Geo-Localization and Risk Estimation using Deep Learning: Utility poles

Student: Jamie Voynow, University of Vermont
jvoynow@uvm.edu, (856) 304-4930

Mentor: Dr. Safwan Wshah, University of Vermont
swshah@uvm.edu

NORTHEAST CYBERTEAM

Background

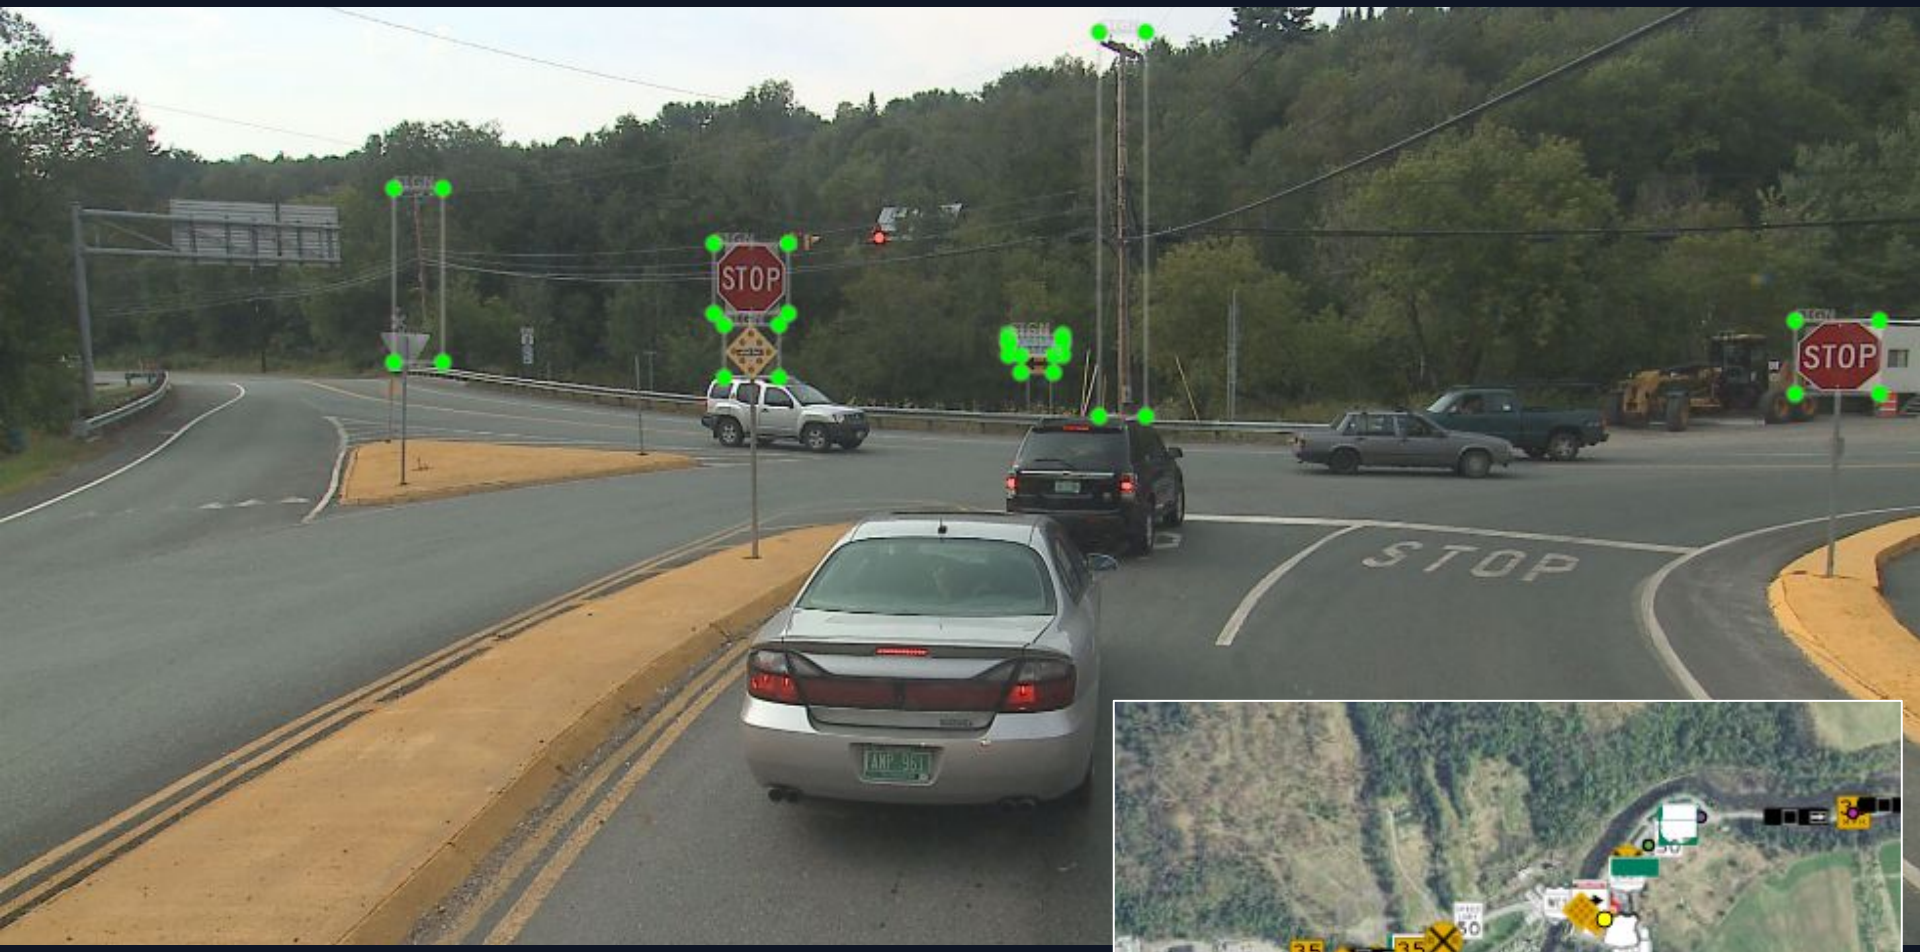

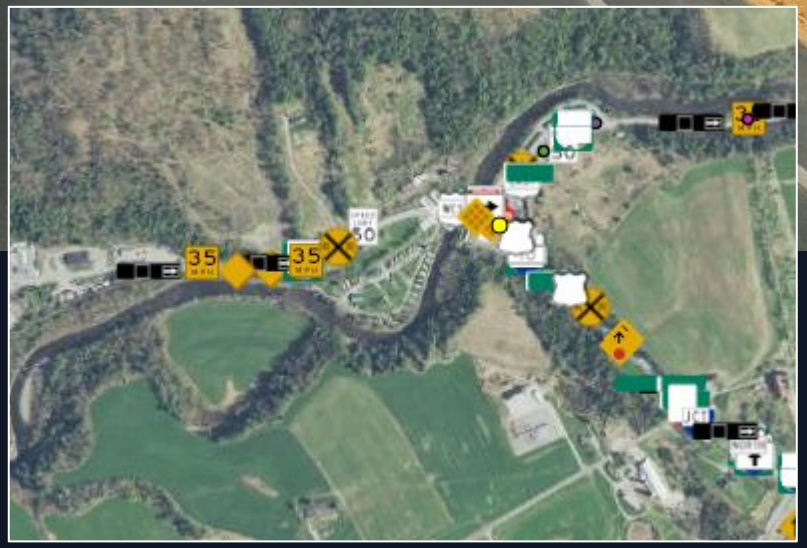

- The Vermont AI Lab (VaiL) and VTrans collaboration
- Previous Projects (Sign Detection, Lane Line QA)
- Next asset class: Utility Poles
- The connection between this project and past research

Motivation

Poles
Utility poles exist on roads worldwide
Provide services to humans everywhere
Tens of millions of utility poles exist in the united states

Potential for damage
Wooden structures tend to have issues in nature
Susceptible to damage from weather, insects/animals.

The burden
When a pole is damaged (and collapses), the burden is felt by the community, the local government, and the utility companies.

Role of Deep Learning
How can we protect companies and communities?
What is needed in order to achieve these goals?

