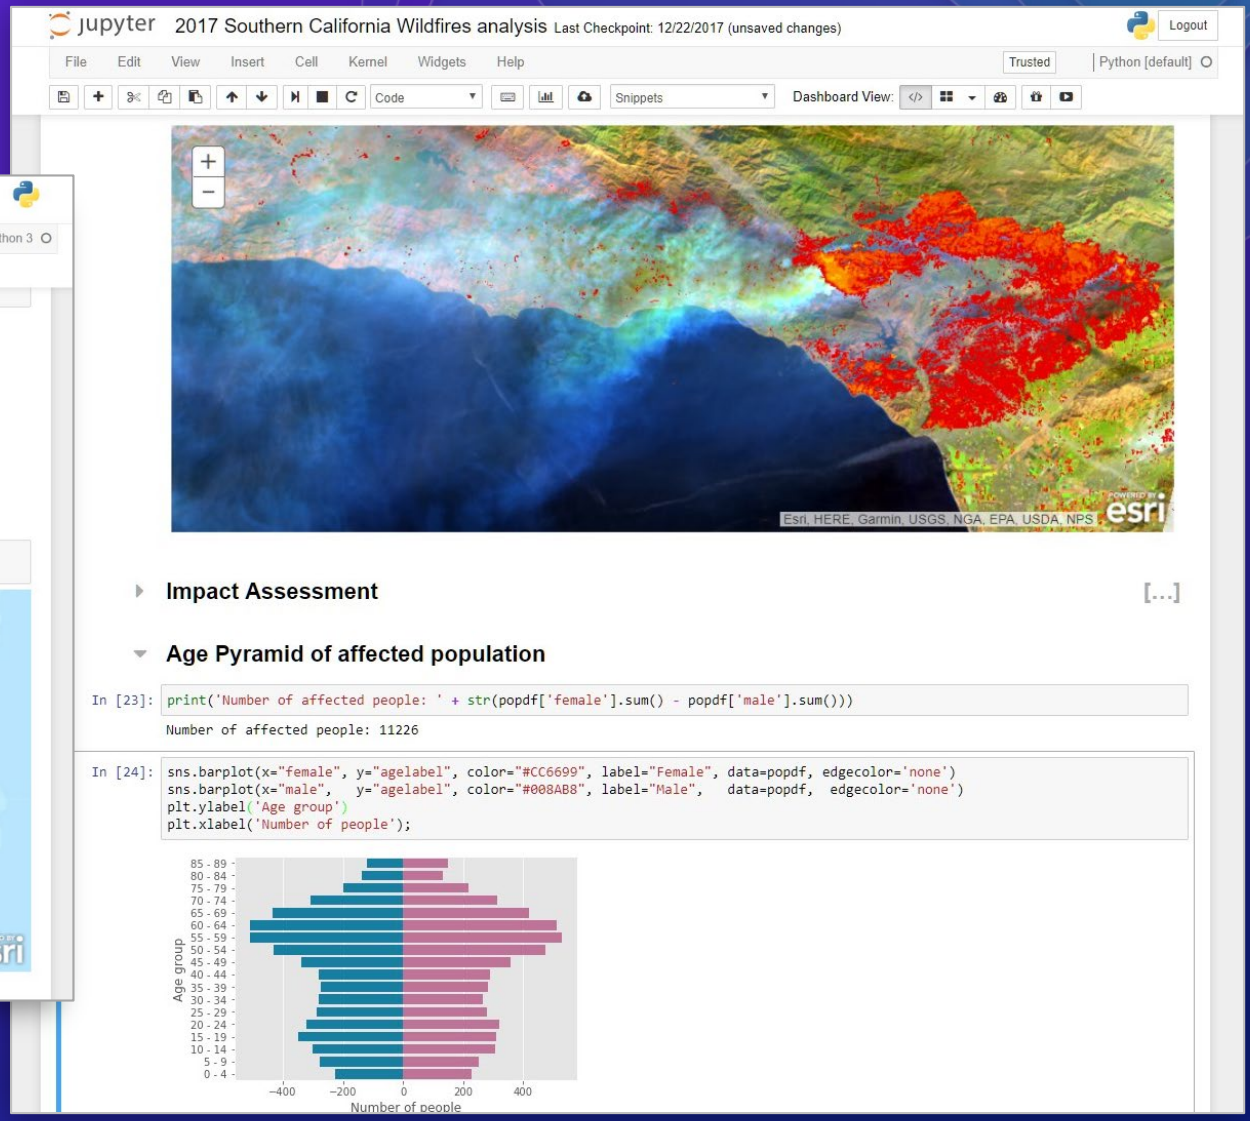

ArcGIS Notebooks sits at the intersection of ArcGIS and open data science

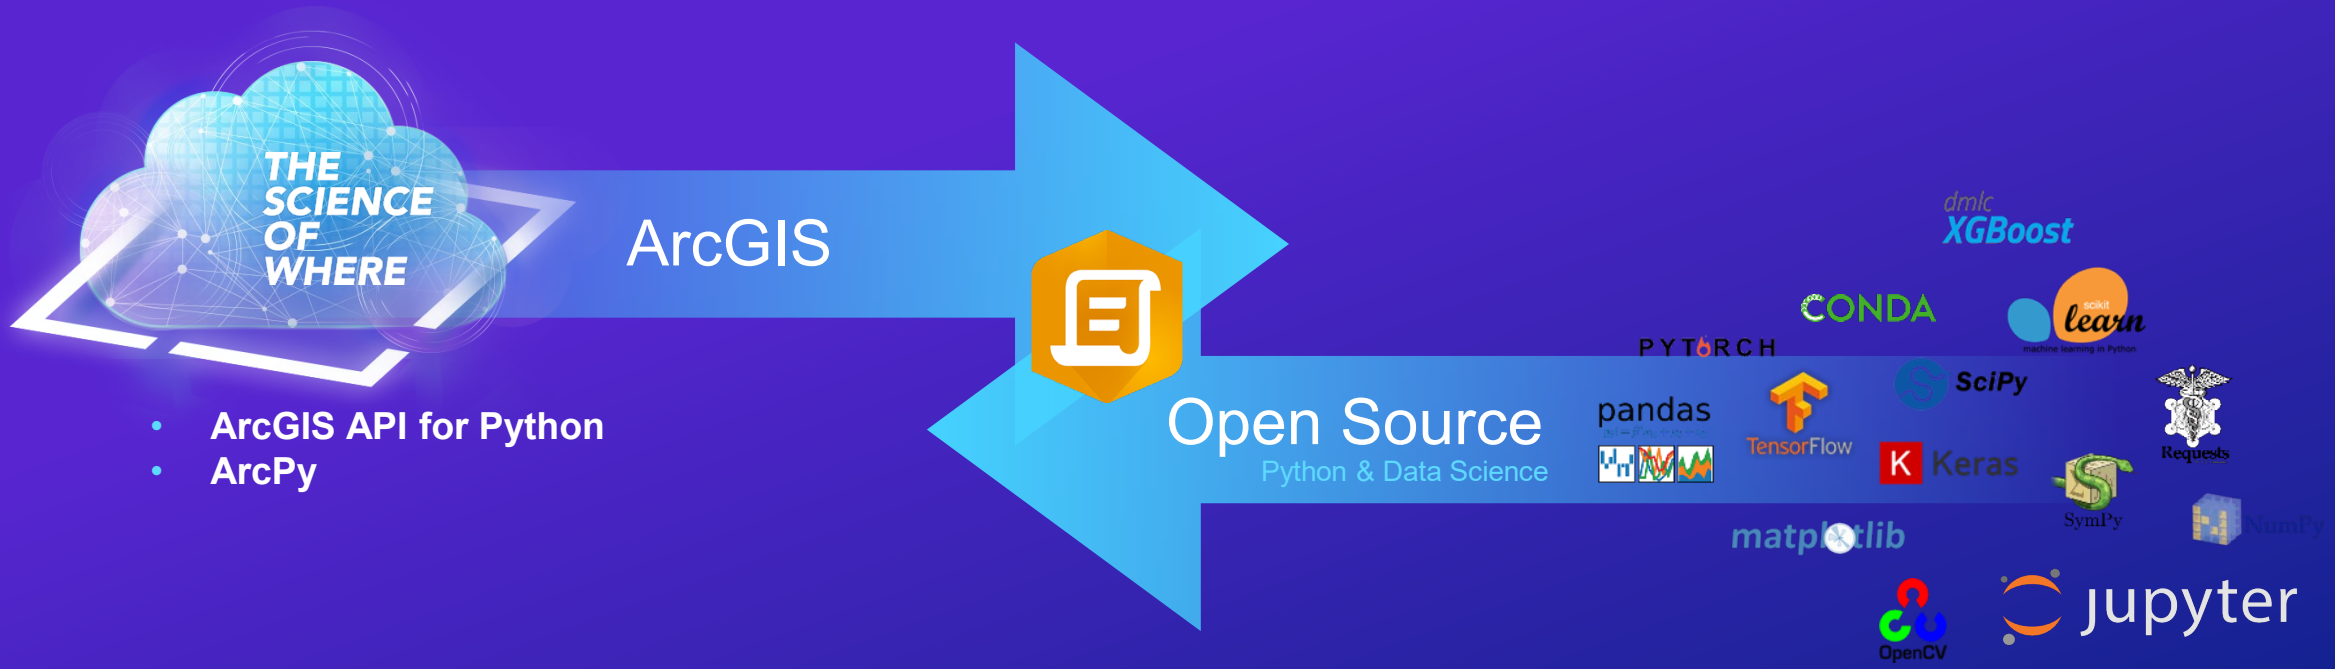

- ArcGIS API for Python
- ArcPy

ArcGIS Pro 2.5 seamlessly integrates Python Notebooks

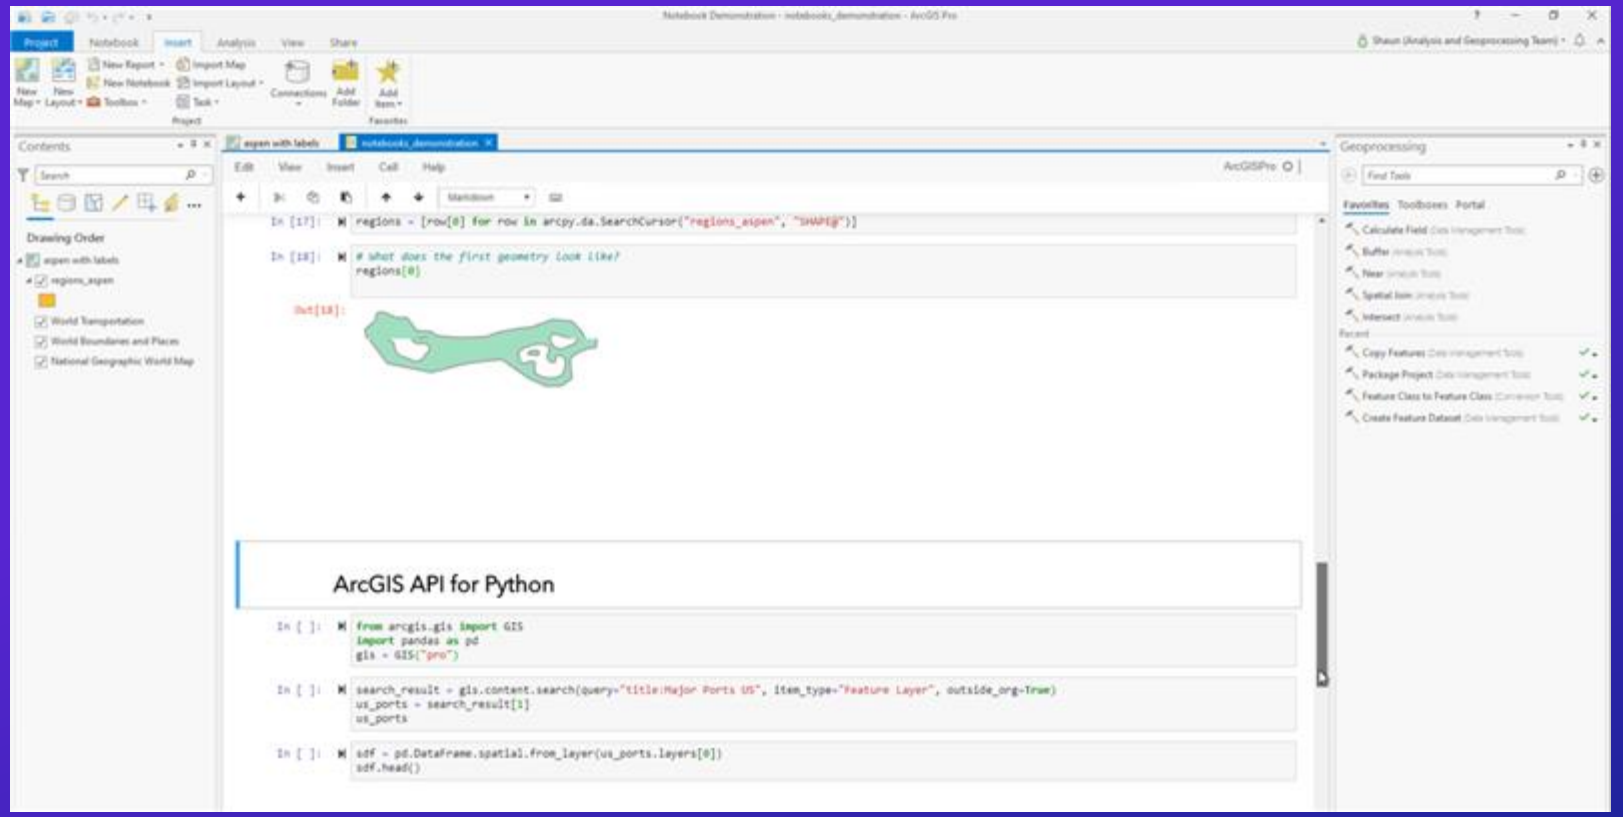

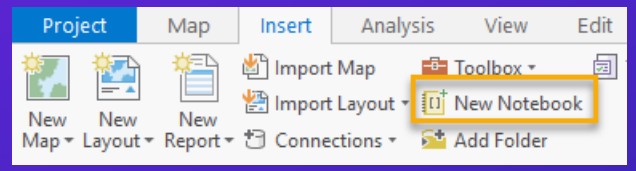

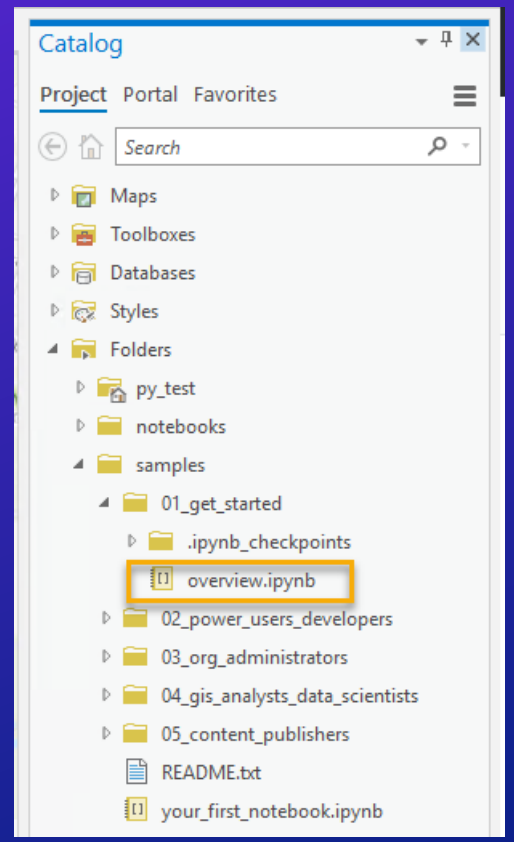

Scripting and Automation / **ArcGIS API for Python (1.6.1)**

ArcGIS API for Python

Install the API

Version 1.6.1 · May 16, 2019

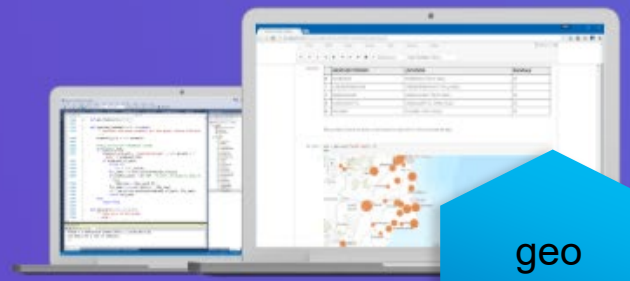

Home | Guide | Sample Notebooks | API Reference | Community

A powerful Python library for spatial analysis, mapping and GIS

ArcGIS API for Python is a Python library for working with maps and geospatial data, powered by web GIS. It provides simple and efficient tools for sophisticated vector and raster analysis, geocoding, map making, routing and directions, as well as for organizing and managing a GIS with users, groups and information items. In addition to working with your own data, the library enables access to ready to use maps and curated geographic data from Esri and other authoritative sources. It also integrates well with the scientific Python ecosystem and includes rich support for Pandas and Jupyter notebook.

Install the API | Get started | View samples

Understand your GIS

This "hello world" style notebook shows how to get started with the GIS and visualize its contents.

Manage your GIS

The ArcGIS API for Python provides APIs and samples for ArcGIS Online administrators to manage their online

Perform Spatial Analysis

Call sophisticated spatial analysis tools that work with online content, using a few lines of code.

ArcGIS API for Python

arcgis.learn module

The `arcgis.learn` module in **ArcGIS API for Python** enables Python developers and data scientists to **easily train and use deep learning models with a simple, intuitive API.**

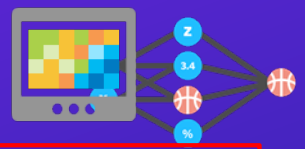

Train Models

Before

8. Conver Label to One Hot Vector

9. Define IoU Metric

10. Define Custom Loss Function

11. Set Parameters before Training

Land Cover Classes

- Dozens of lines of Code
- Installing External DL Frameworks
- HARD!

After

- 3-5 lines
- No Installation (ArcGIS Pro & Notebooks)
- EASY

Train SingleShotDetector Model

```
from arcgis.learn import SingleShotDetector

ssd = SingleShotDetector(data, grids=[9], zooms=[1.0], ratios=[[1.0, 1.0]])
```

```
ssd.fit(10, lr=slice(1e-3, 1e-2))
```

ArcGIS API for Python

Not just “training”!

Data Preparation

- arcgis.learn.export_training_data
- arcgis.learn.prepare_data

Training DL Models

- arcgis.learn.SingleShotDetector
