Deep Learning in 3D

Object extraction from point clouds

AMLD 2020

Dmitry Kudinov, dkudinov@esri.com

Sr. Principal Data Scientist, Esri

3D Models of Cities: Valuable and Expensive

The Digital Twin story

Lower the cost of GIS-grade content acquisition by providing workflows for [semi-]automated objects extraction from point clouds, imagery, meshes.

What is LiDAR data?

Massive 3D point collections with additional attributes like:

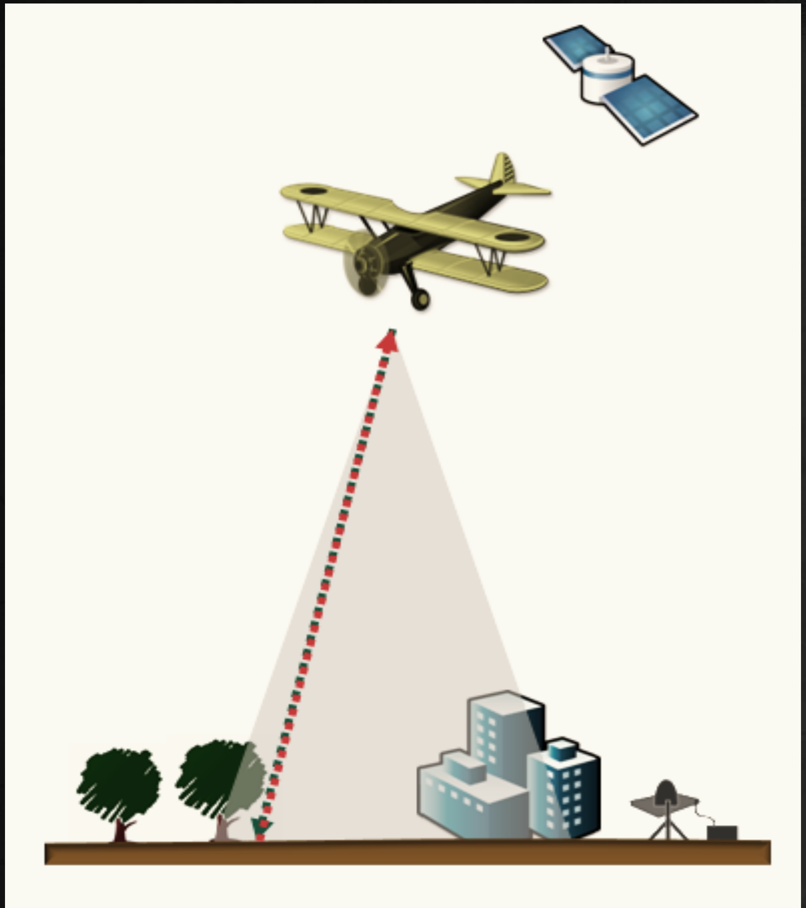

- Intensity,
- Number of Returns,
- Scan Angle,
- RGB,…

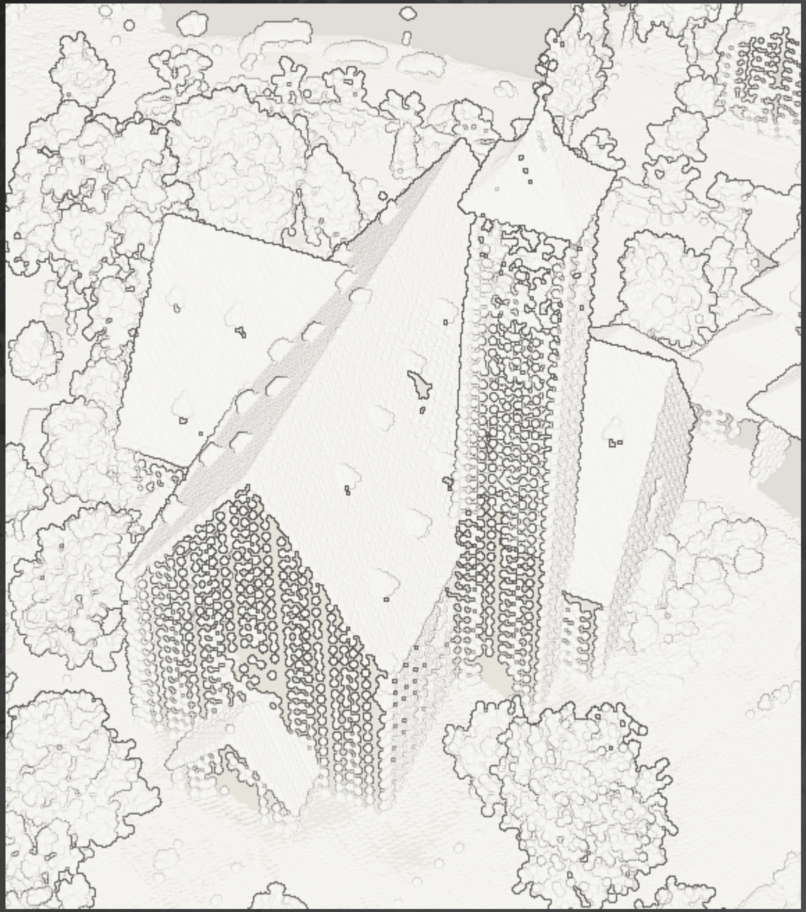

What is LiDAR data?

Massive 3D point collections with additional attributes like:

- Intensity,
- Number of Returns,
- Scan Angle,
- RGB,…

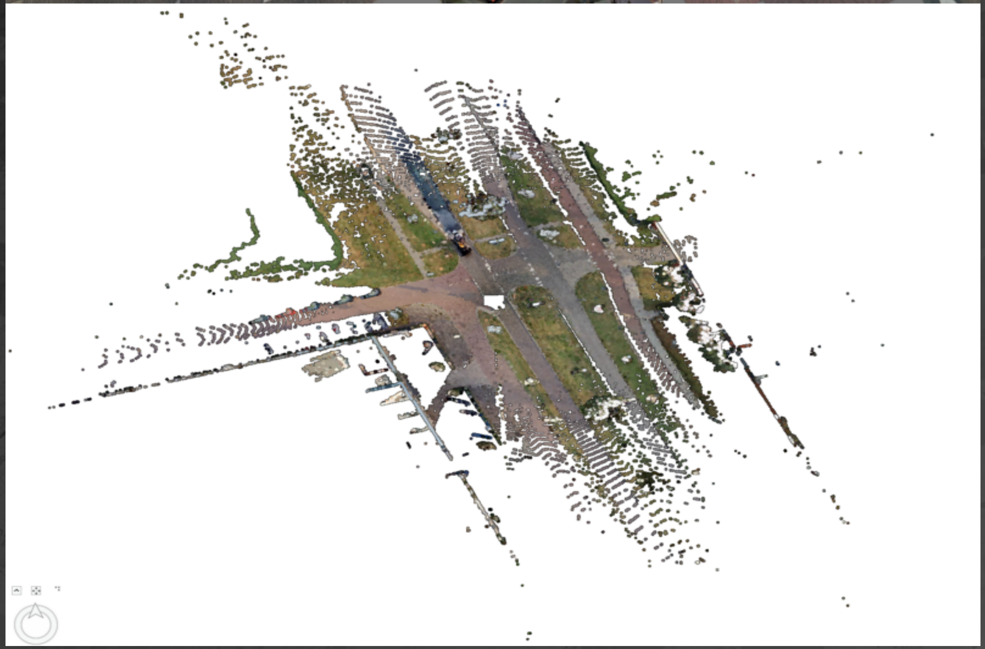

What can be extracted from point clouds?

- 3D building models
  - schematic, or
  - detailed (up to RGB+D textures)
- Electric grid
  - overhead conductors, poles, stay wires, transformers, etc.
- Trees
  - canopy polygons, precise height
- Street furniture
  - Signs, traffic lights, dividers, bus stops, fire hydrants, etc.

3D Building models

from airborne LiDAR

Building models: Realism vs Cubism

1. High fidelity models of historical buildings and cityscape features which are considered stable and never / rarely undergo any modifications.
   - Manually crafted models,
   - Often have designated budgets for creation,
   - Rarely updated.

2. Schematic-like models of commercial, industrial, residential zones which develop and change often.
   - Have the largest area,
   - Need to be re-evaluated periodically for taxation and regulatory purposes,
   - Must be evaluated first and fast in case of a natural disaster, e.g. earthquake,
   - The process must be quick, accurate enough, and cost effective.