NORTHEAST CYBERTEAM

System Overview

Database

Sampling Data

Neural Network

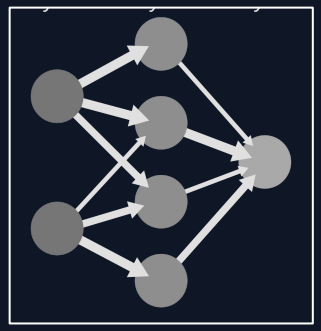

Architecture from "Focal loss for dense object detection" https://arxiv.org/pdf/1708.02002.pdf

Output for each image

- Bounding Box 1 + GPS + Quality Estimation
- Bounding Box 2 + GPS + Quality Estimation
- Bounding Box n + GPS + Quality Estimation

Item(s) with quality estimate

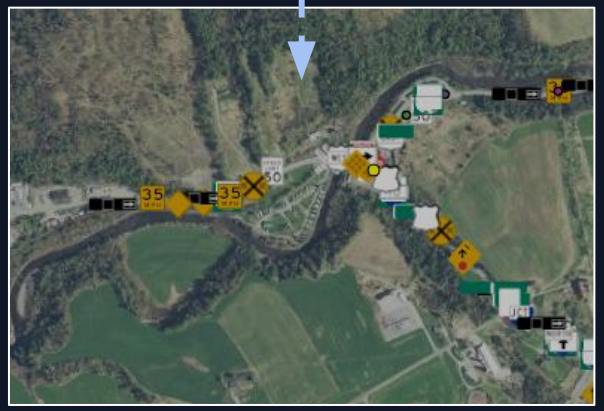

Inside Database

Labeling Images

Back into Database

NORTHEAST CYBERTEAM

Developments to date (3/17)

- Dataset
- Labeling Tool
- Labeled Data
- Data Extraction Pipeline

NORTHEAST CYBERTEAM

Developments to date (4/28)

- Dataset
- Labeling Tool
- Labeled Data
- Data Extraction Pipeline
  - Supported by DeepGreen
- Baseline Model
  - Supported by DeepGreen

NORTHEAST CYBERTEAM

Developments to date (6/16)

- Dataset
- Labeling Tool
- Labeled Data
- Data Extraction Pipeline
  - Supported by DeepGreen
- Baseline Model
  - Supported by DeepGreen
- Data Analysis Tool
- Final Model

Increases in:
- labeled images from 470 to 901
- labeled objects from 445 to 1382

NORTHEAST CYBERTEAM

Developments to date (7/20)

- Dataset
- Labeling Tool
- Labeled Data
- Data Extraction Pipeline
  - Supported by DeepGreen
- Baseline Model
  - Supported by DeepGreen
- Data Analysis Tool
- Final Model

Increases in:
- labeled images from 901 to 1083
- labeled objects from 1382 to 1754

Note about last month of research:
Many hours were spent debugging various functionality of the final model. Because the software system that the model relies on has been deprecated for over a year, major challenges were faced that were not anticipated. In hindsight, a solution could have been achieved using a different model, with up to date code/support.

NORTHEAST CYBERTEAM

	im_num	obj_num	lat	long	xmin	ymin	xmax	ymax	im_path
1492	827	1192	43.604302	-73.003041	367	116	466	578	2018/B004_E/20180616.164219/000005102079.jpg
1493	827	1193	43.604104	-73.003106	1372	94	1495	608	2018/B004_E/20180616.164219/000005102079.jpg
1494	827	1194	43.604302	-73.002604	678	343	752	552	2018/B004_E/20180616.164219/000005102079.jpg