- arcgis.learn.UnetClassifier
- arcgis.learn.FeatureClassifier
- arcgis.learn.PSPNetClassifier
- arcgis.learn.RetinaNet
- arcgis.learn.MaskRCNN

Model Management

- arcgis.learn.list_models
- arcgis.learn.Model
  - Model.install
  - Model.uninstall
  - Model.query_info

Inference APIs

- arcgis.learn.detect_objects
- arcgis.learn.classify_pixels
- arcgis.learn.classify_objects

Advantages of arcgis.learn

- Closely integrated with ArcGIS
  - Directly consumes exported training data from ArcGIS (no custom preprocessing)
  - Saved models (DLPKs) are directly usable in ArcGIS
    - No custom postprocessing of model output
    - Image space to map space conversion automatically handled
  - Preserves symbology
- Consistent API (prepare_data(), fit(), save() to train model, show_batch/show_results to visualize)
- Performs data augmentations suitable for satellite imagery
- Extensible using fast.ai transforms, custom loss functions, model backbones
- Fast.ai goodies: Automatic learning rate finder, transfer learning, early stopping, checkpointing, one-cycle learning
- Model metrics, sample results and training details are stored along with the model
- Padding support, multi-gpu training, CPU/GPU support, predict on videos, multispectral imagery*…

Things you can do today with arcgis.learn

Object Detection, Pixel Classification, Feature Classification, Instance Segmentation

Damaged Structures

Building Footprints

Land Cover

Catfish

Brick Kilns

Roads

Oil Pads

Palm trees

Refugee Camps

Sinkholes

Swimming Pools

Training models using `arcgis.learn`

Types of Deep Learning Models & their applications to GIS

Image Classification

Assign a label to a given image

Object Classification

Assign a label to a given feature

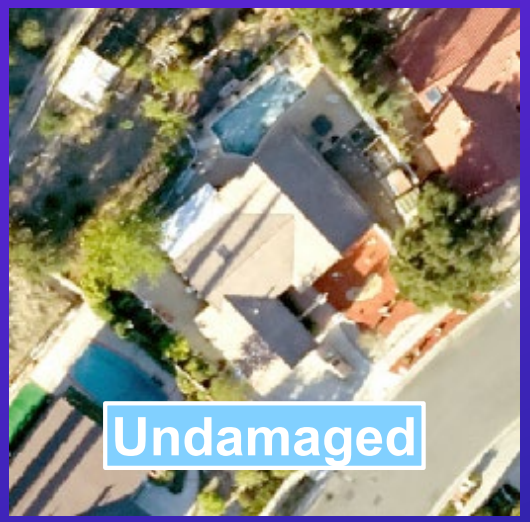

Models (from torchvision):

- Inception
- ResNet
- VGG…

Applications:

- Damaged building classification
- Clean or ‘green’ pools…

Demo

Building Damage classification

Semantic Segmentation

Assign a label to each pixel

Cat

Ground

Sky

Models:

- UNetClassifier
- PSPNetClassifier

Pixel Classification

Turf/Grass

Building

Water

Applications:

- Land Cover Classification
- Pervious/Impervious mapping…

Demo

Building Footprint Extraction

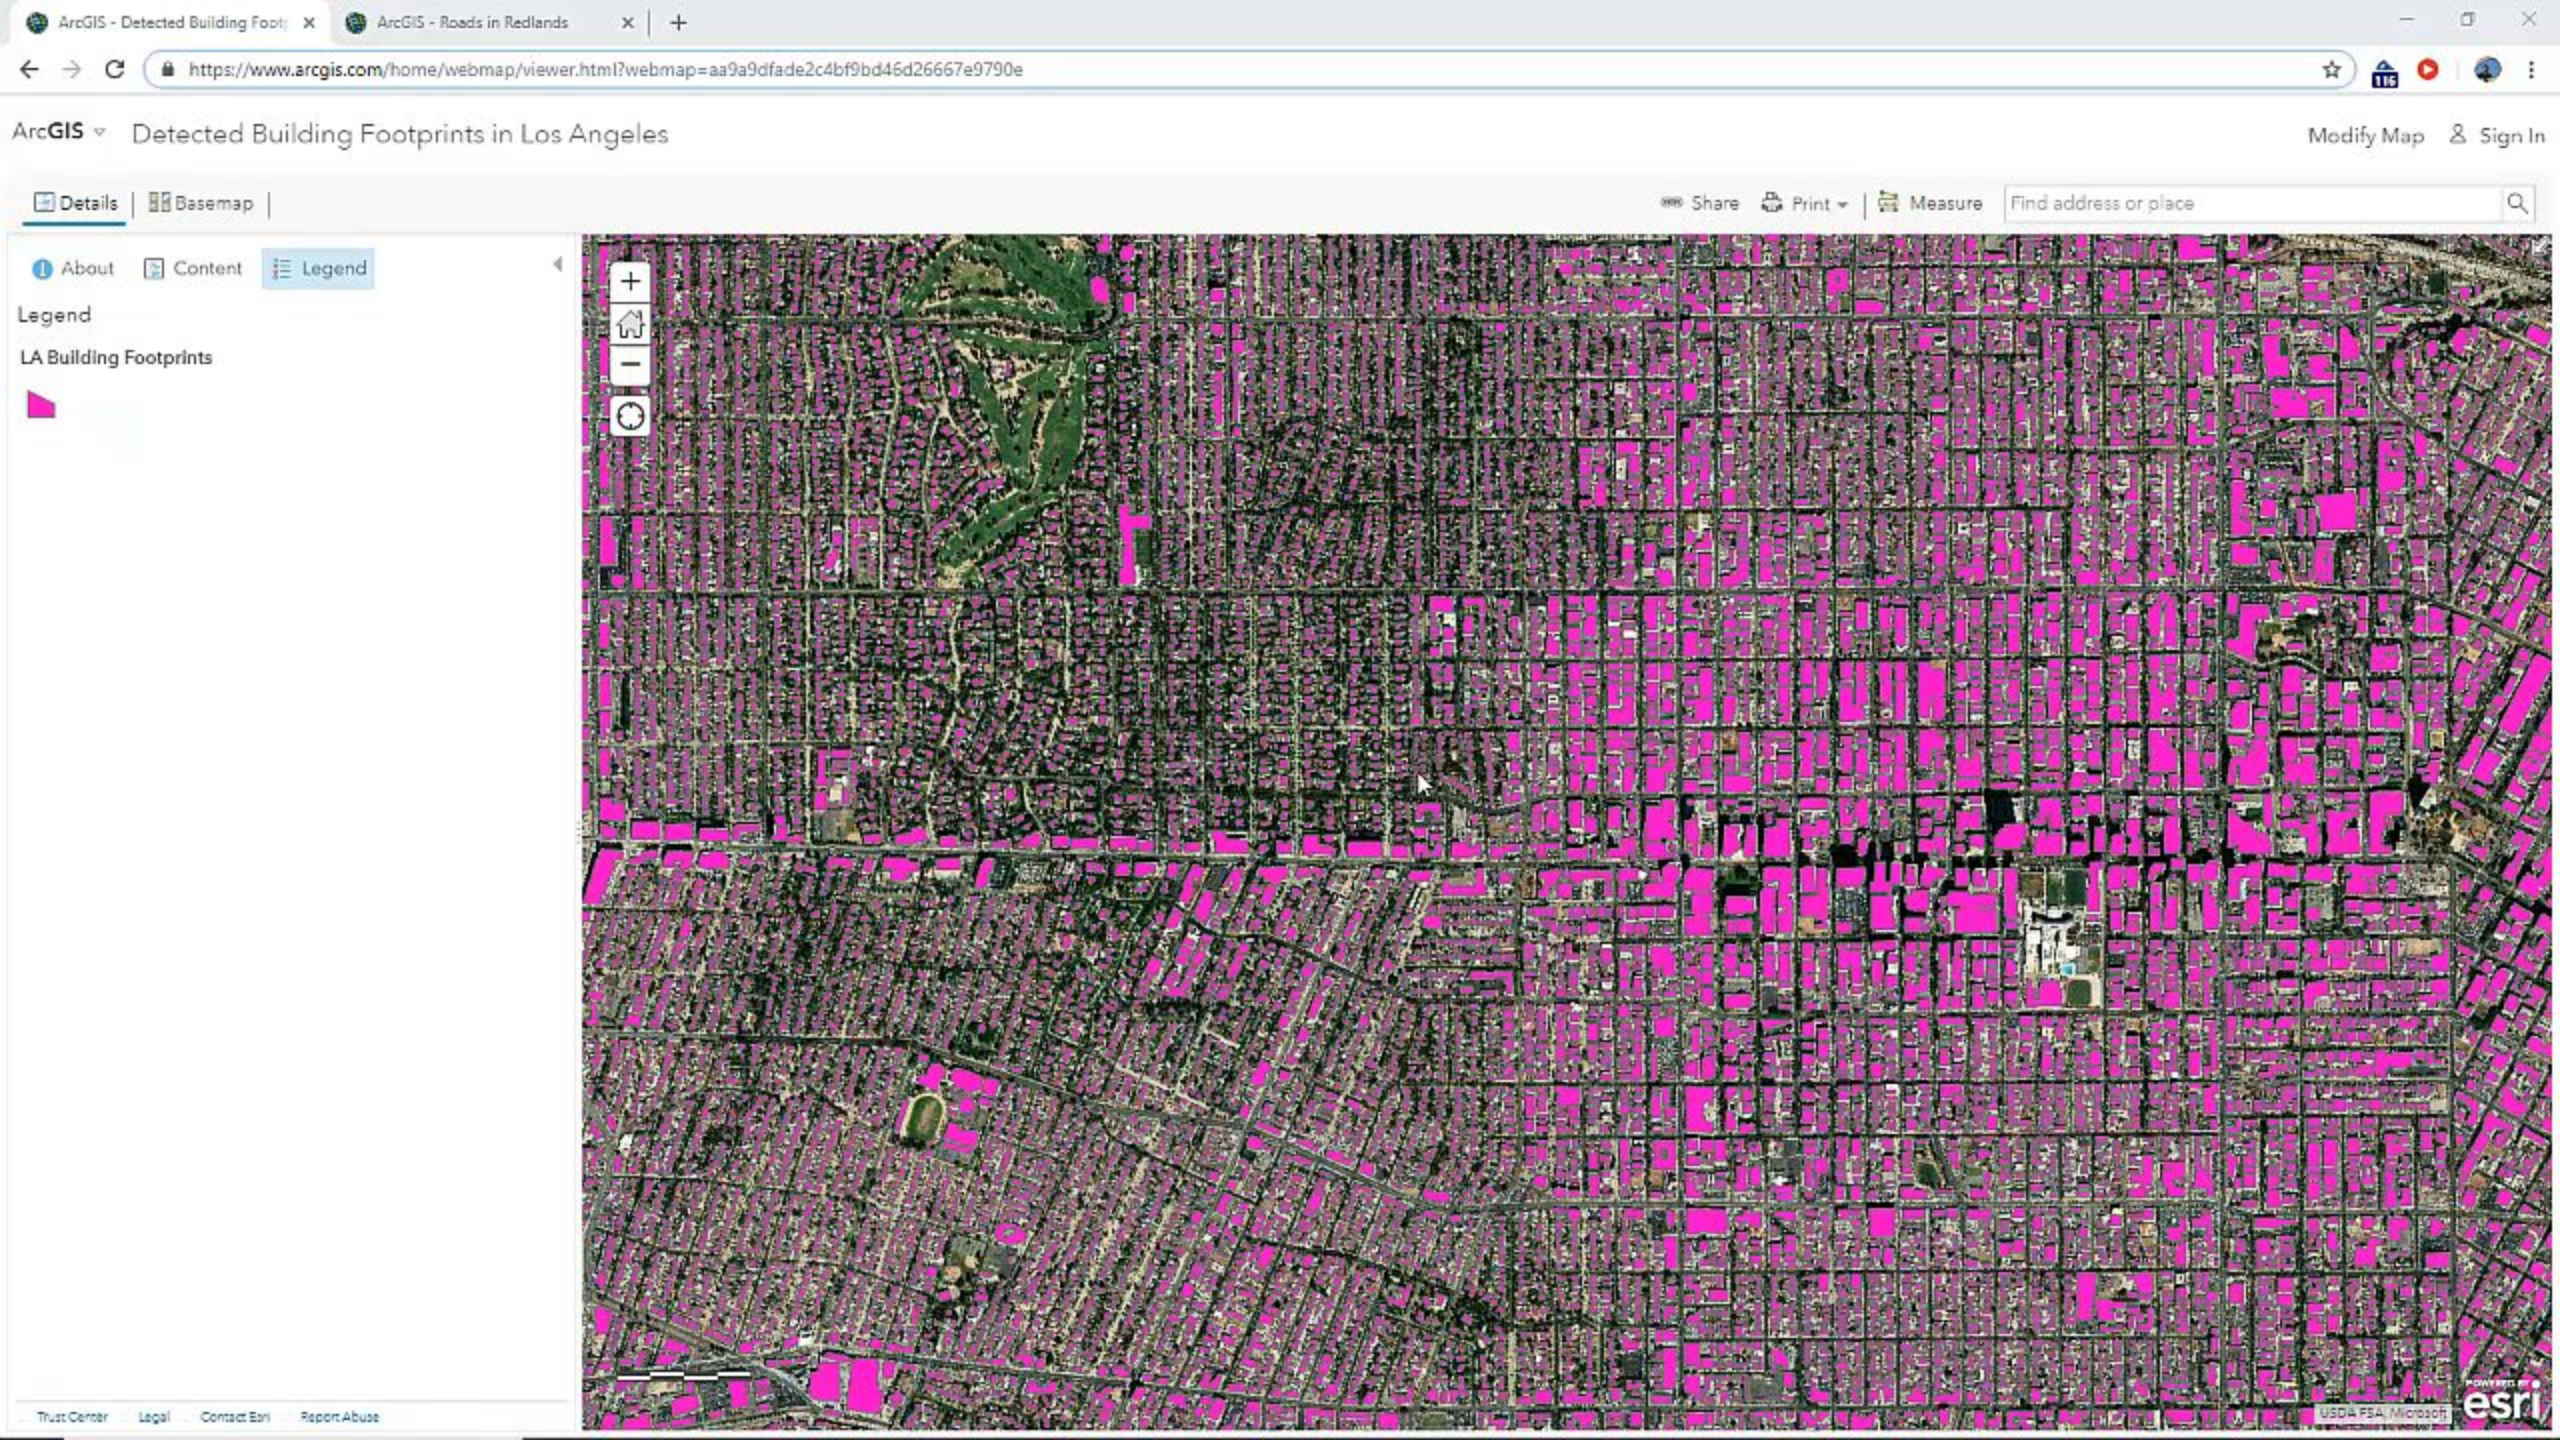
ArcGIS - Detected Building Foot
ArcGIS - Roads in Redlands
https://www.arcgis.com/home/webmap/viewer.html?webmap=aa9a9dfade2c4bf9bd46d26667e9790e
ArcGIS
Detected Building Footprints in Los Angeles
Modify Map
Sign In
Details
Basemap
Share
Print
Measure
Find address or place
About
Content
Legend
Legend
LA Building Footprints
Trust Center
Legal
Contact Esri
Report Abuse
USDA FSA, Microsoft
esri