PoC: 3D Building model reconstruction from aerial LiDAR

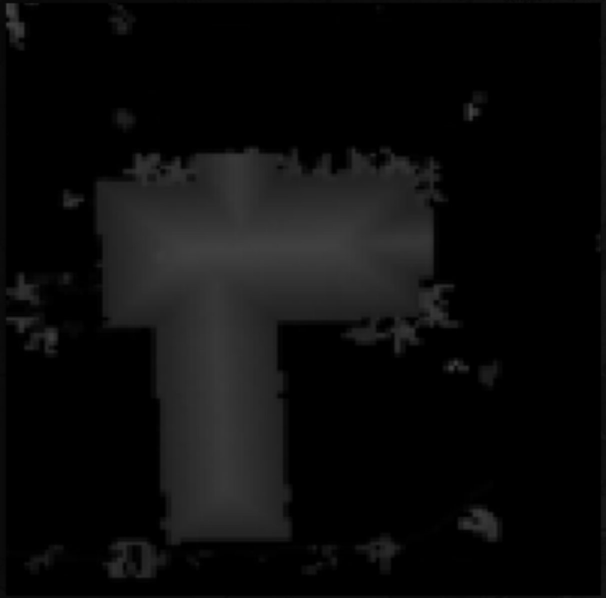

Rasterized Aerial LiDAR

Manually digitized Hip (purple) and Gable (orange) segments

3D reconstruction of building using manually digitized segments

PoC: 3D Building model reconstruction from aerial LiDAR

- Manually digitizing roof segments:
  - Over 3,000 man hours were spent on digitizing about 213,000 polygons covering the area of 200 square miles.
  - ~70 polygons / man hour.

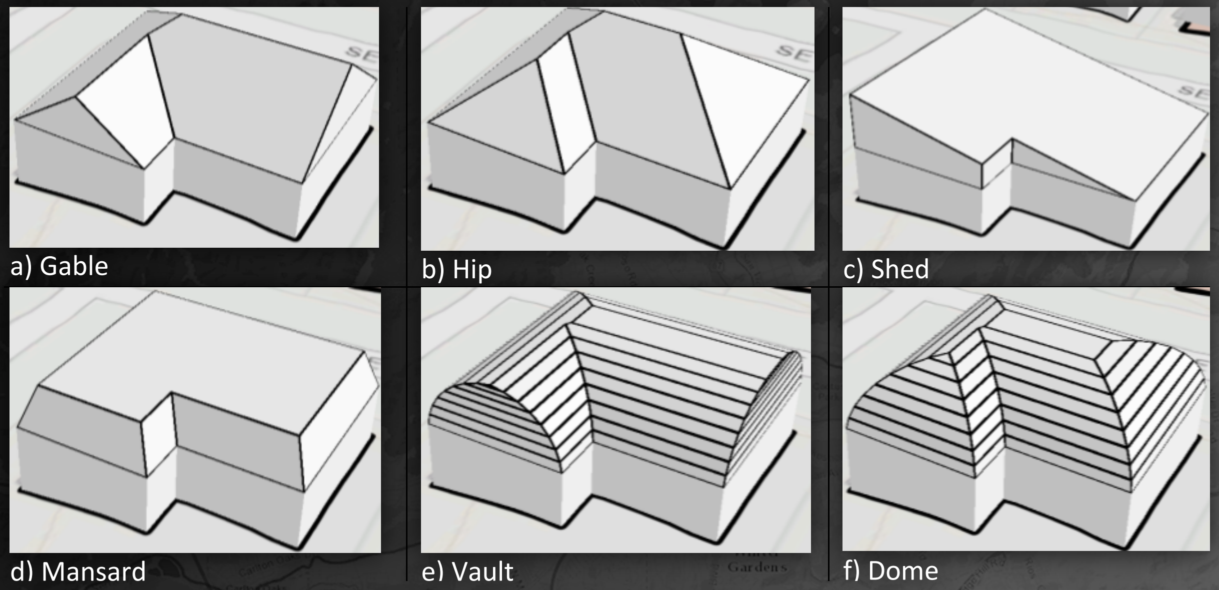

PoC: 3D Building model reconstruction from aerial LiDAR

- Using Mask R-CNN to digitize roof segments
- Not as accurate as humans, but much faster: 60,000 polygons / hour.
- Regularize Building Footprints helps with accuracy.