R
R
EAST
Rutland

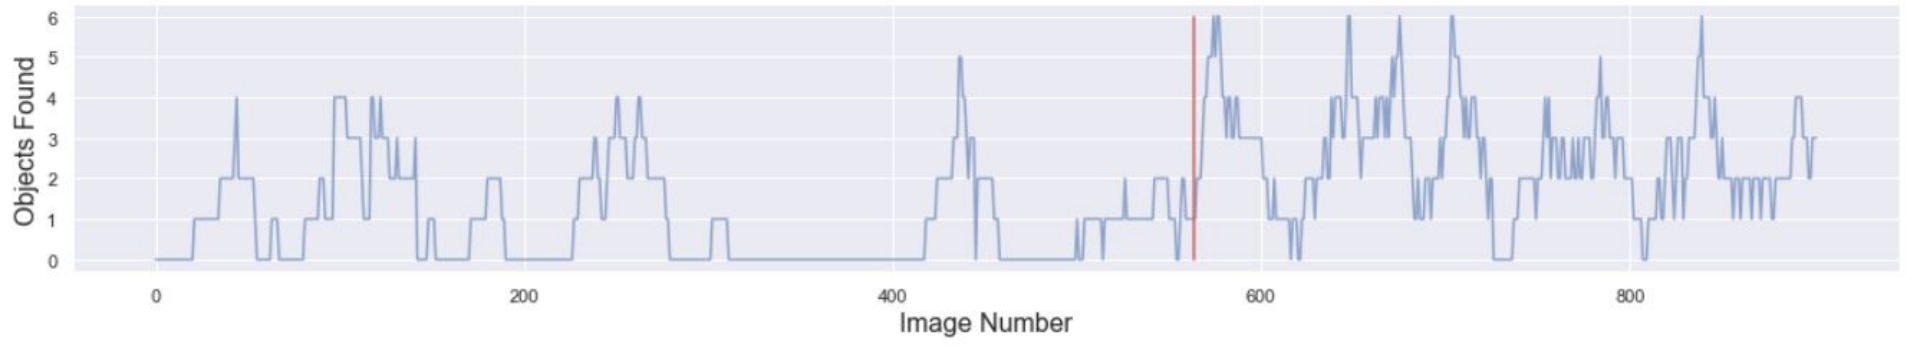
Objects Found
0
1
2
3
4
5
6
0
200
400
600
800
Image Number