Object Detection

Find objects and their location (bounding boxes)

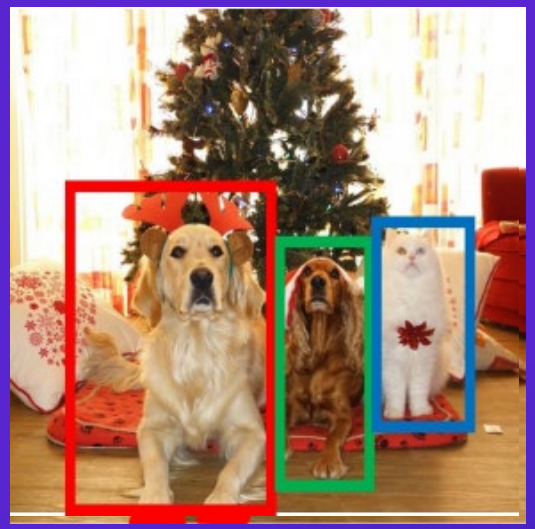

Models:

- SingleShotDetector
- RetinaNet

Applications:

- Detect trees, cars, airplanes, …

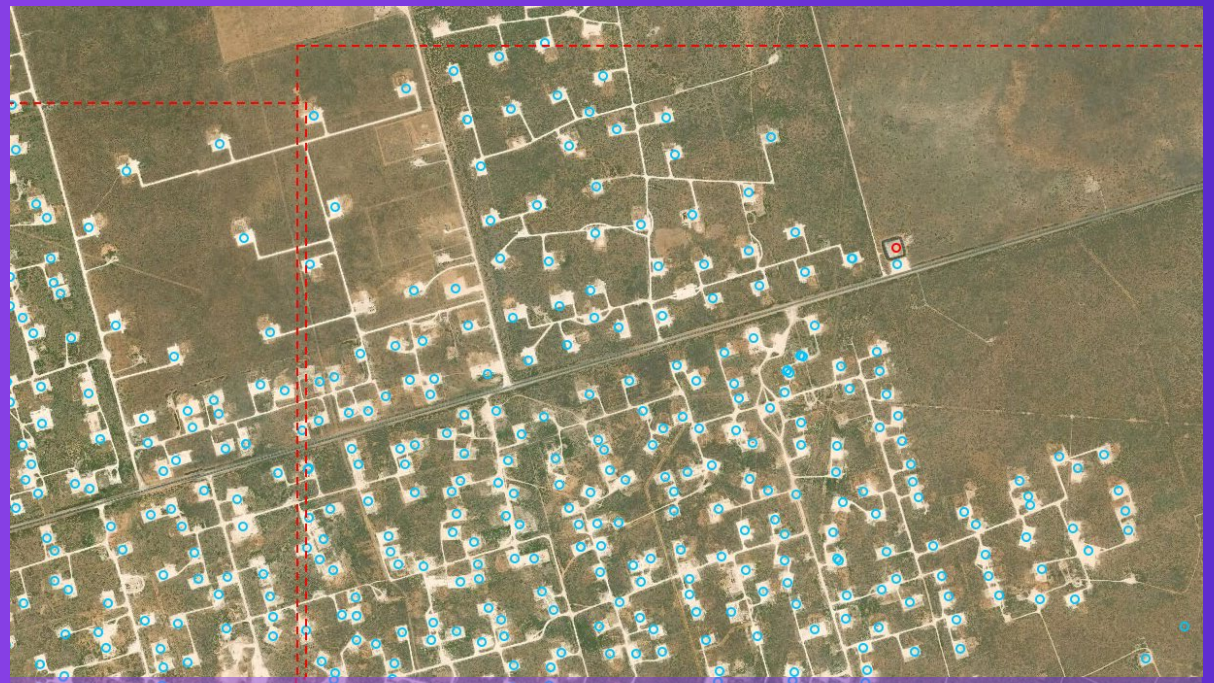

Demo

Detecting Well Pads

Instance Segmentation

Find objects and their *precise* locations (masks or polygonal features)

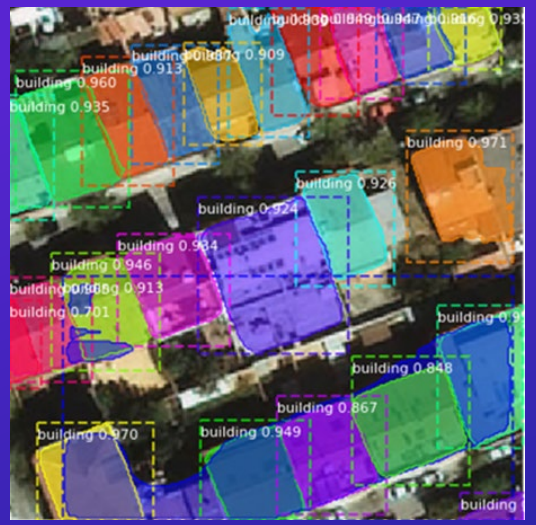

Models:

- MaskRCNN

Applications:

- Building footprint extraction

Demo

Contiguous Building Footprints

Text / NLP

- arcgis.learn – model for extracting location and other entities
  - EntityExtractor model

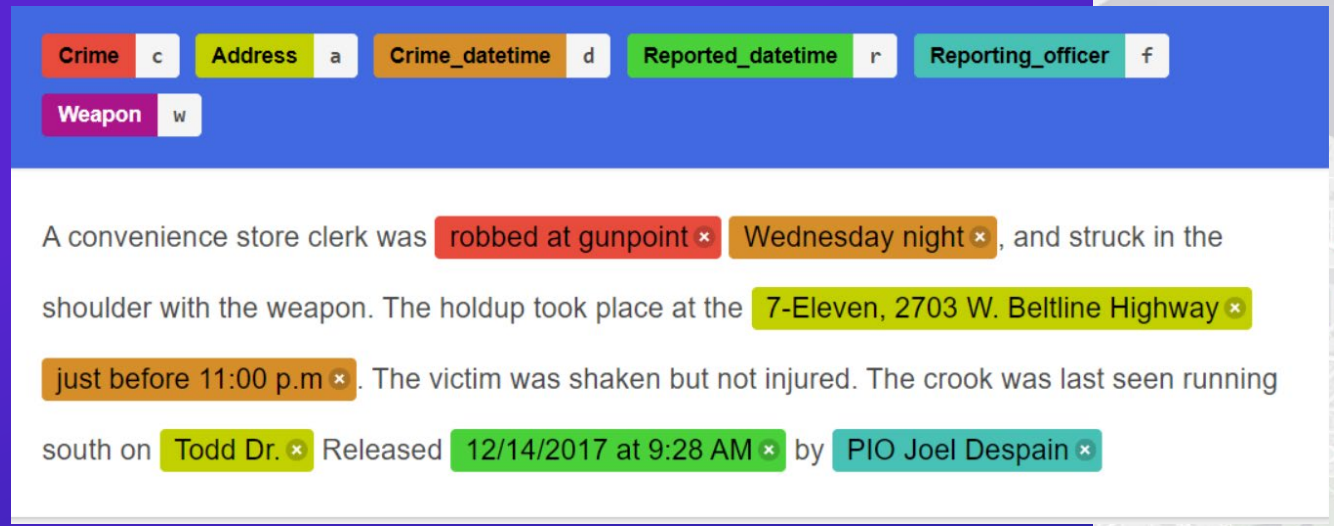

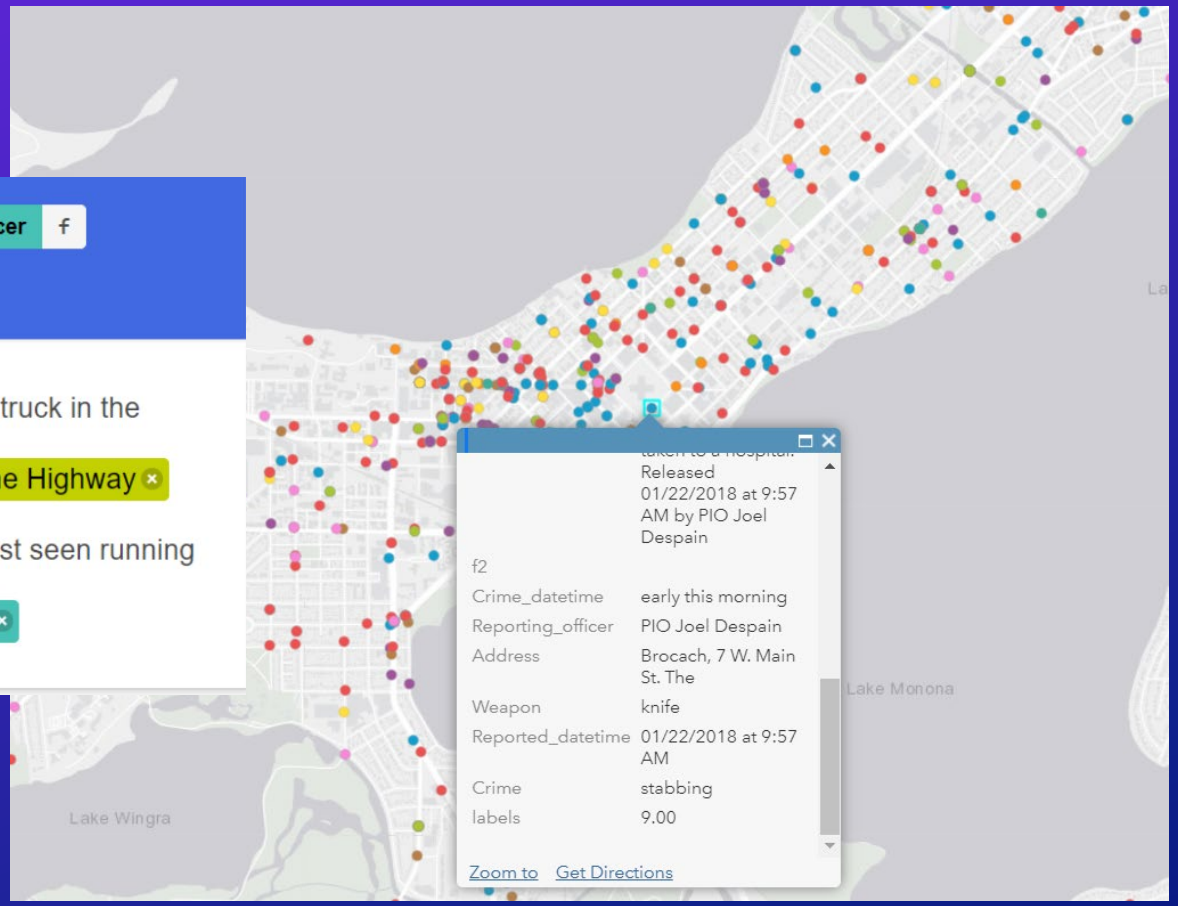

ArcGIS Enterprise for Scaling Deep Learning

- Leverage Raster Analytics to scale inferencing
- All desktop inferencing tools are accessible through enterprise
- Clients to invoke distributed inferencing – map viewer, pro, notebooks
- Multi GPU support
- Requires the ArcGIS Image Server license