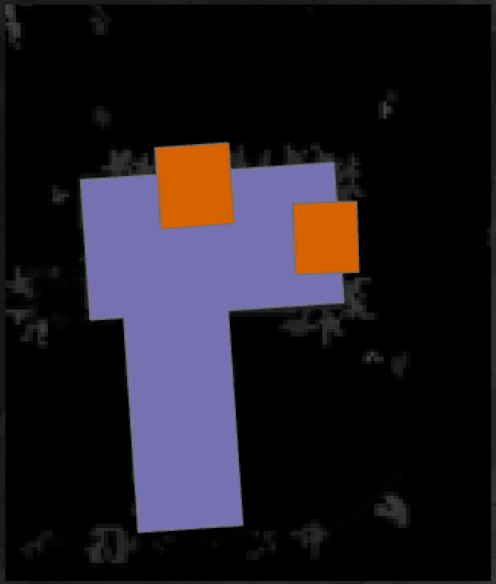

Manually digitized "ground truth" data from the Test set

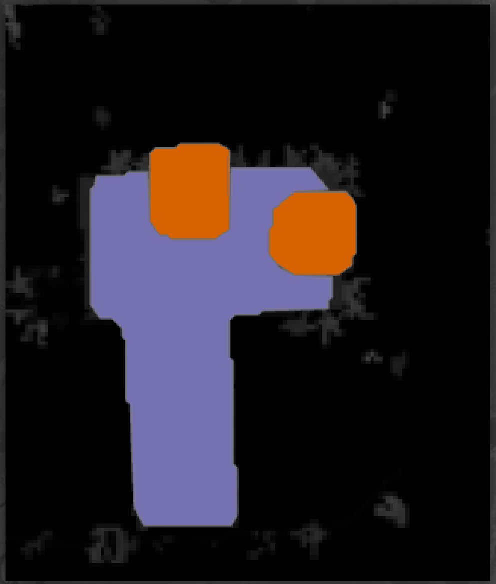

Prediction produced by the neural network

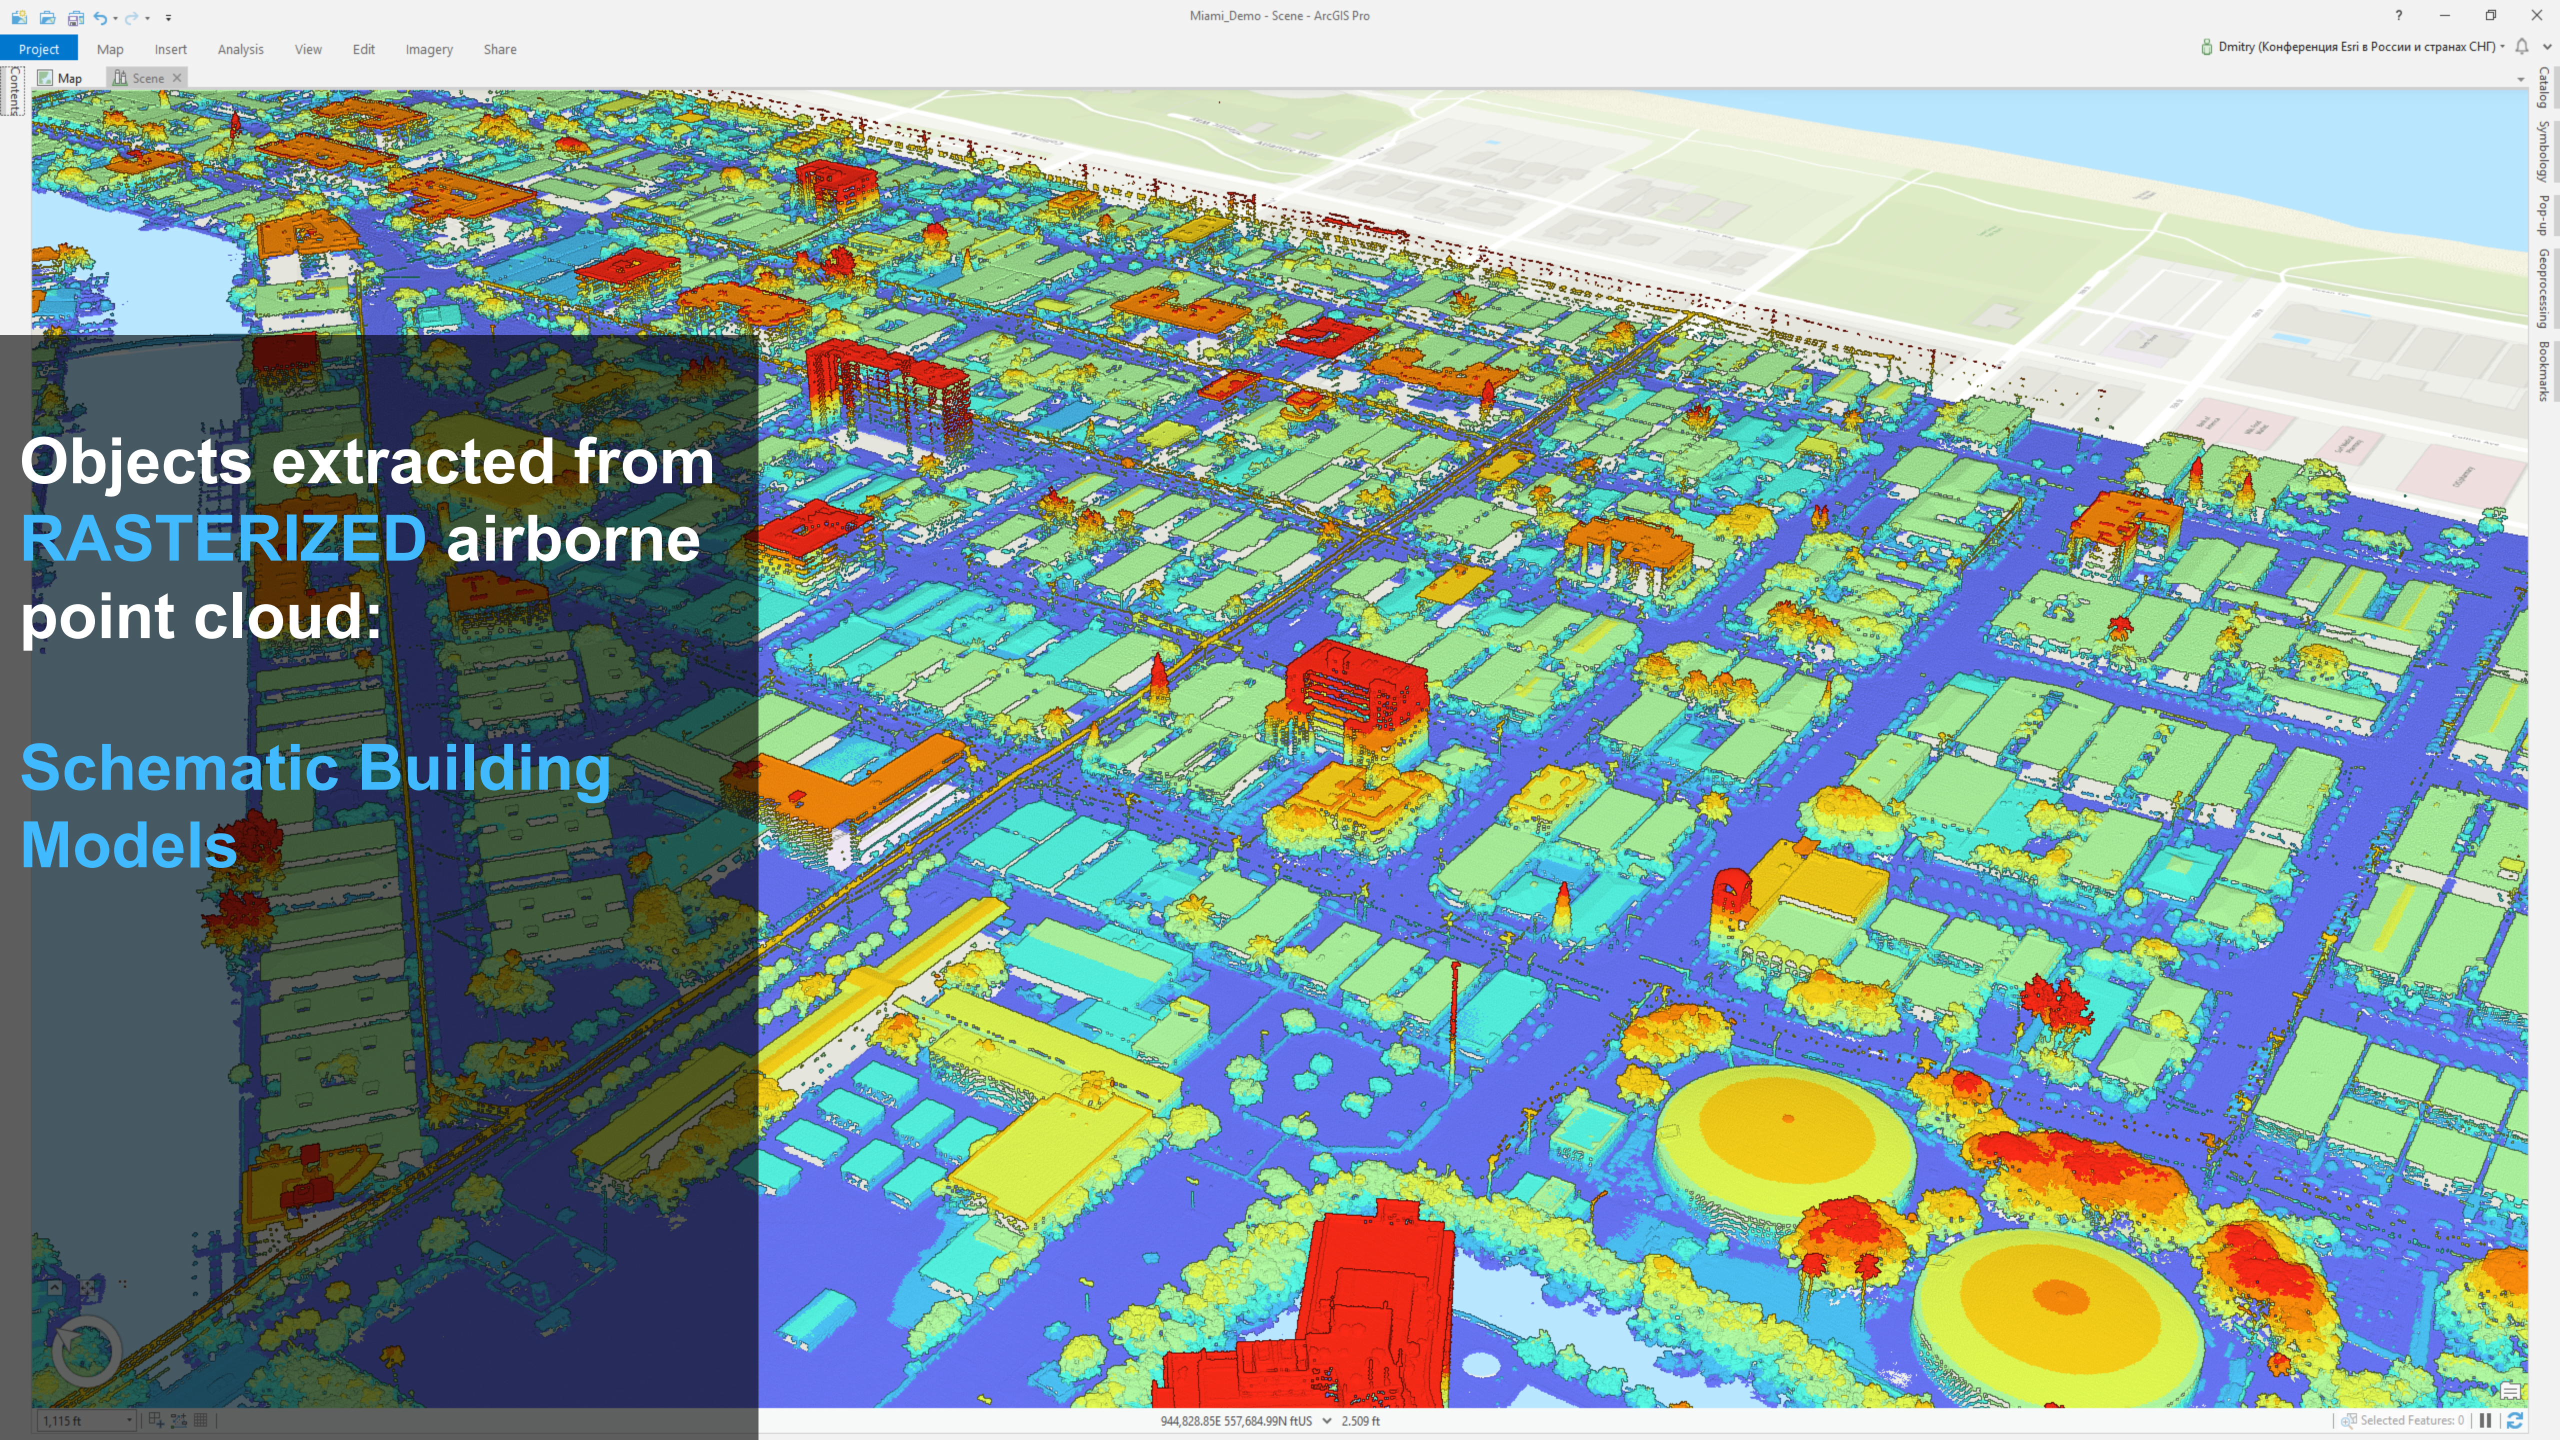
Objects extracted from
RASTERIZED airborne
point cloud:
Schematic Building
Models
Miami_Demo - Scene - ArcGIS Pro
Project
Map
Insert
Analysis
View
Edit
Imagery
Share
Map
Scene
Dmitry (Конференция Esri в России и странах СНГ)
Catalog
Symbology
Pop-up
Geoprocessing
Bookmarks
Contents
1,115 ft
944,828.85E 557,684.99N ftUS
2.509 ft
Selected Features: 0

Objects extracted from **RASTERIZED** airborne point cloud:

Schematic Building Models

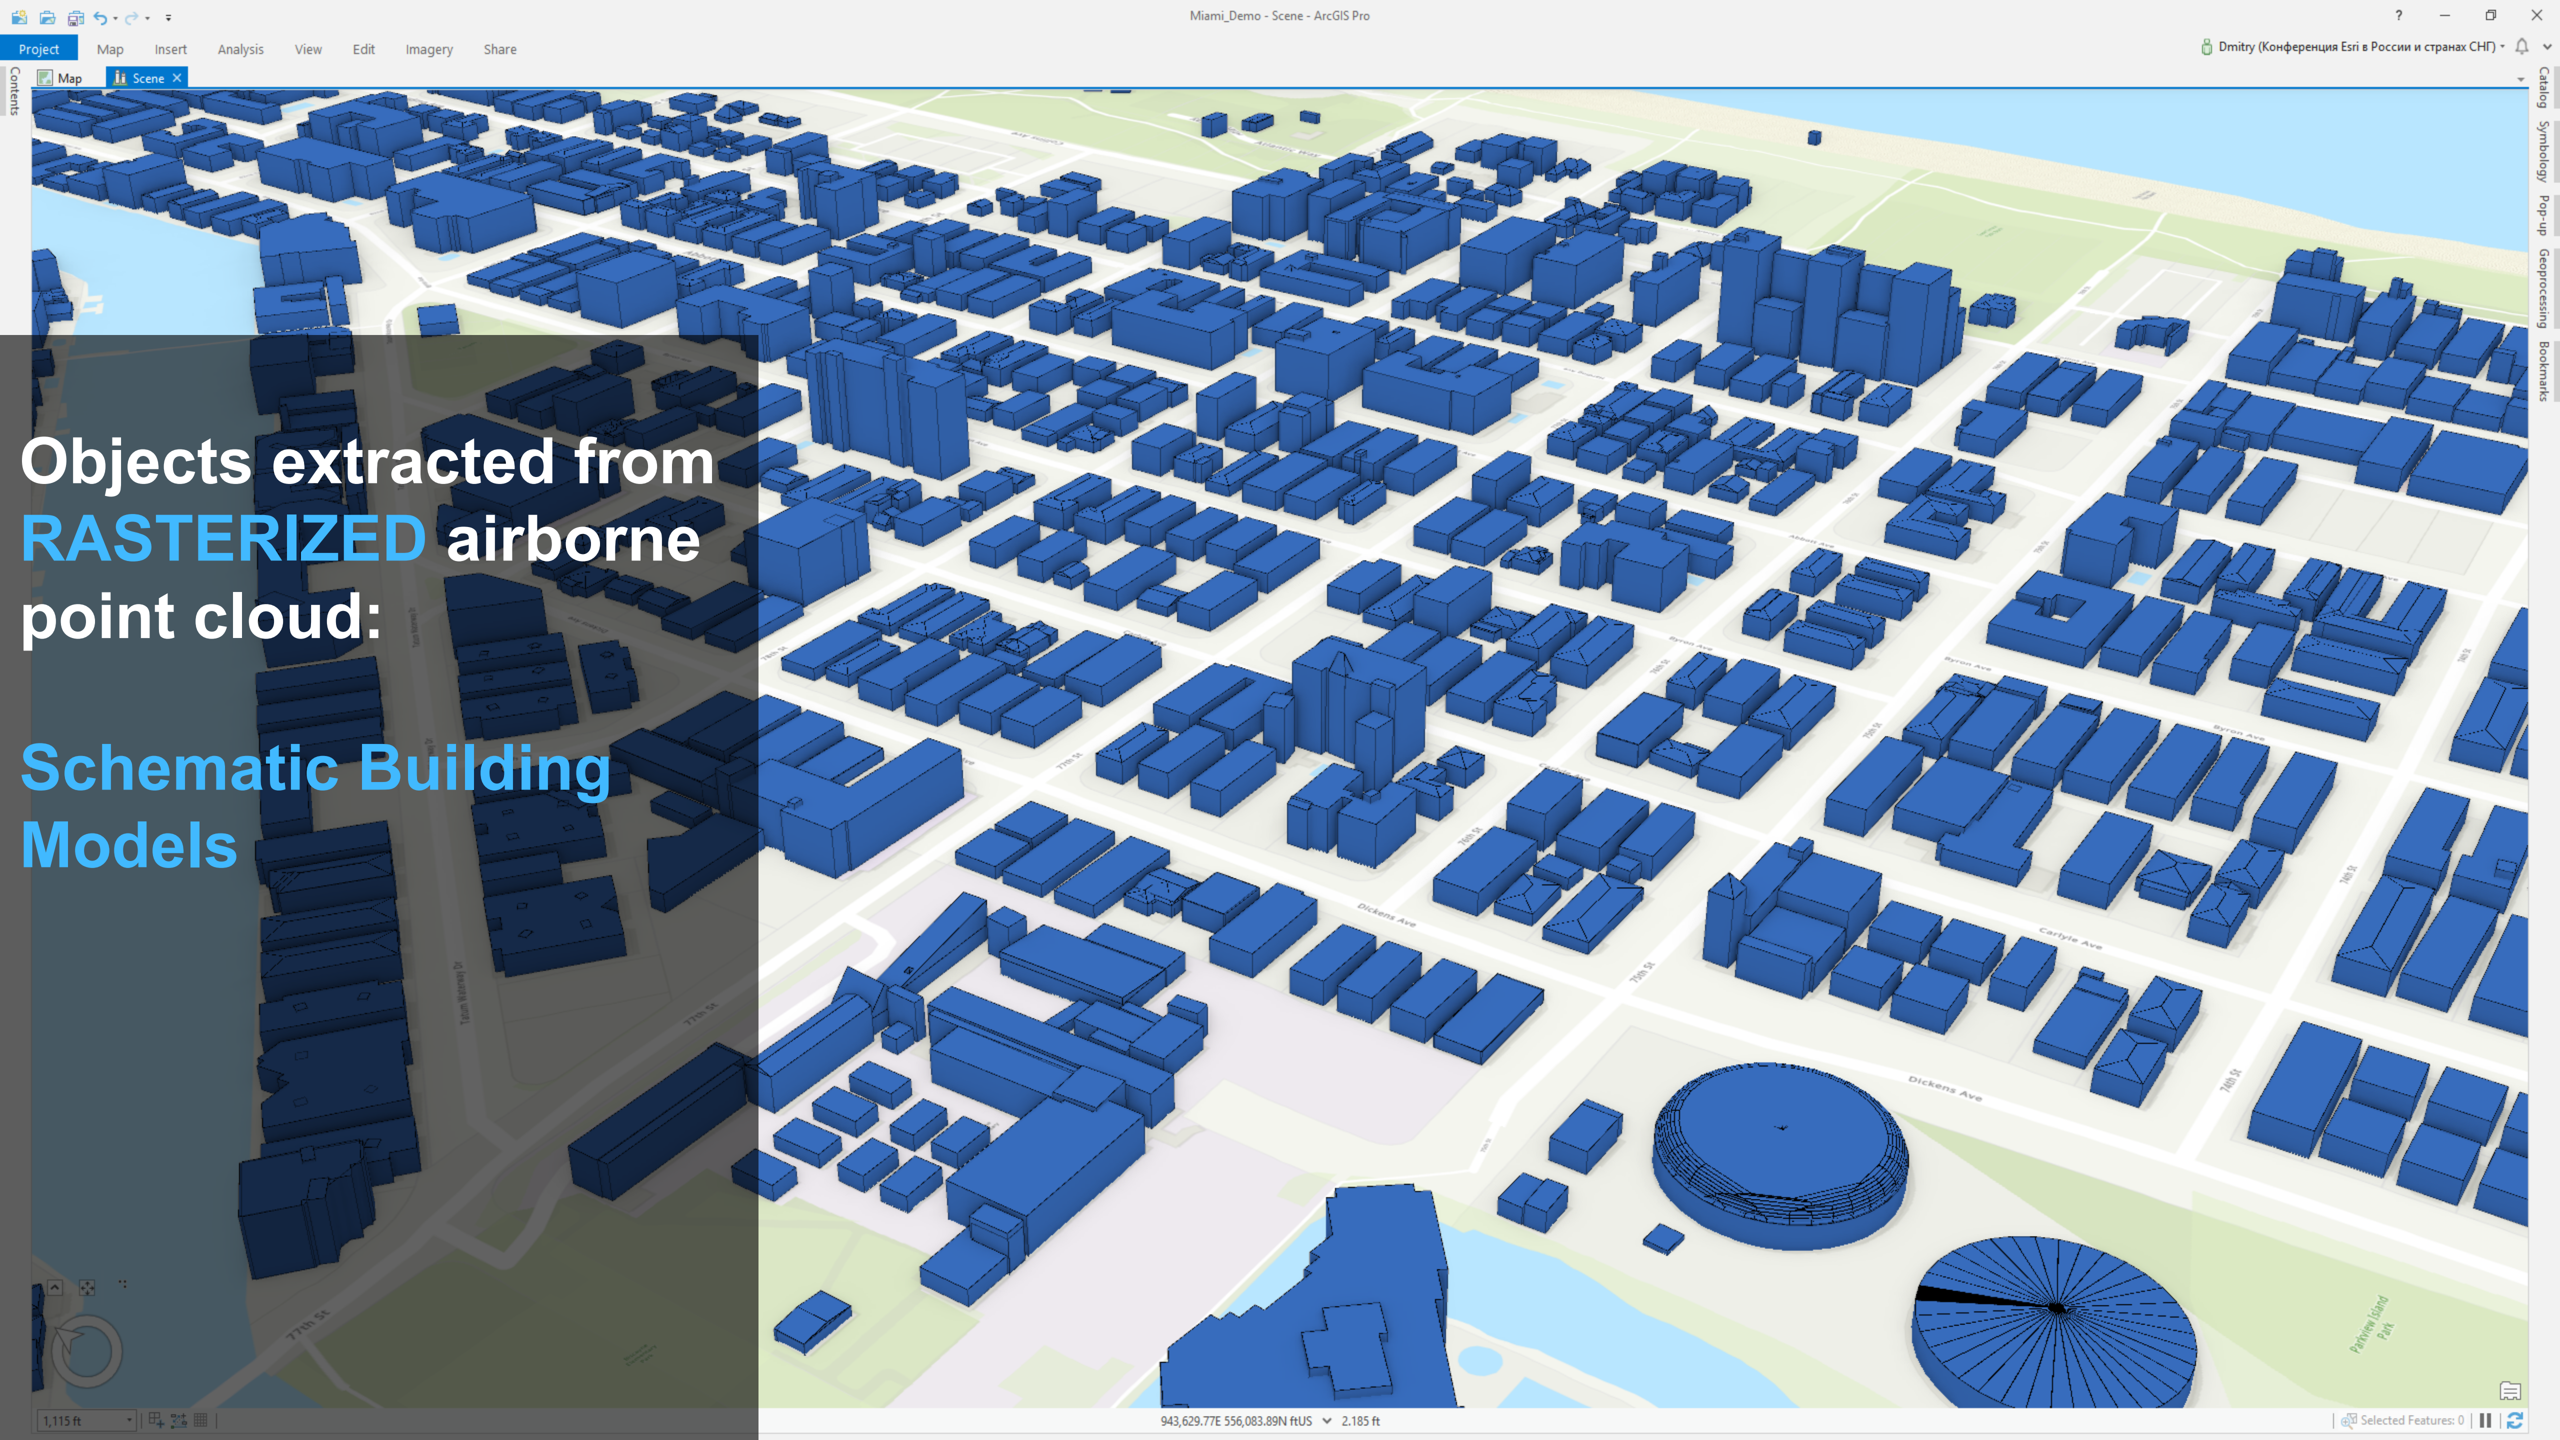

Buildings: Working directly with point clouds

Today ArcGIS allows for reconstruction of buildings directly from point clouds using released GP Tools.

The “LASBuildingMultipatch” workflow:

1. ClassifyLASGround
2. ClassifyLASBuilding
3. LASPointStatisticsAsRaster
   - with LAS layer filtered on class 6 (building)
   - using the ‘Most Frequent Class Code’ option
4. RasterToPolygon
5. EliminatePolygonPart
6. RegularizeBuildingFootprint
7. LASDatasetToRaster
   - with input LAS layer filtered on class 2 points to make DEM