NORTHEAST CYBERTEAM

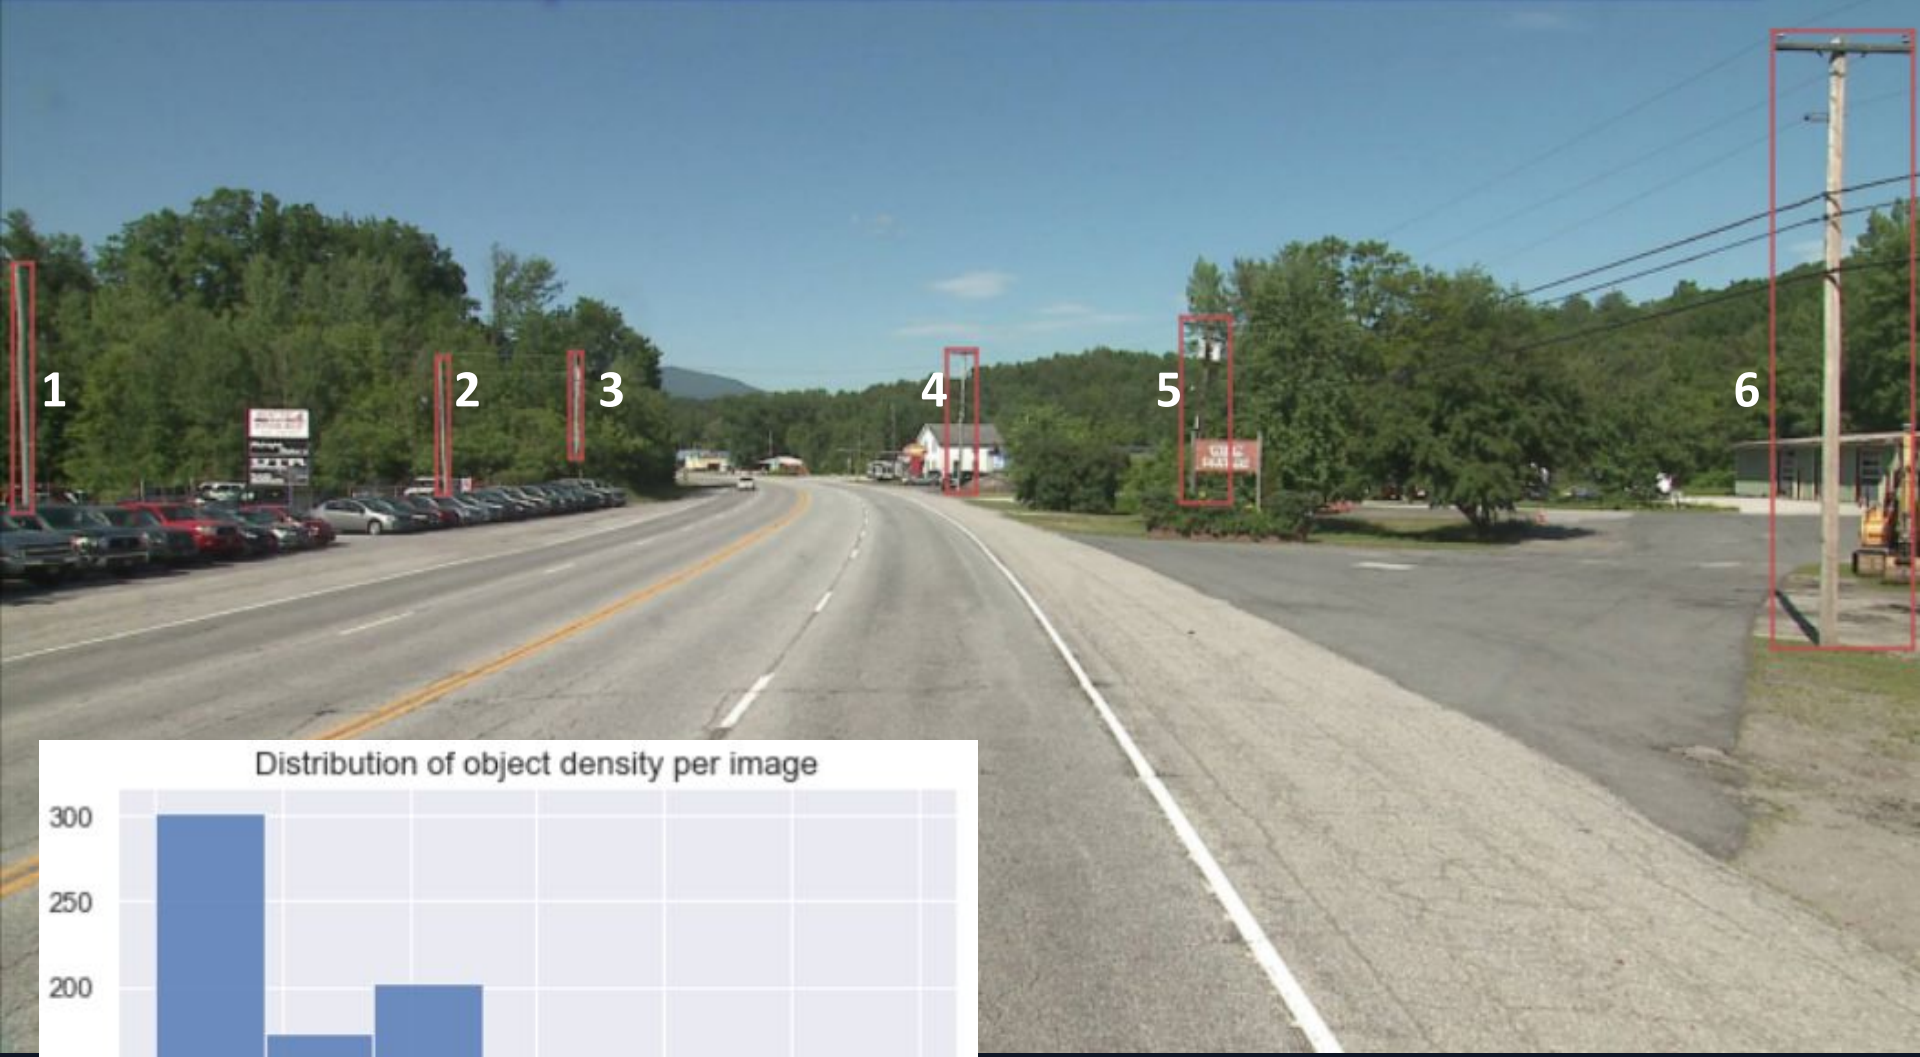

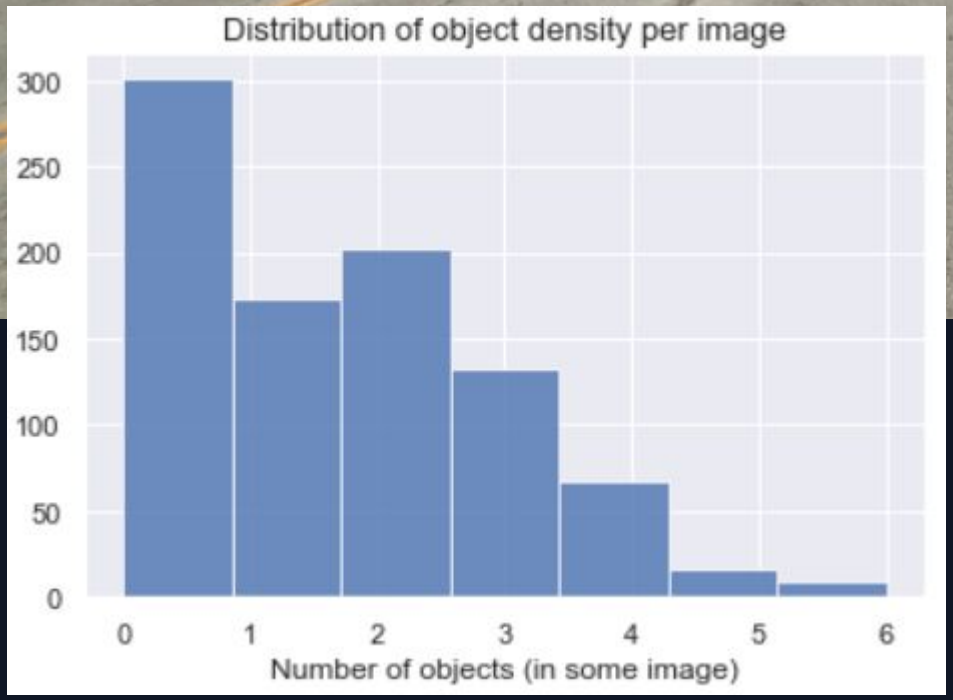

NORTHEAST CYBERTEAM

Final Model Iteration

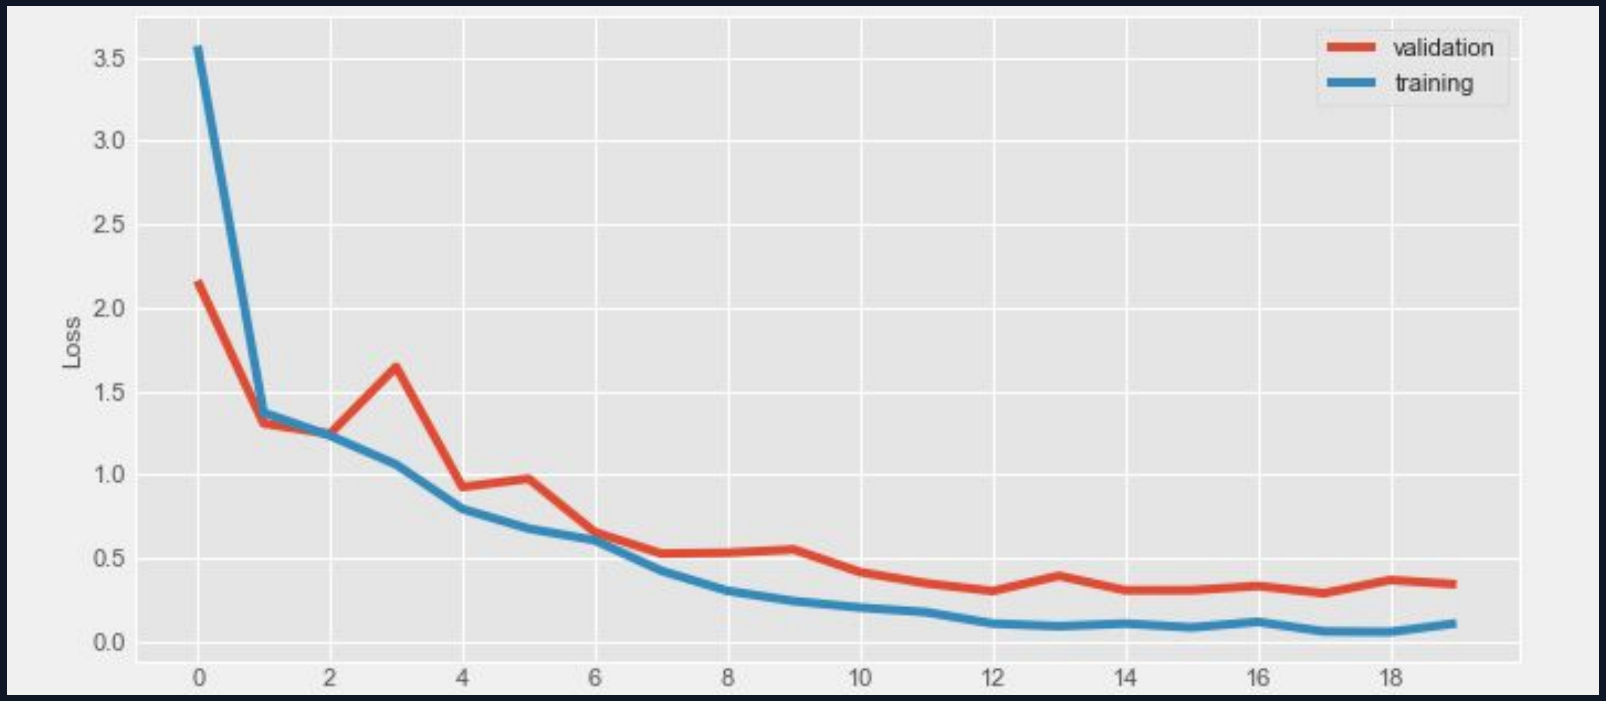

Interpretation:

- Predicting number of utility pole objects in a given image
- Achieving validation loss of <.5 MSE
- Trained on Deep Green for 20 iterations

NORTHEAST CYBERTEAM

Model Output Holdout Data (example 1)

Here we have a test image containing only one utility pole. Our model correctly identified this attribute of this image.

Model Output Holdout Data (example 2)

The model learned to only count objects sufficiently close to the camera. As seen in this image, there are as many as 5 poles that exist in the distance. Our algorithm learned to find only the closest and clearest objects, in this case: we predicted two which was correct.

Model Output Holdout Data (example 3)