ArcGIS Enterprise Deep Learning Tools / Services

Pro + Server Tools	
ExportTrainingDataforDeepLearning	*Uses a remote sensing image to convert labeled vector or raster data into deep learning training datasets. The output is a folder of image chips and a folder of metadata files.*
DetectObjectsUsingDeepLearning	*Runs a trained deep learning model on an input raster to produce a feature class containing the objects it finds. The features can be bounding boxes or polygons around the objects found, or points at the centers of the objects.*
ClassifyPixelsUsingDeepLearning	*Runs a trained deep learning model on an input raster to produce a classified raster with each valid pixel having a class label assigned.*
ClassifyObjectsUsingDeepLearning	*Runs a trained deep learning model on an input raster and feature class to produce a classified feature class*
TrainDeepLearningModel	*Enables training deep learning models*
QueryDeepLearningModelInfo	*Extracts the model specific settings from the model package item or model definition file.*
InstallDeepLearningModel	*Installs the model package item from portal to the Raster Analysis Image Server.*
Server Only Tools	
UninstallDeepLearningModel	*Uninstalls the model package from portal to the Raster Analysis Image Server*
ListDeepLearningModels	*Lists all the installed model packages on the Raster Analysis Image Server*

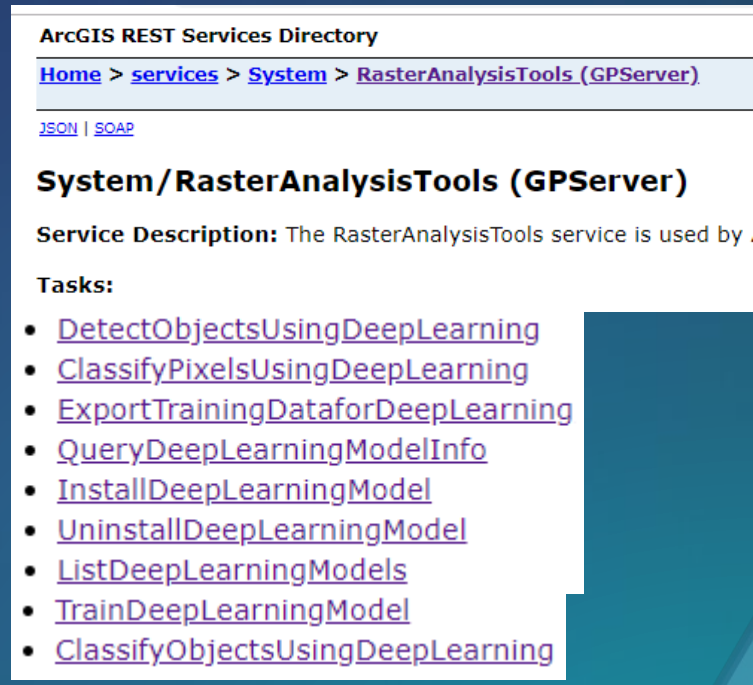

ArcGIS Enterprise Deep Learning – System Architecture

Scenario – Detecting Building Footprints at Scale

- Aim: Determine benchmark on a fixed dataset with a fixed inference method
  - Denver World Imagery
    - 4,727 blobs
    - 13,375,880,336 bytes
    - 1,669,001,963 m^2 (0.6m)
    - 2048x2048 cells per blob (TIF)
    - Custom Mask R-CNN for detecting building footprints
  - Florida aerial imagery from NOAA
  - xView baseline model for detecting building footprints
- Images stored in Azure Blob Storage
- Speedup achievable from scaling up # of Raster Analysis Servers
  - Observe relationship

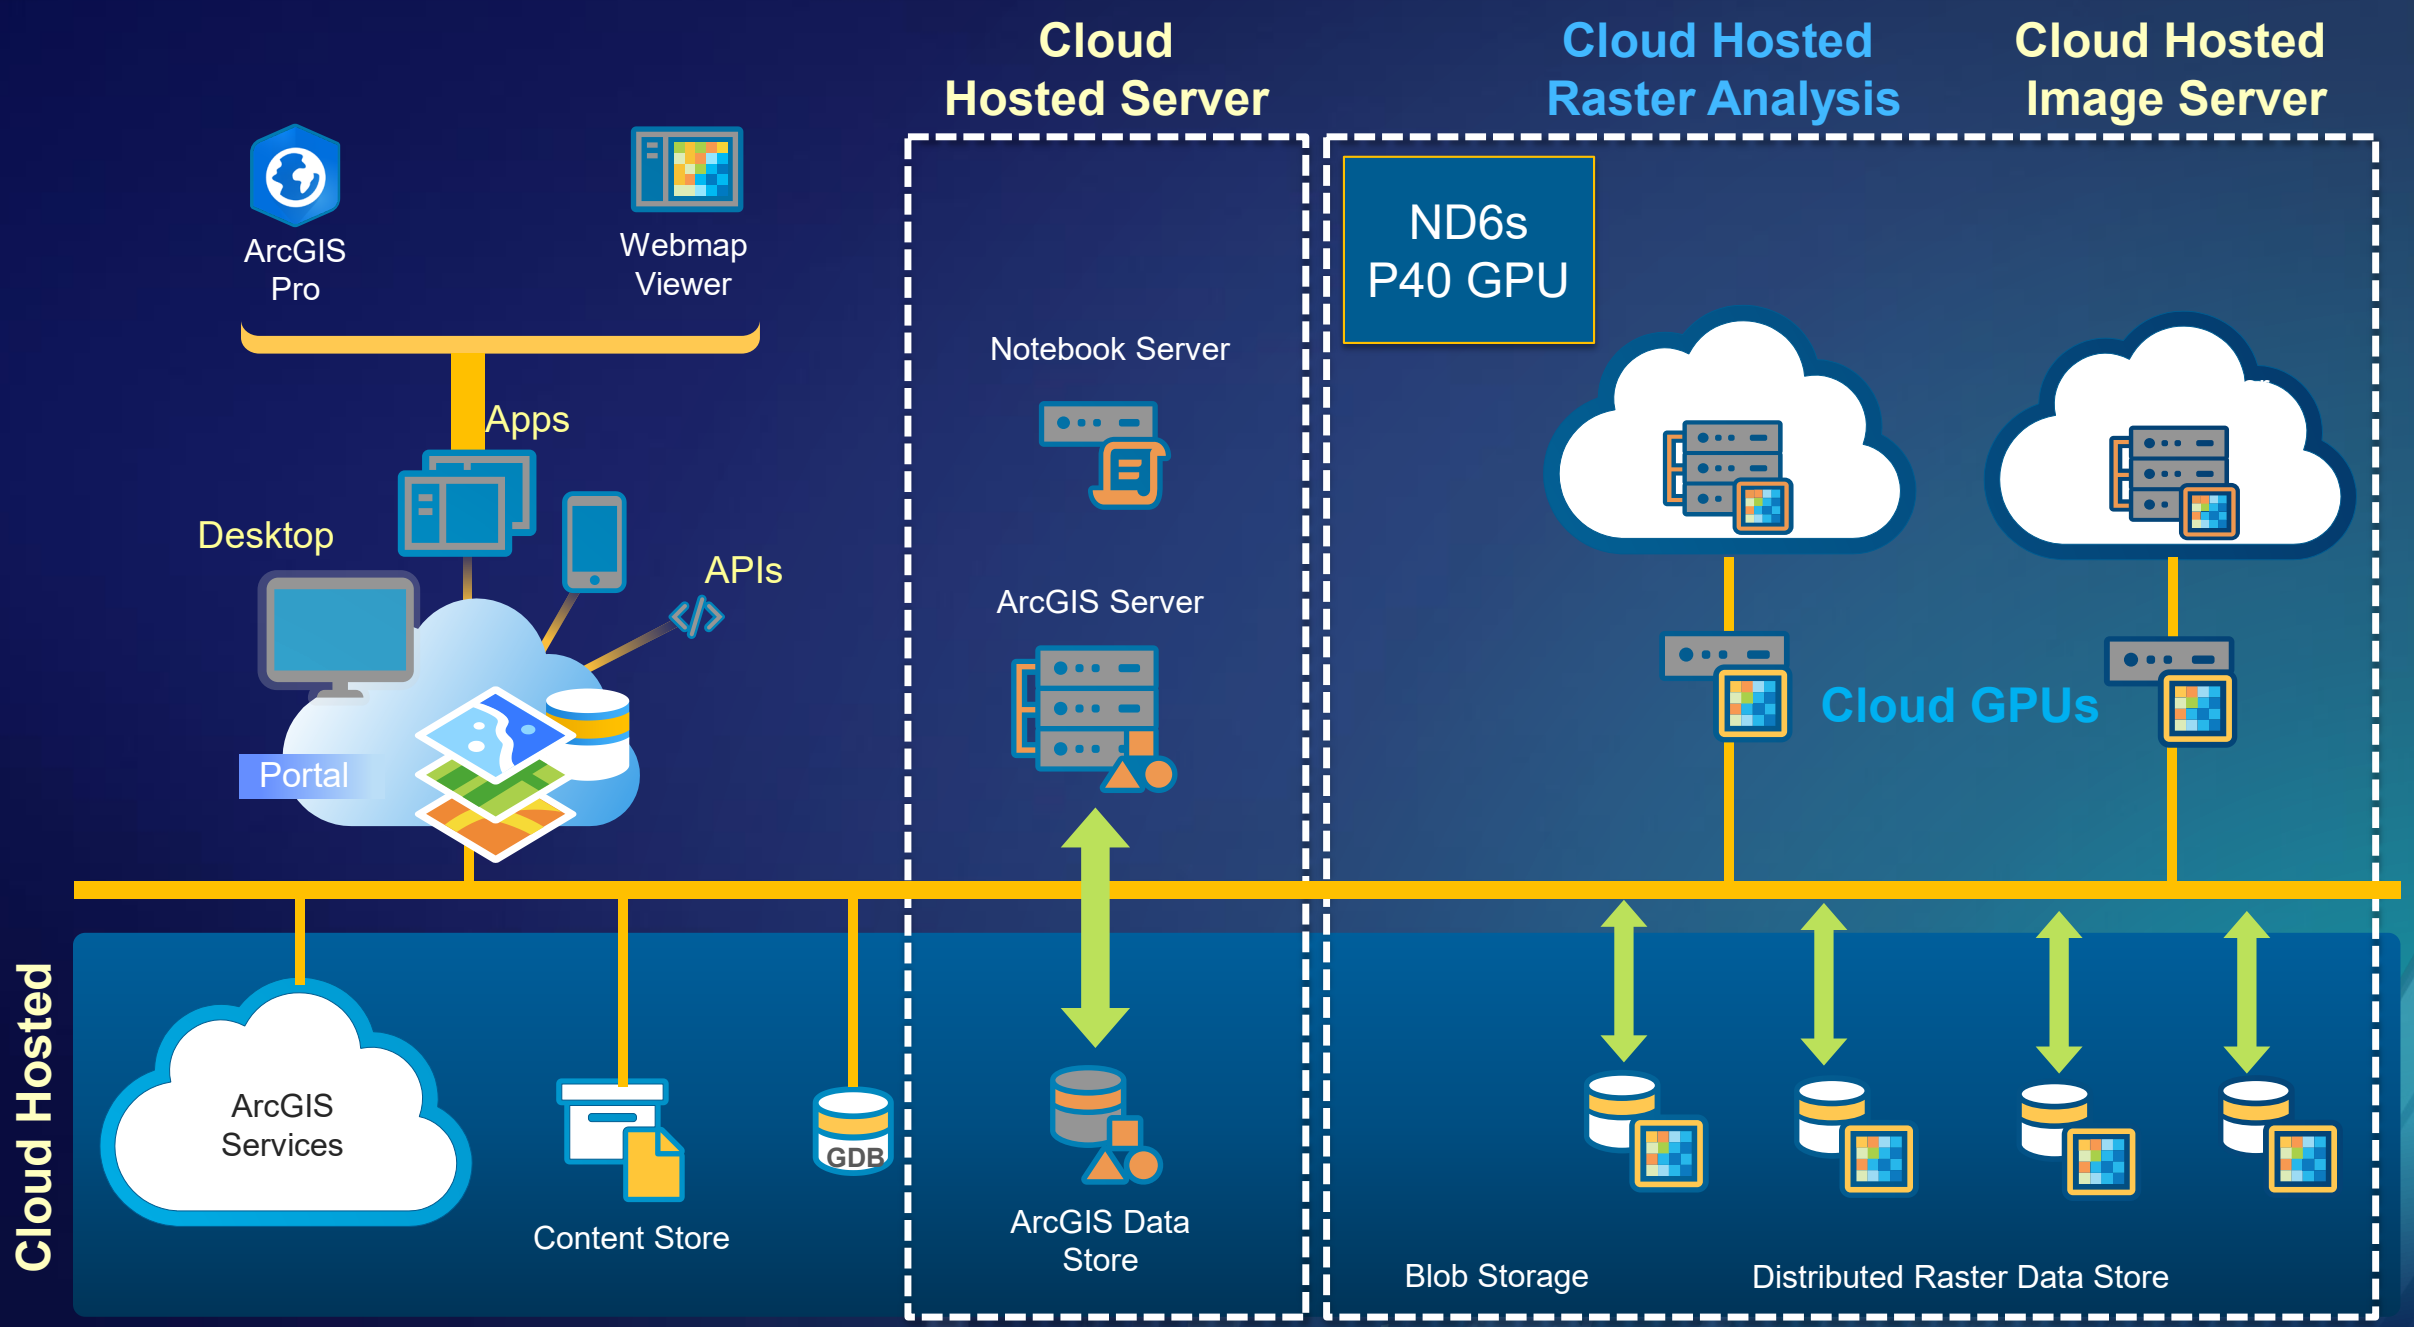
Cloud Hosted Server
Cloud Hosted Raster Analysis
Cloud Hosted Image Server
ArcGIS Pro
Webmap Viewer
Apps
Desktop
APIs
Portal
Notebook Server
ArcGIS Server
ND6s P40 GPU
Cloud GPUs
Cloud Hosted
ArcGIS Services
Content Store
GDB
ArcGIS Data Store
Blob Storage
Distributed Raster Data Store