8. LASBuildingMultipatch

Buildings: Working directly with point clouds

Models contain many faces and are not suitable for manual editing.

Noise level depends on accuracy of point cloud classification.

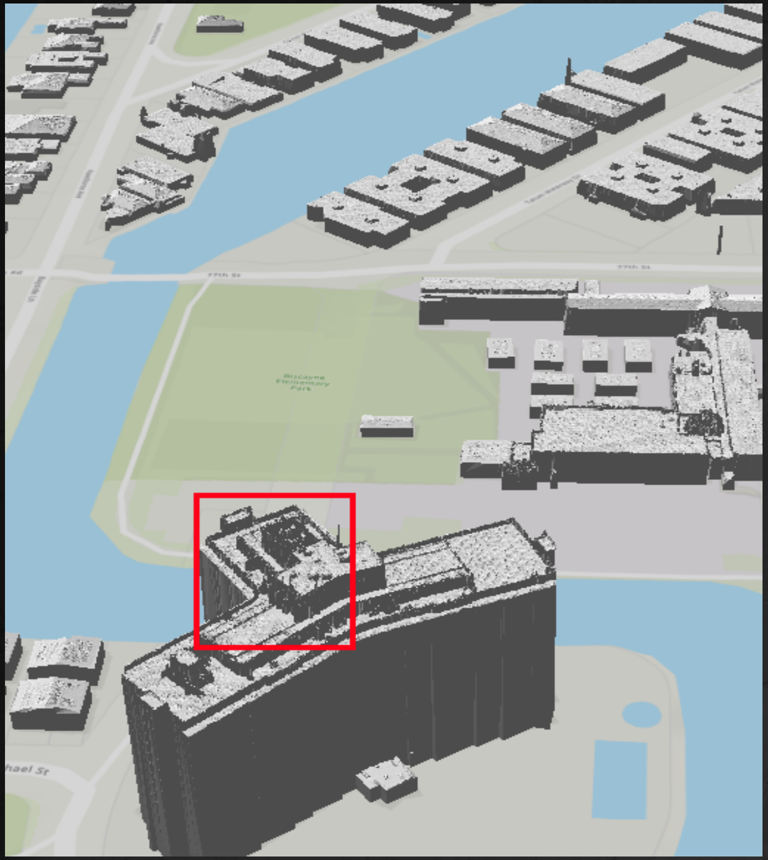

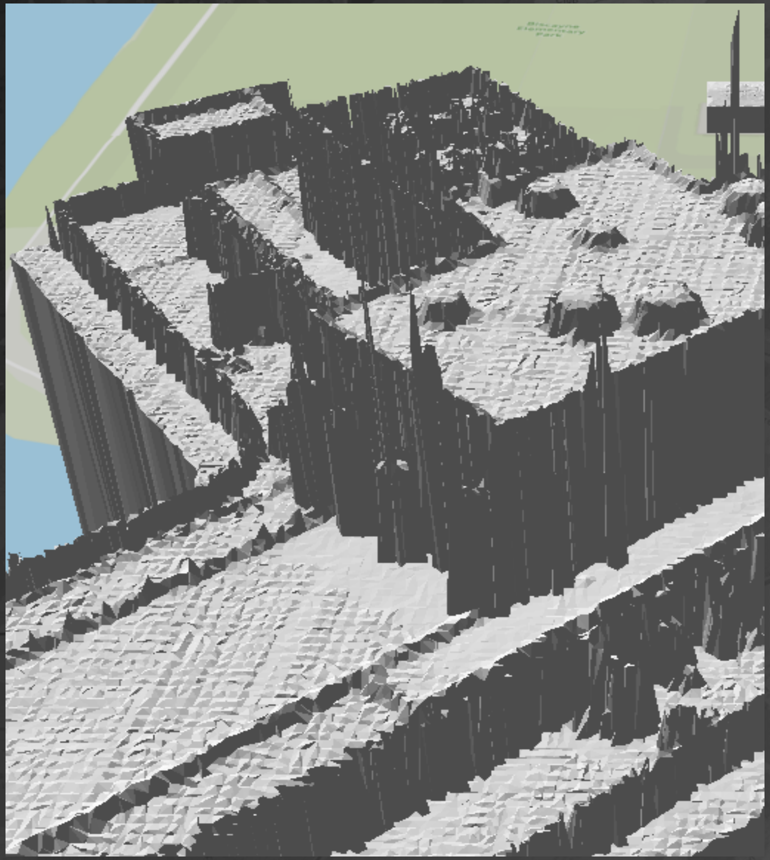

Can we use Deep Learning to label point clouds?

- Point clouds are irregular and unordered, cannot apply convolution ops directly.
- Good news: multiple developments, DL architectures, and papers in recent years: PointNet++, Graph Convolutional networks, Deep Sets, PointCNN, etc.