Here our model incorrectly counts a distant third pole, that was deemed too far and too blurry to be labeled. Future iterations could consider some criteria for including/excluding objects based on distance or blurriness.

What Have I Learned?

Python Specific Topics:

- Using packages (os, sys) to interact with the operating/file system
  - I now am able to utilize these packages in my full time job
- Using packages to interact with image data
  - The pillow package makes it easy to work with .jpg files
- Using packages to interact with data structured hierarchically (.xml files)
- Preparing python scripts to be executed on distributed computing systems

What Have I Learned?

Software Development Specific Topics:

- Using version control
  - Pulling massive repositories into local memory
- Deploying code to GPU/Deep Green
  - Working in linux, moving files on and off of external hardware
- Working with open-source software
  - Using the retinanet and vtrans_ai software systems
- Building on top of massive software systems
- Debugging package level errors
  - Getting down to the fine details package installation by version
  - Searching deep in package code for debugging purposes
- Package management with anaconda environments

What Have I Learned?

Bigger Picture Topics:

- Adhering to schedule is crucial for completing large projects
  - I fell behind my schedule early on
  - My biggest challenges did not become apparent until the end of my project
- There are multiple ways to solve every problem
  - Don't continue on a dead end just because you think you are too invested
- Hand labeling data is a time consuming task
- Asking for help is important, especially when you are surrounded by experts
- Working with open source software can challenging
  - But as detailed above, is a great learning experience!

NORTHEAST CYBERTEAM

Geo-Localization and Risk Estimation using Deep Learning: Utility poles

Thank you for this opportunity!