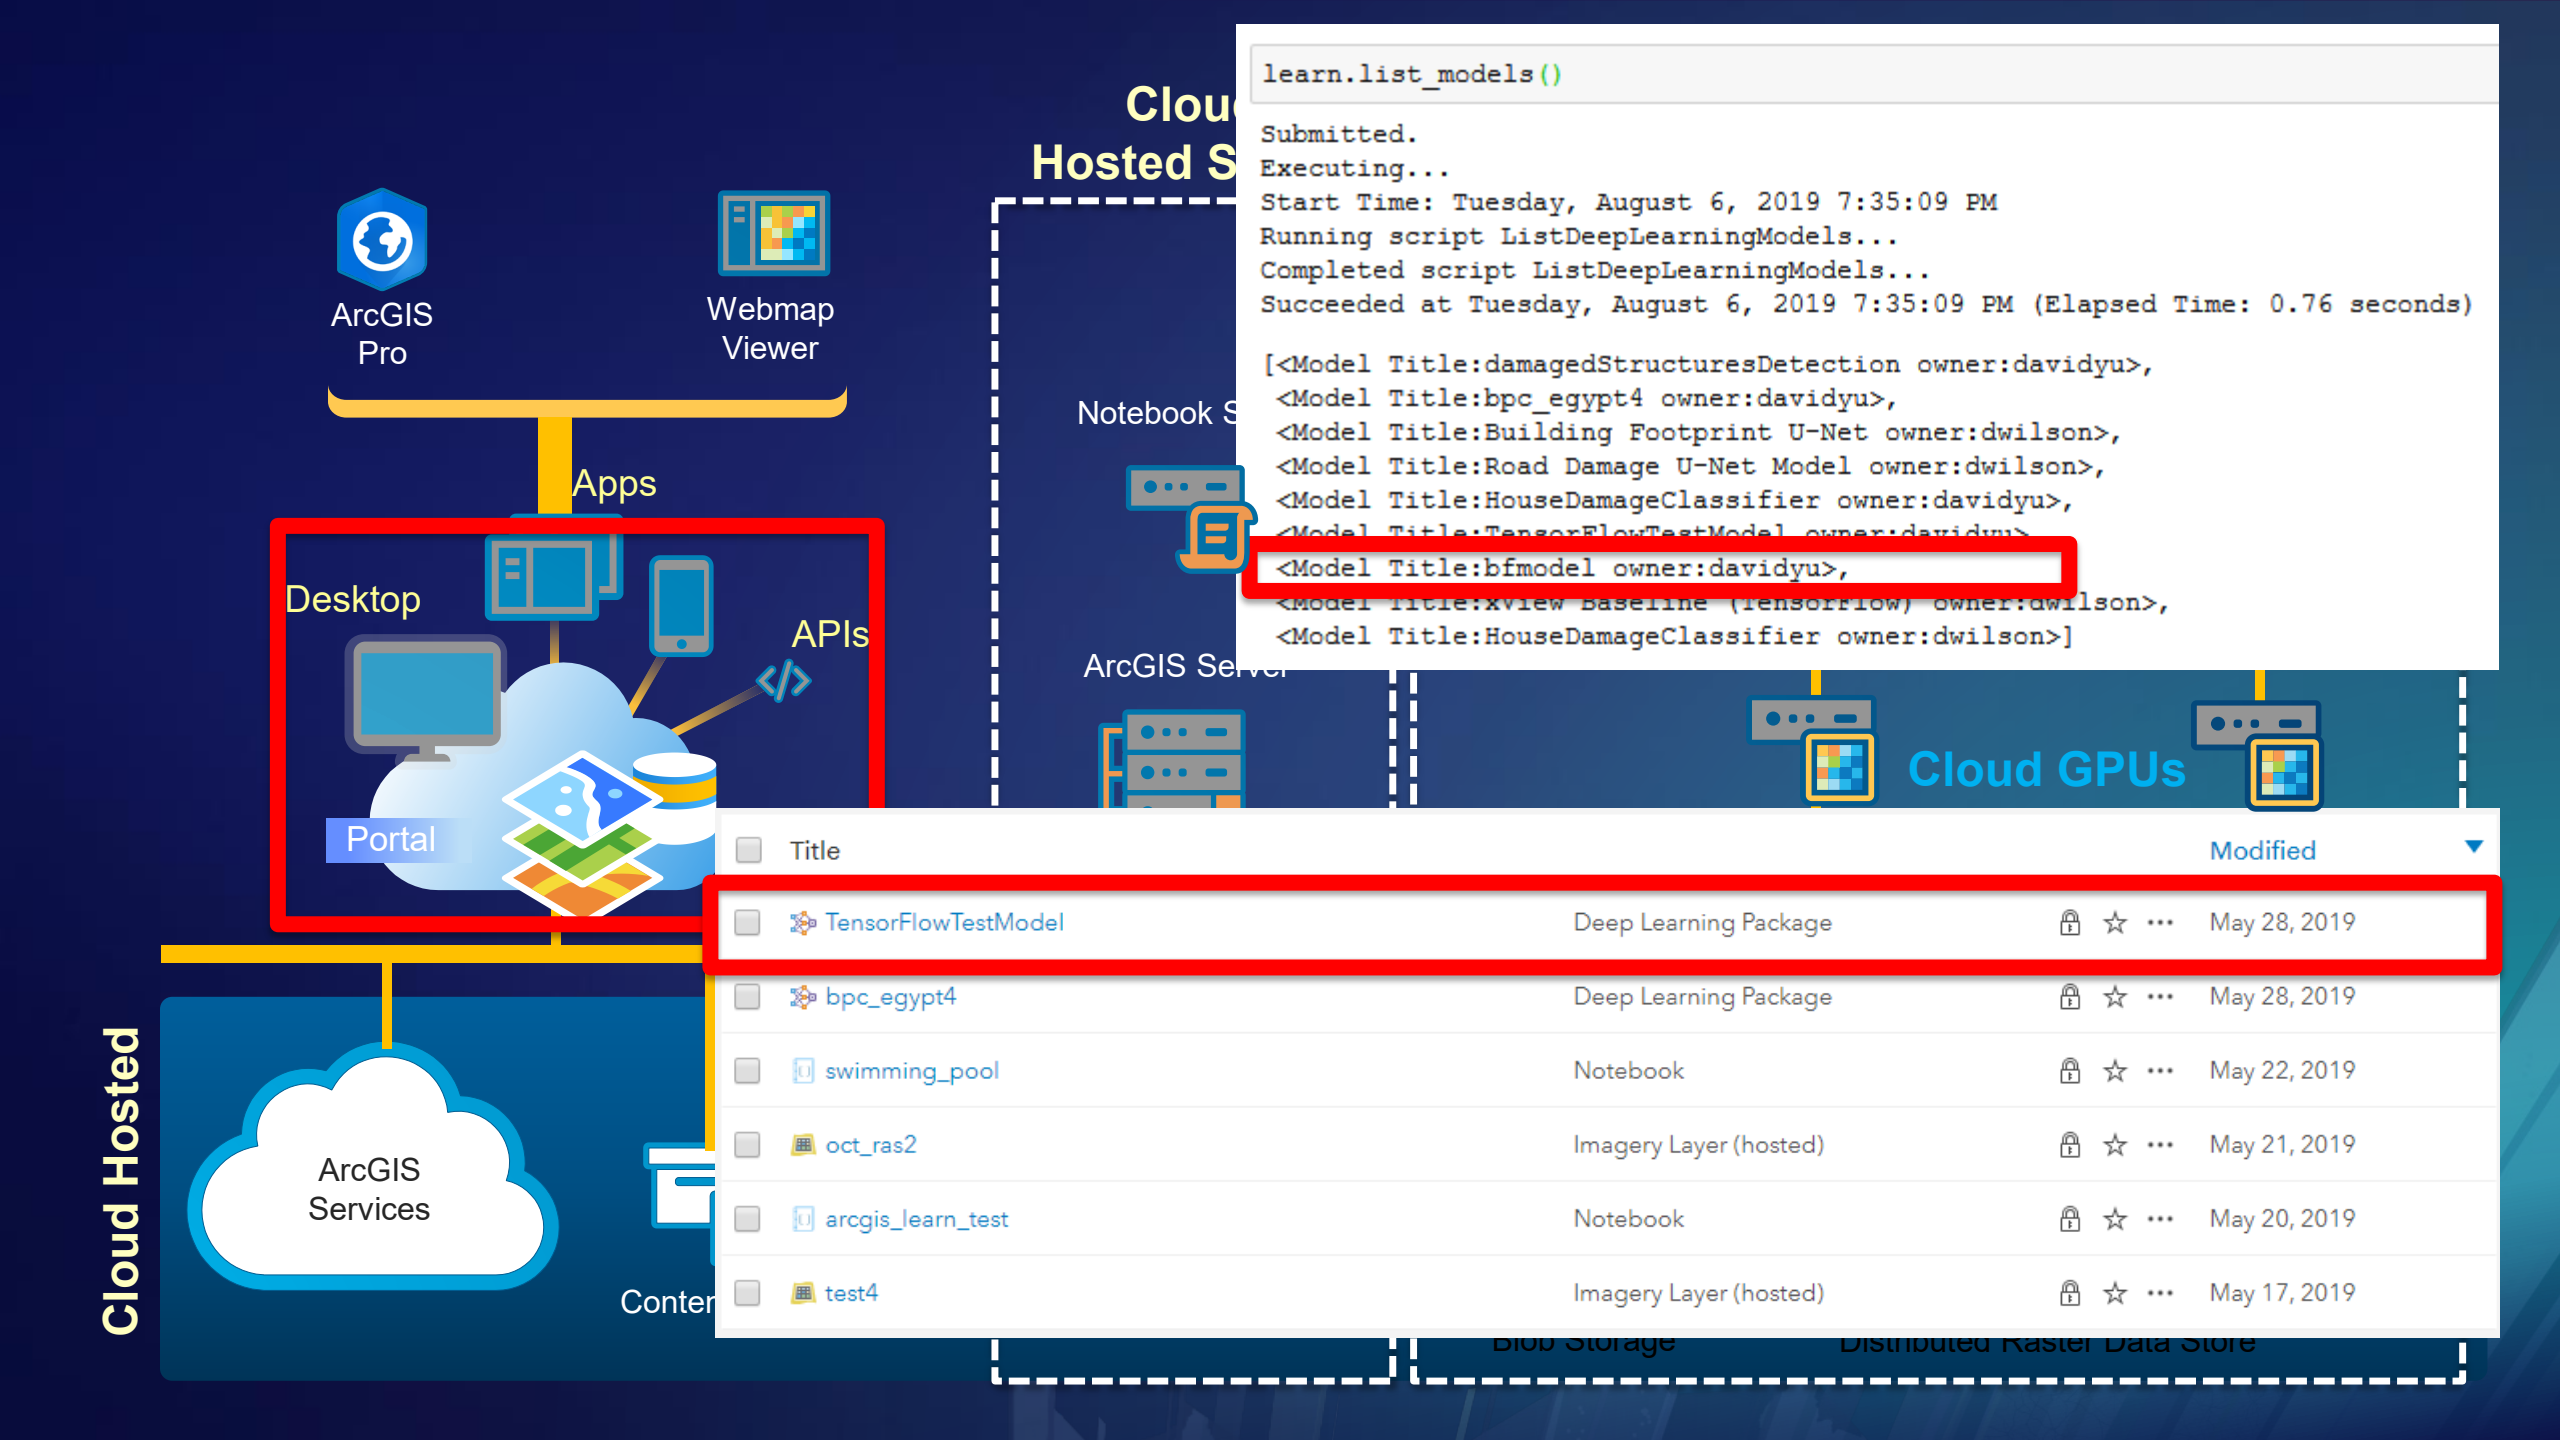
learn.list_models()
Submitted.
Executing...
Start Time: Tuesday, August 6, 2019 7:35:09 PM
Running script ListDeepLearningModels...
Completed script ListDeepLearningModels...
Succeeded at Tuesday, August 6, 2019 7:35:09 PM (Elapsed Time: 0.76 seconds)
[<Model Title:damagedStructuresDetection owner:davidyu>,
<Model Title:bpc_egypt4 owner:davidyu>,
<Model Title:Building Footprint U-Net owner:dwilson>,
<Model Title:Road Damage U-Net Model owner:dwilson>,
<Model Title:HouseDamageClassifier owner:davidyu>,
<Model Title:bfmodel owner:davidyu>,
<Model Title:HouseDamageClassifier owner:dwilson>]
ArcGIS Pro
Webmap Viewer
Apps
Desktop
APIs
Portal
Notebook S
ArcGIS Server
Cloud GPUs
Cloud Hosted
ArcGIS Services
Title
Modified
TensorFlowTestModel
Deep Learning Package
May 28, 2019
bpc_egypt4
Deep Learning Package
May 28, 2019
swimming_pool
Notebook
May 22, 2019
oct_ras2
Imagery Layer (hosted)
May 21, 2019
arcgis_learn_test
Notebook
May 20, 2019
test4
Imagery Layer (hosted)
May 17, 2019