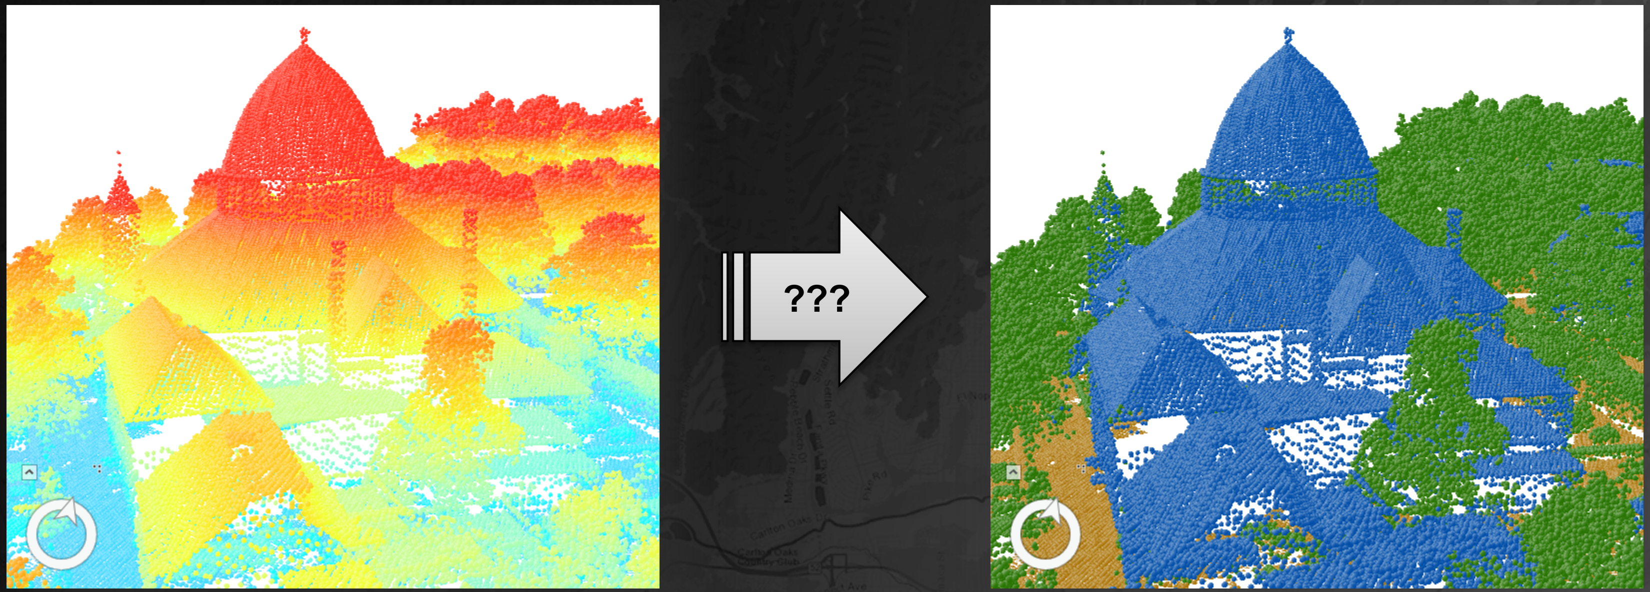

Buildings: Working directly with point clouds

- PointCNN trained to classify buildings in airborne point cloud

LASBuildingMultipatch workflow: with & w/o PointCNN-labeled building points

- Lower noise level in resulting building models

- PointCNN after 6.5 hours (GV100) of training on XYZ-geometry only.