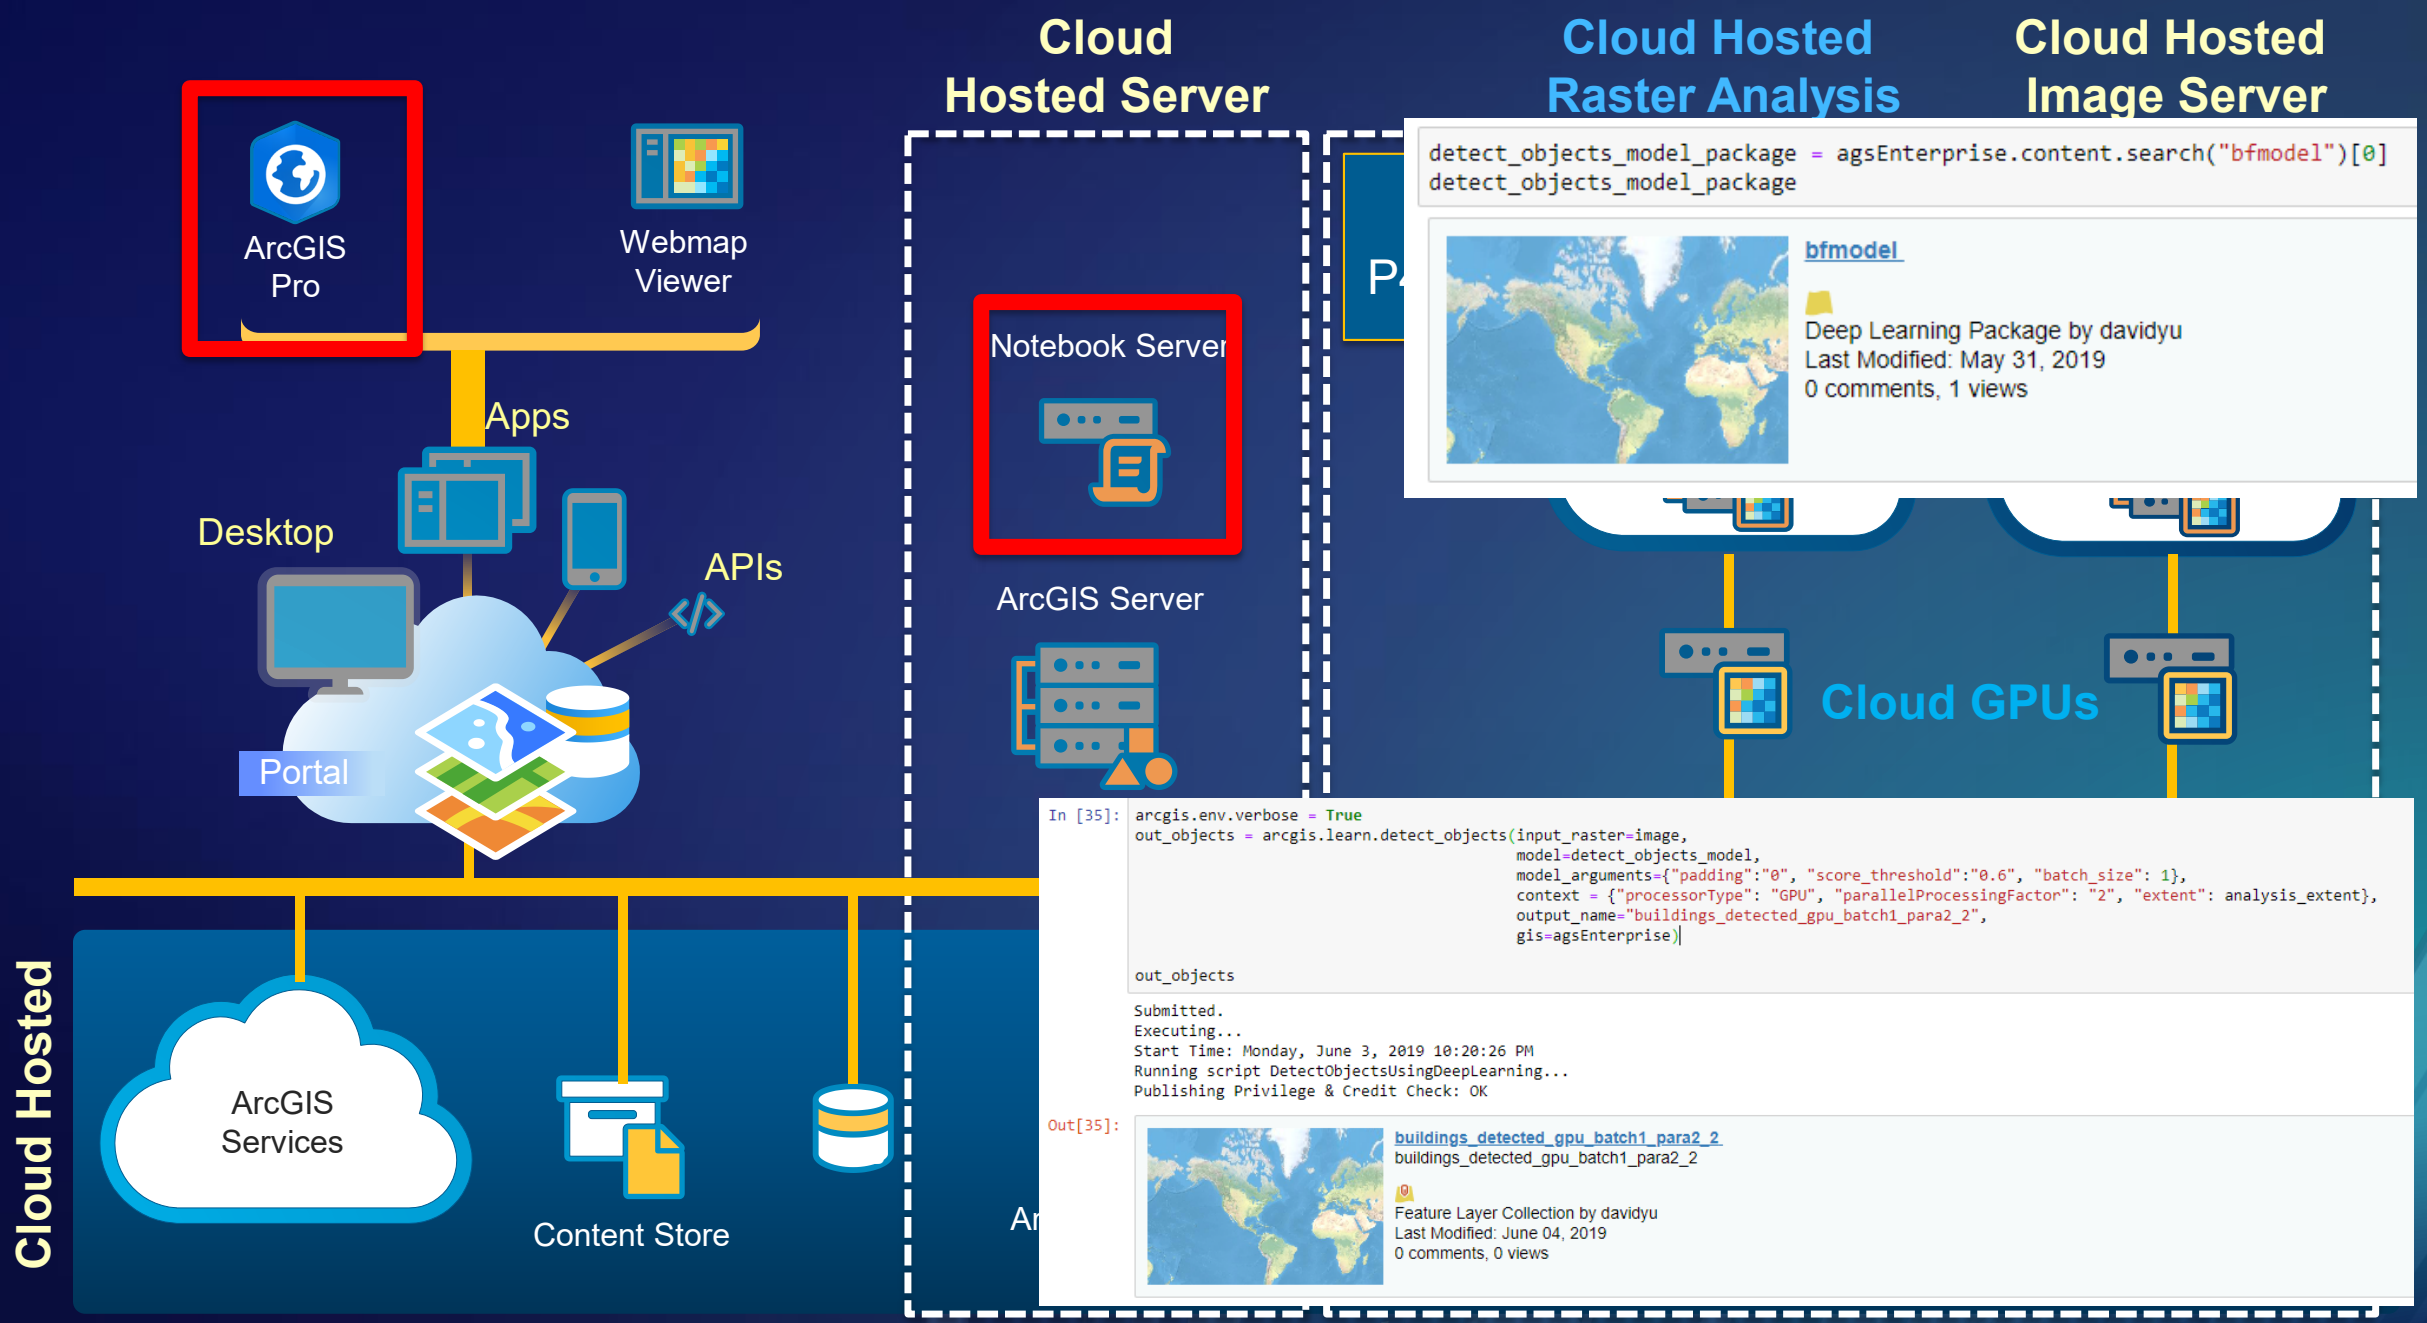
Cloud Hosted Server
Cloud Hosted Raster Analysis
Cloud Hosted Image Server
ArcGIS Pro
Webmap Viewer
Apps
Desktop
APIs
Portal
Notebook Server
ArcGIS Server
Cloud GPUs
Cloud Hosted
ArcGIS Services
Content Store
detect_objects_model_package = agsEnterprise.content.search("bfmodel")[0]
detect_objects_model_package
bfmodel
Deep Learning Package by davidyu
Last Modified: May 31, 2019
0 comments, 1 views
In [35]:
arcgis.env.verbose = True
out_objects = arcgis.learn.detect_objects(input_raster=image,
model=detect_objects_model,
model_arguments={"padding":"0", "score_threshold":"0.6", "batch_size": 1},
context = {"processorType": "GPU", "parallelProcessingFactor": "2", "extent": analysis_extent},
output_name="buildings_detected_gpu_batch1_para2_2",
gis=agsEnterprise)
out_objects
Submitted.
Executing...
Start Time: Monday, June 3, 2019 10:20:26 PM
Running script DetectObjectsUsingDeepLearning...
Publishing Privilege & Credit Check: OK
Out[35]:
buildings_detected_gpu_batch1_para2_2
buildings_detected_gpu_batch1_para2_2
Feature Layer Collection by davidyu
Last Modified: June 04, 2019
0 comments, 0 views

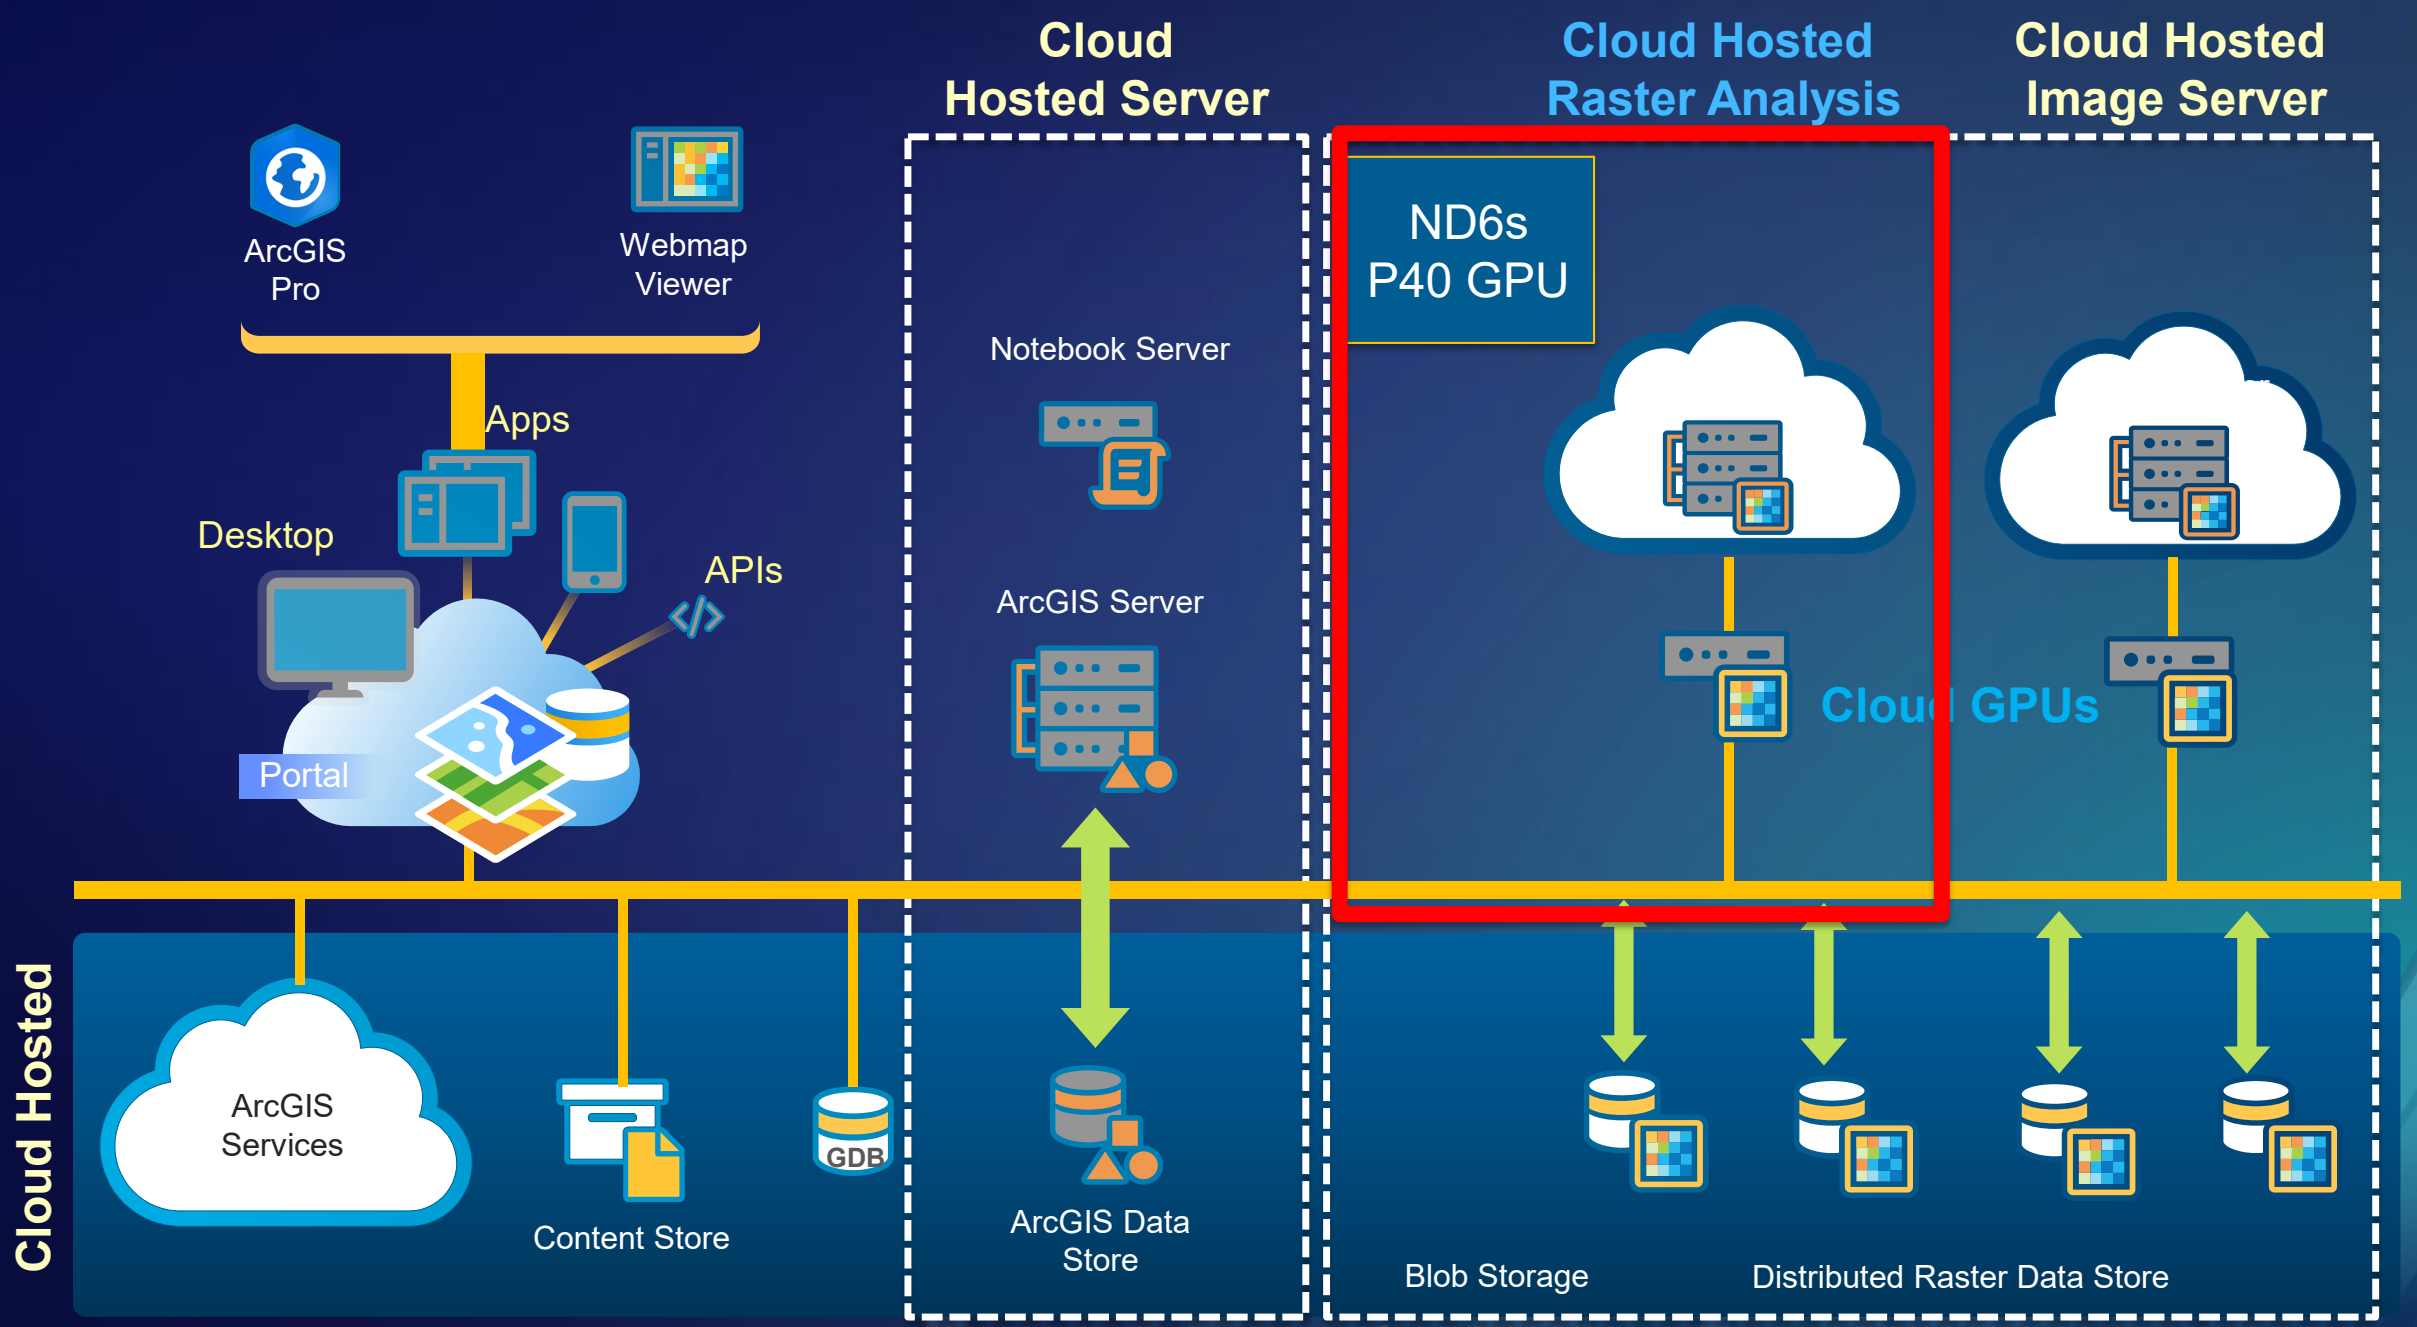
Cloud Hosted Server
Cloud Hosted Raster Analysis
Cloud Hosted Image Server
ArcGIS Pro
Webmap Viewer
Apps
Desktop
APIs
Portal
Notebook Server
ArcGIS Server
ND6s P40 GPU
Cloud GPUs
Cloud Hosted
ArcGIS Services
Content Store
GDB
ArcGIS Data Store
Blob Storage
Distributed Raster Data Store

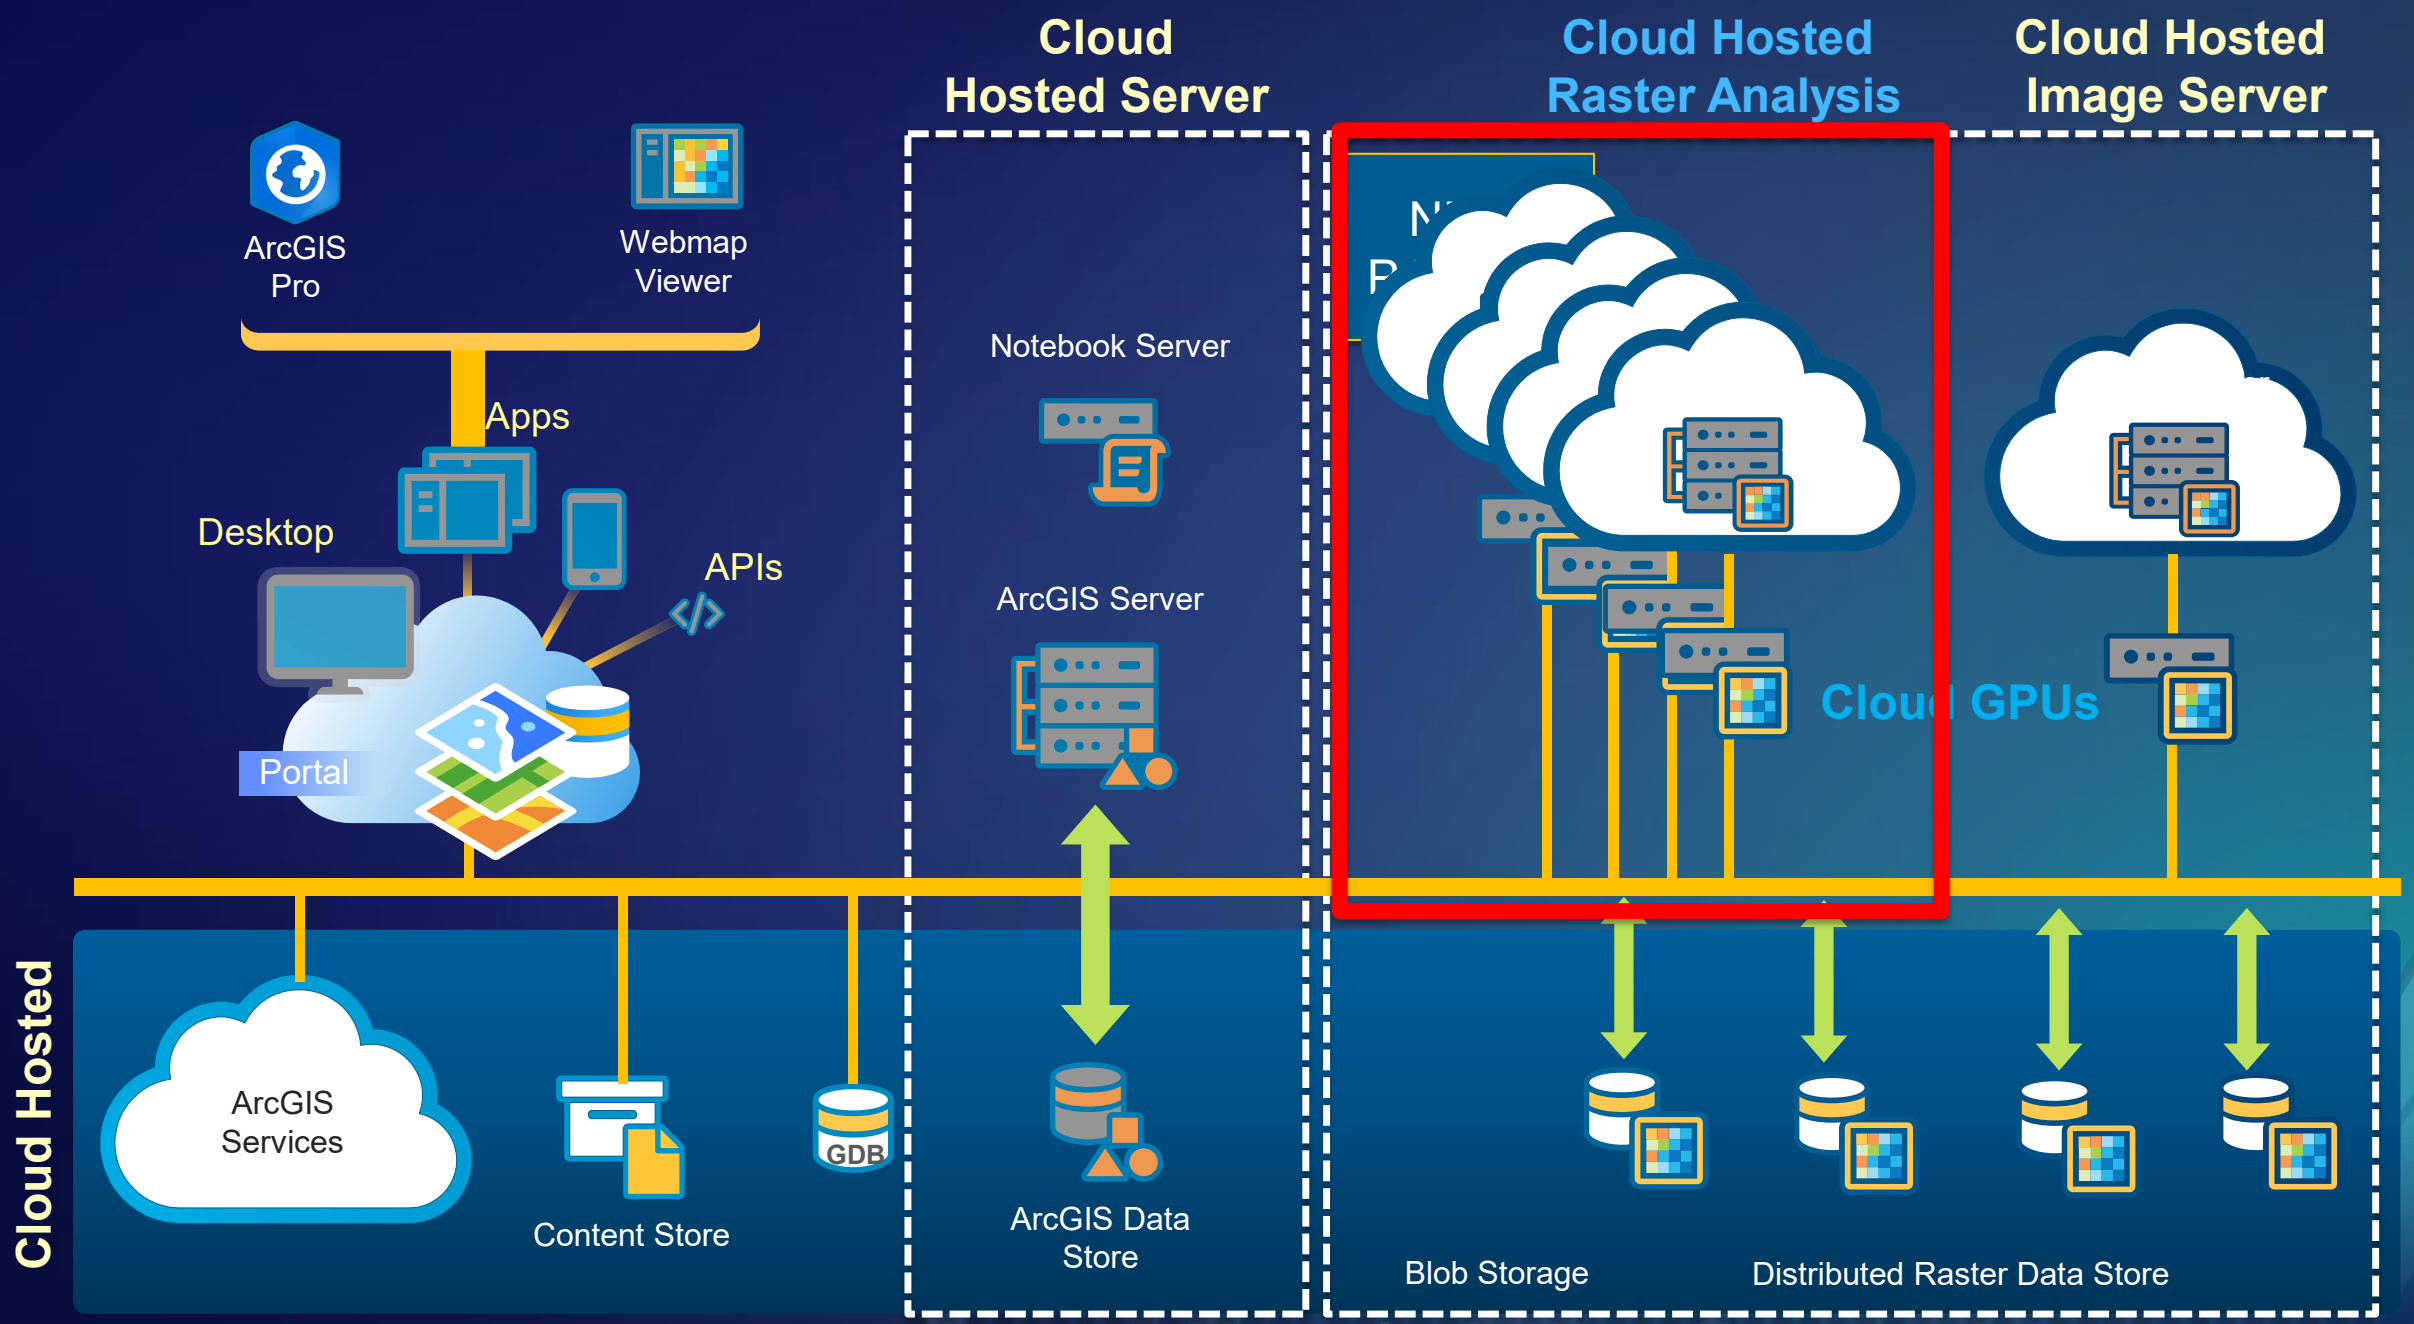
Cloud Hosted Server
Cloud Hosted Raster Analysis
Cloud Hosted Image Server
ArcGIS Pro
Webmap Viewer
Apps
Desktop
APIs
Portal
Notebook Server
ArcGIS Server
Cloud GPUs
Cloud Hosted
ArcGIS Services
Content Store
GDB
ArcGIS Data Store
Blob Storage
Distributed Raster Data Store