PointCNN: Wires & Poles

from airborne LiDAR

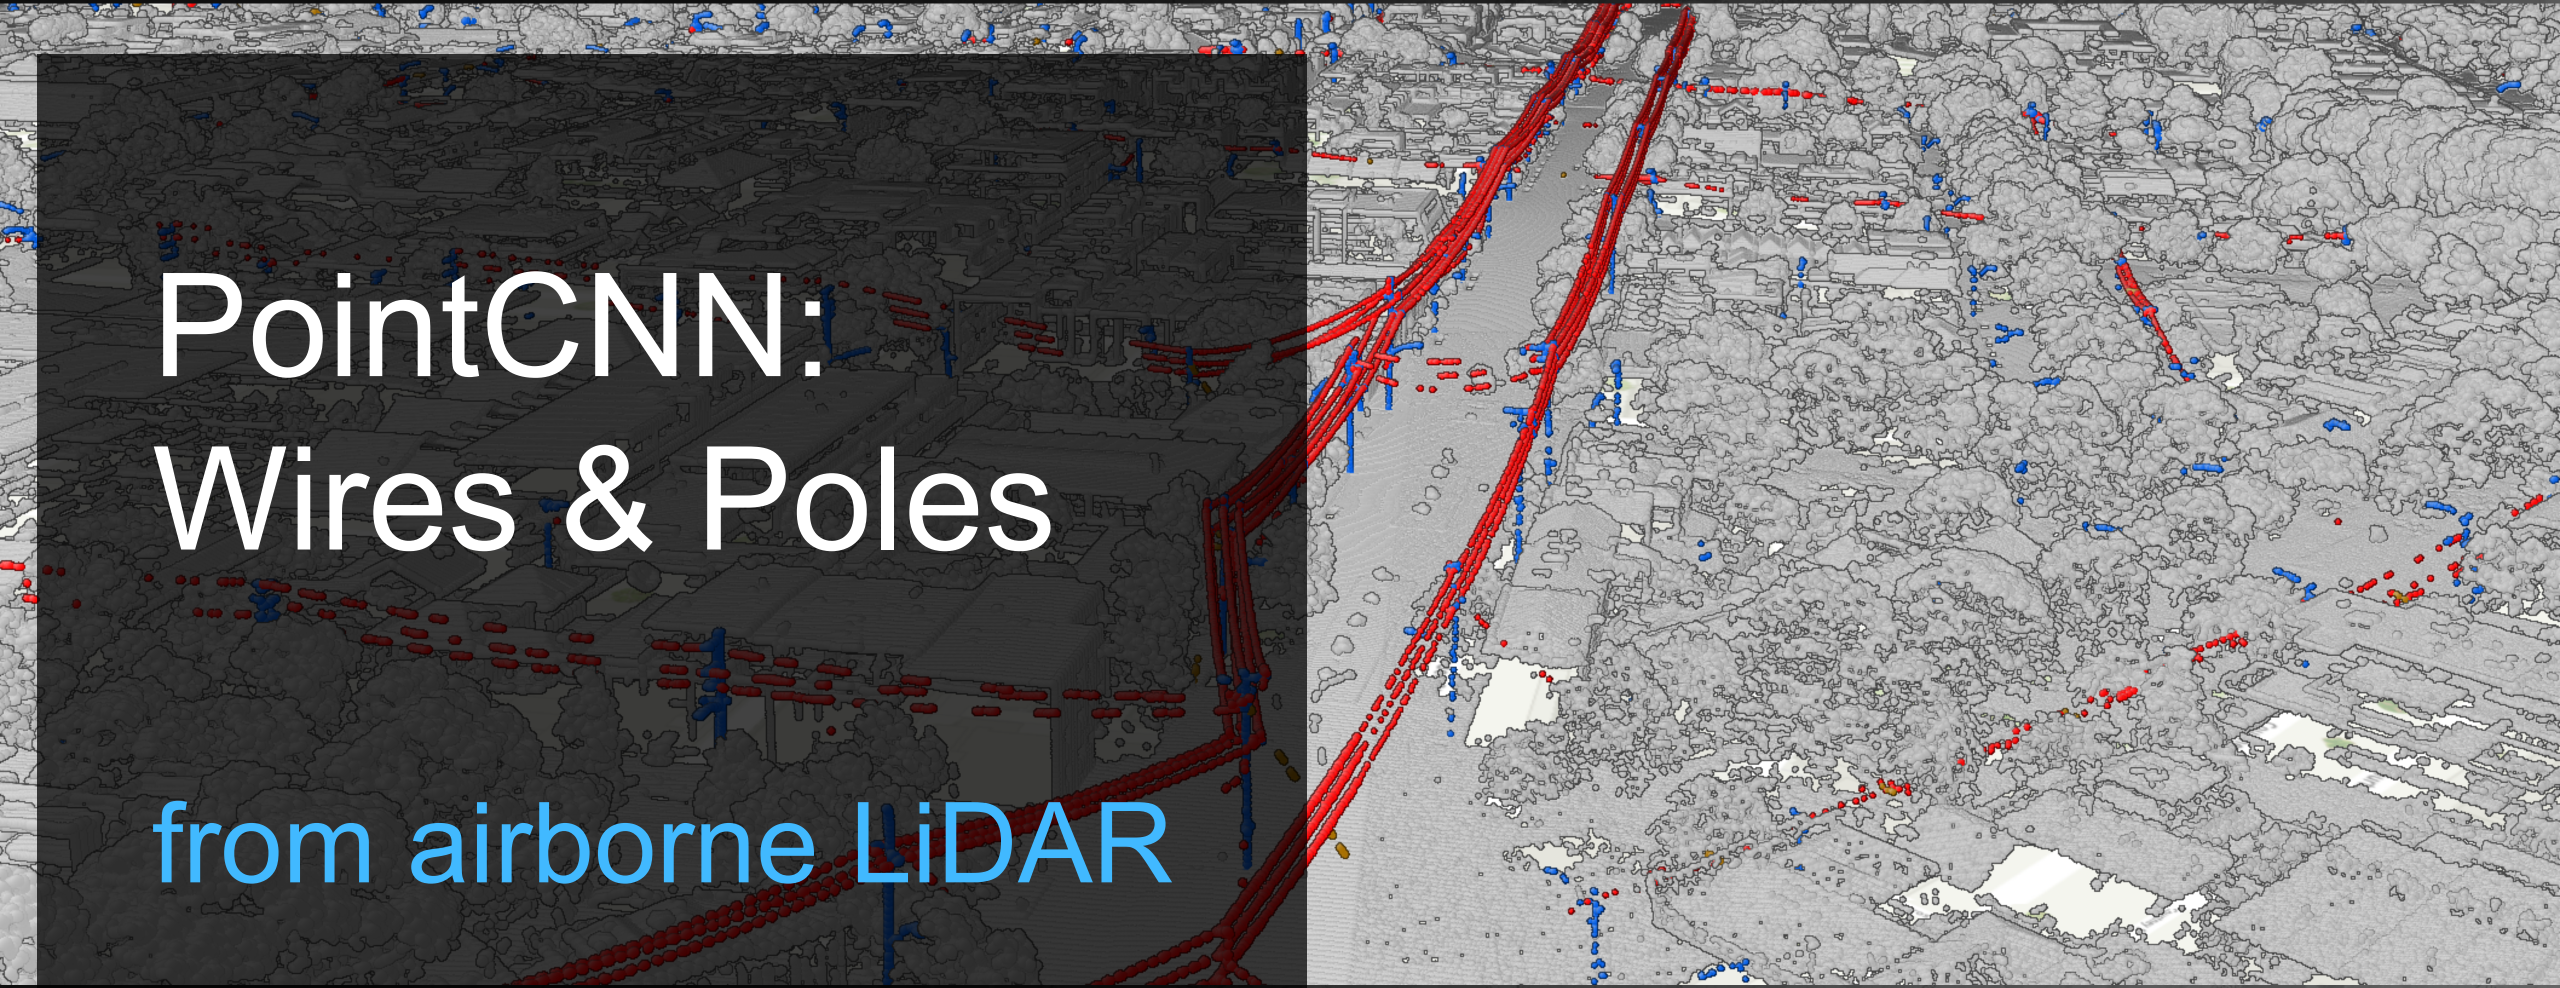

THE SCIENCE OF WHERE

PoC: Wires detection in airborne LiDAR

AAM Group, Australia: collecting airborne LiDAR point clouds to detect the power lines and any easement encroachment.

50,000 man hours per year of manual labeling

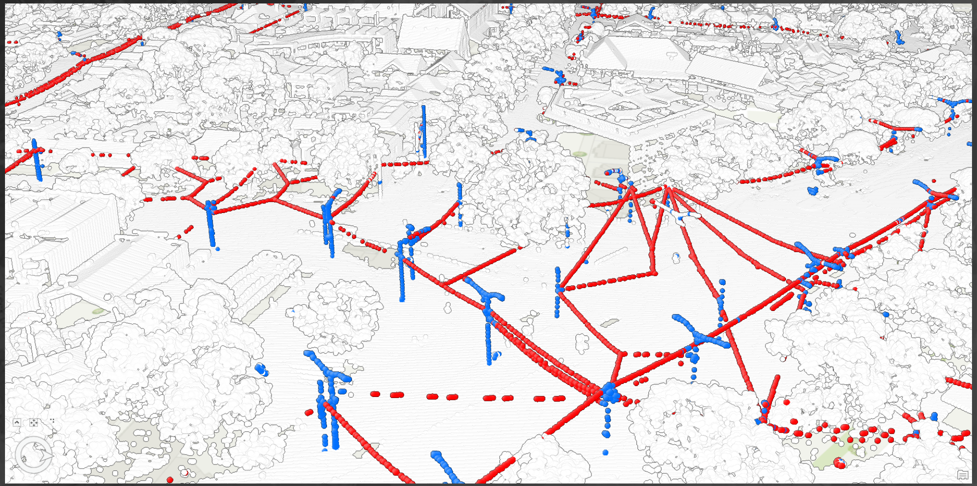

PoC: Wires detection in airborne LiDAR

- AAM Group shared 10B+ of manually labeled points to train a neural network.
- Thousands of miles of transmission lines.