Benchmarks

Denver Imagery

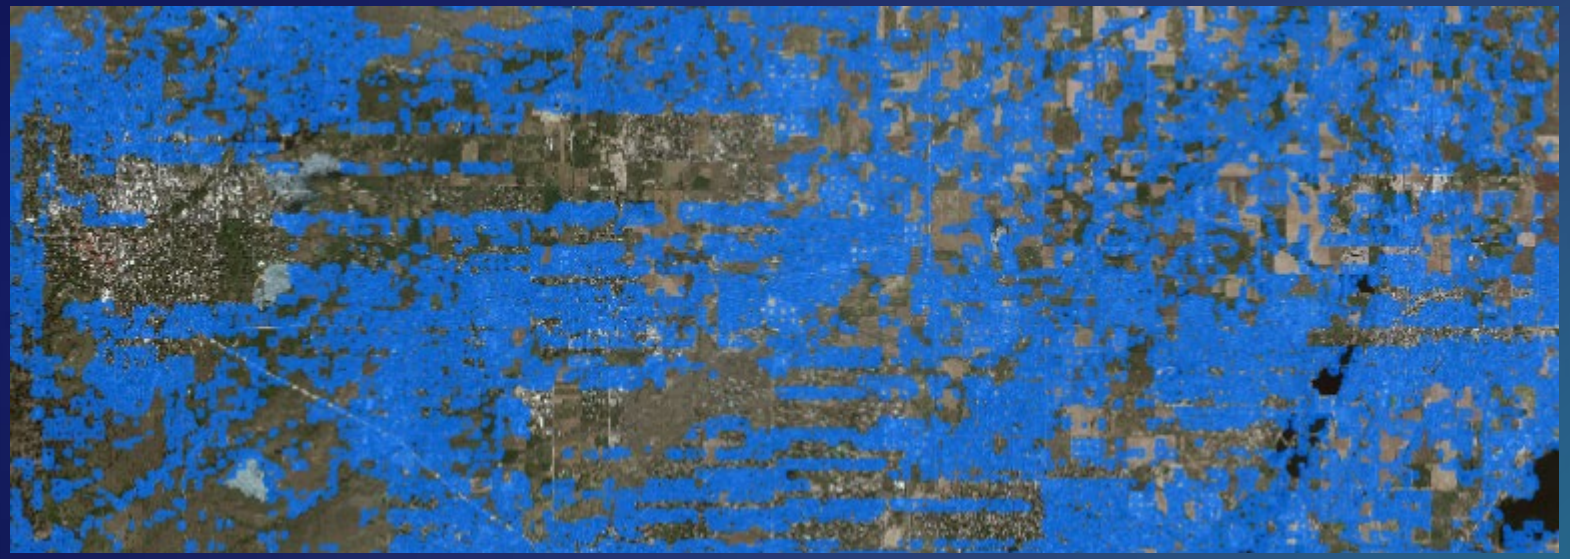

NOAA Imagery

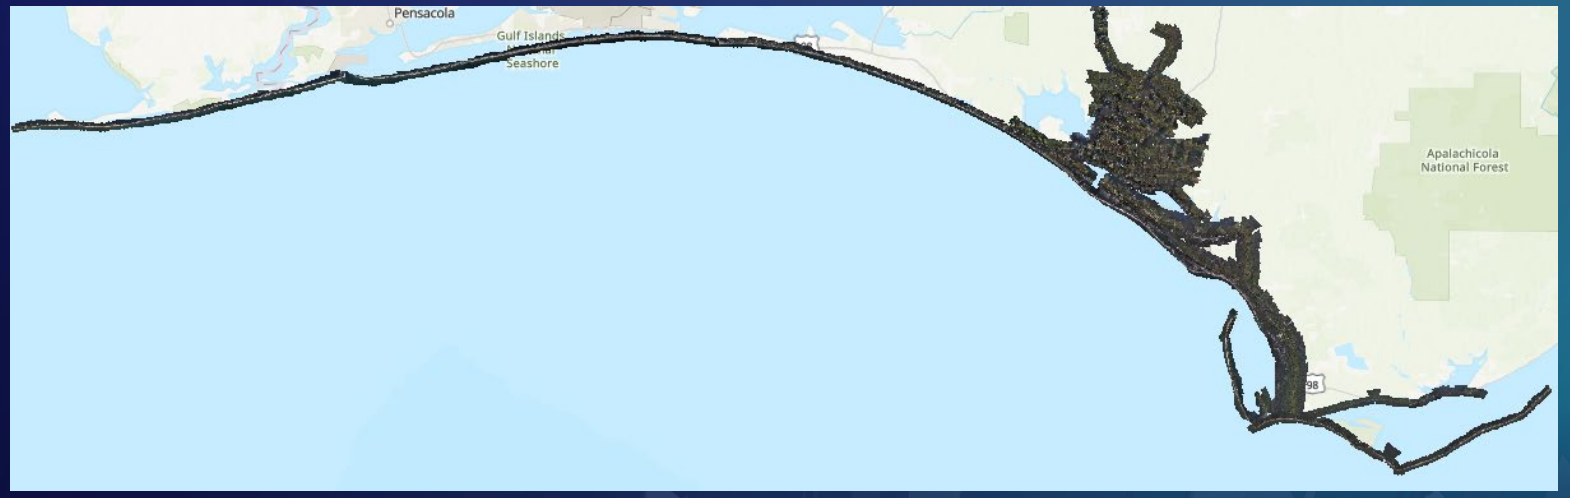

Benchmarks

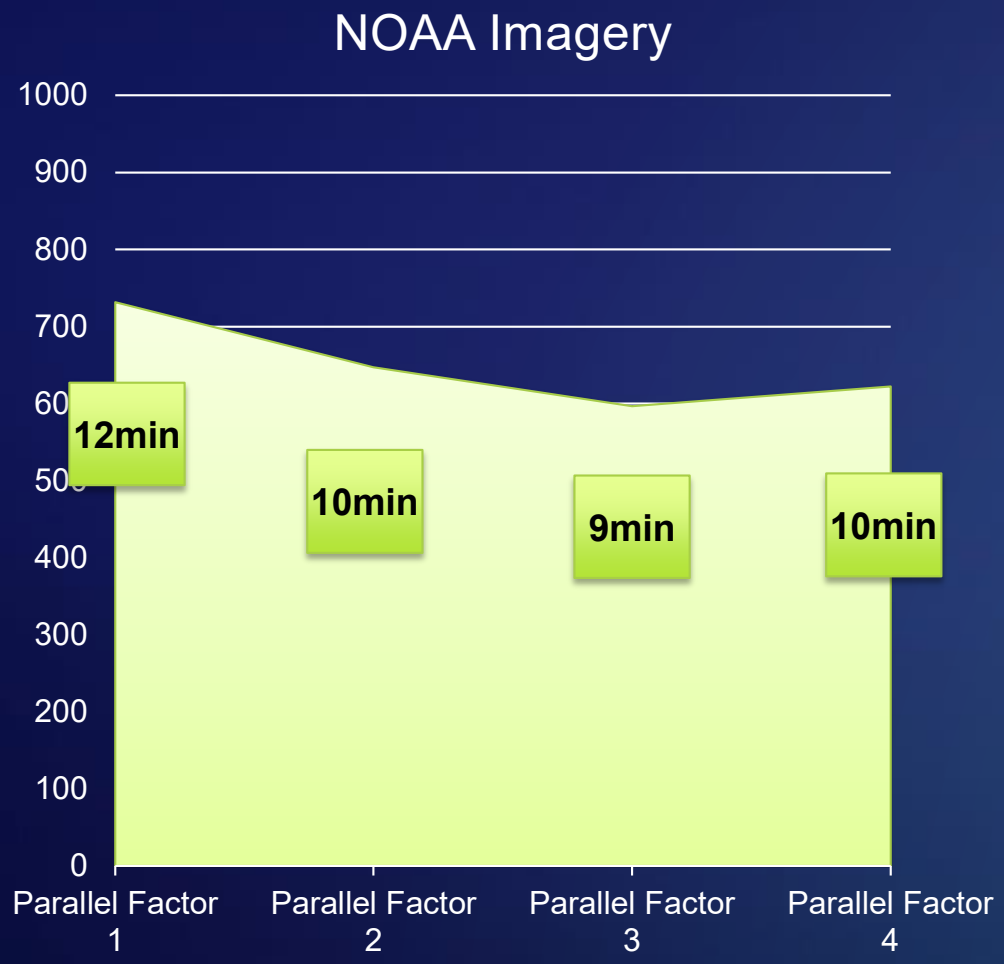
NOAA Imagery
1000
900
800
700
600
500
400
300
200
100
0
12min
10min
9min
10min
Parallel Factor 1
Parallel Factor 2
Parallel Factor 3
Parallel Factor 4

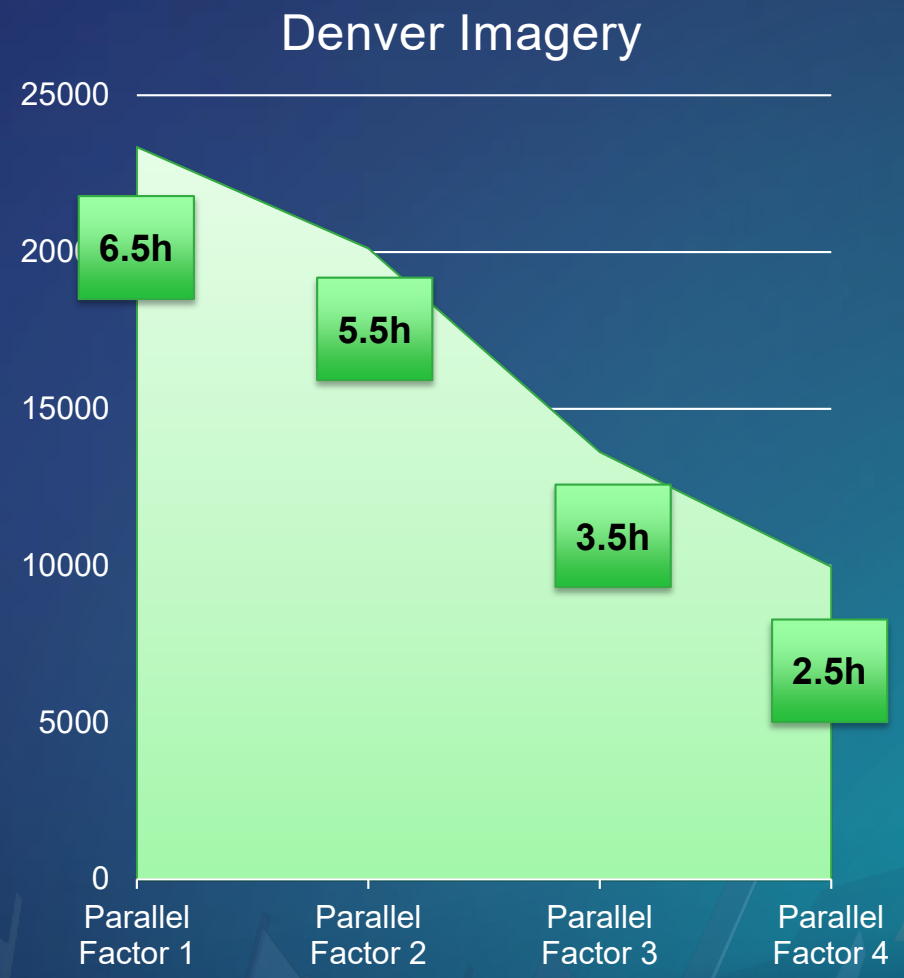
Denver Imagery
25000
20000
15000
10000
5000
0
6.5h
5.5h
3.5h
2.5h
Parallel Factor 1
Parallel Factor 2
Parallel Factor 3
Parallel Factor 4

Recap

- New Geoprocessing tools
  - model training in ArcGIS Pro and Enterprise
  - Inferencing in ArcGIS Pro and Enterprise
- arcgis.learn module enhancements
- Sample Notebooks and Learn lessons for deep learning workflows
- New model types (PSPNet, RetinaNet, MaskRCNN…)
- Multi GPU inference using enterprise

Resources

- Geonet Community https://community.esri.com/groups/arcgis-python-api/pages/overview
- ArcGIS API for Python https://developers.arcgis.com/python/
- API Reference https://developers.arcgis.com/python/api-reference/
- GitHub Repo https://github.com/Esri/arcgis-python-api
- Sample Notebooks https://developers.arcgis.com/python/sample-notebooks/
- GeoAI blogs https://medium.com/geoai


esri
THE
SCIENCE
OF
WHERE