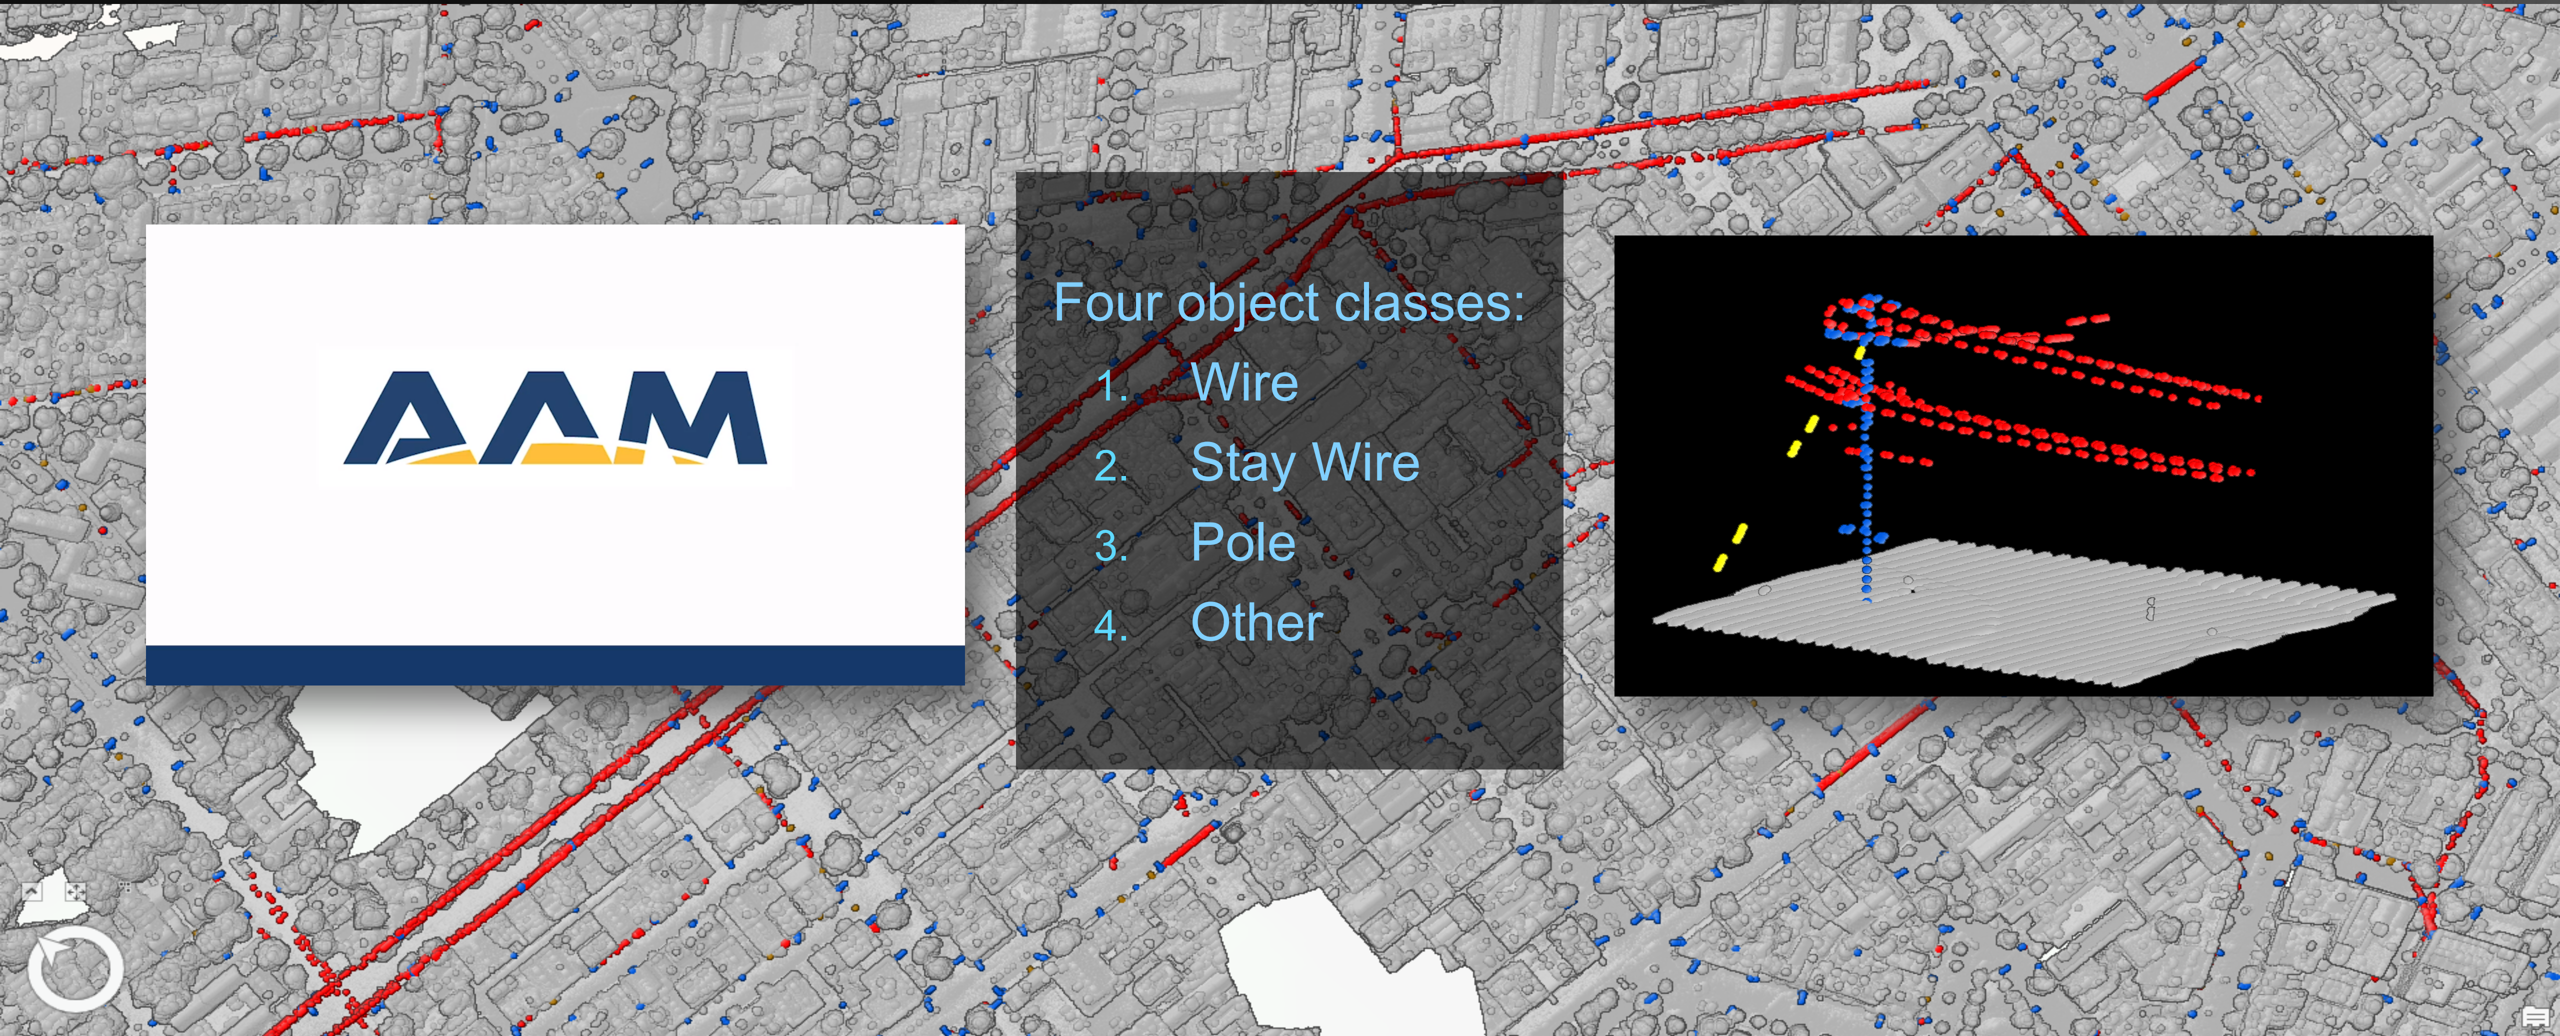

AAM

Four object classes:

1. Wire
2. Stay Wire
3. Pole
4. Other

PoC: Overhead conductor inspection

…but there are some good news too -

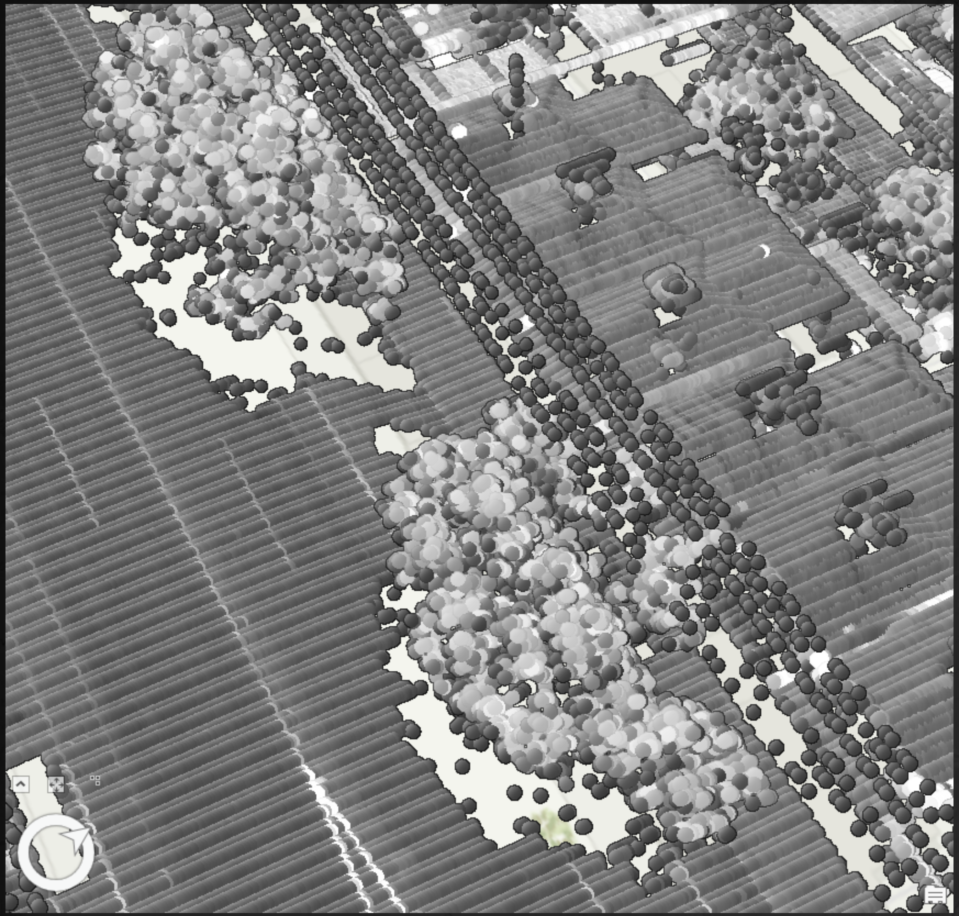

Intensity and the Number of Returns on the Wire points is often different than of the surroundings

This allows for training PointCNN on
XYZ +
Intensity +
Number of Returns

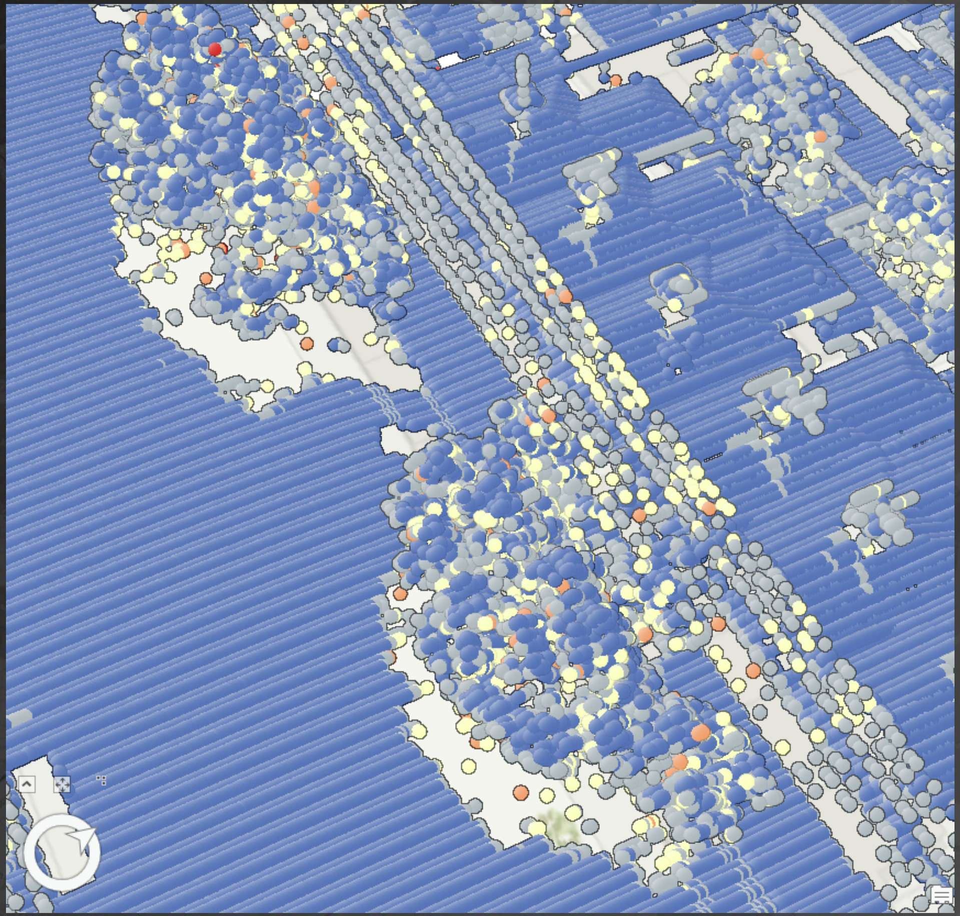

PoC: Overhead conductor inspection

Training & Results:

after training on a single GV100 for ~20 hours

@E232K best RECALL

	OTHER	WIRE	STAY-WIRE	POLE
Precision:	[0.99988538	0.96672749	0.83674406	0.80313546],
Recall:	[0.99987221	0.98060249	0.21455632	0.77497643],
F1 score:	[0.9998788	0.97361556	0.3415365	0.78880471]

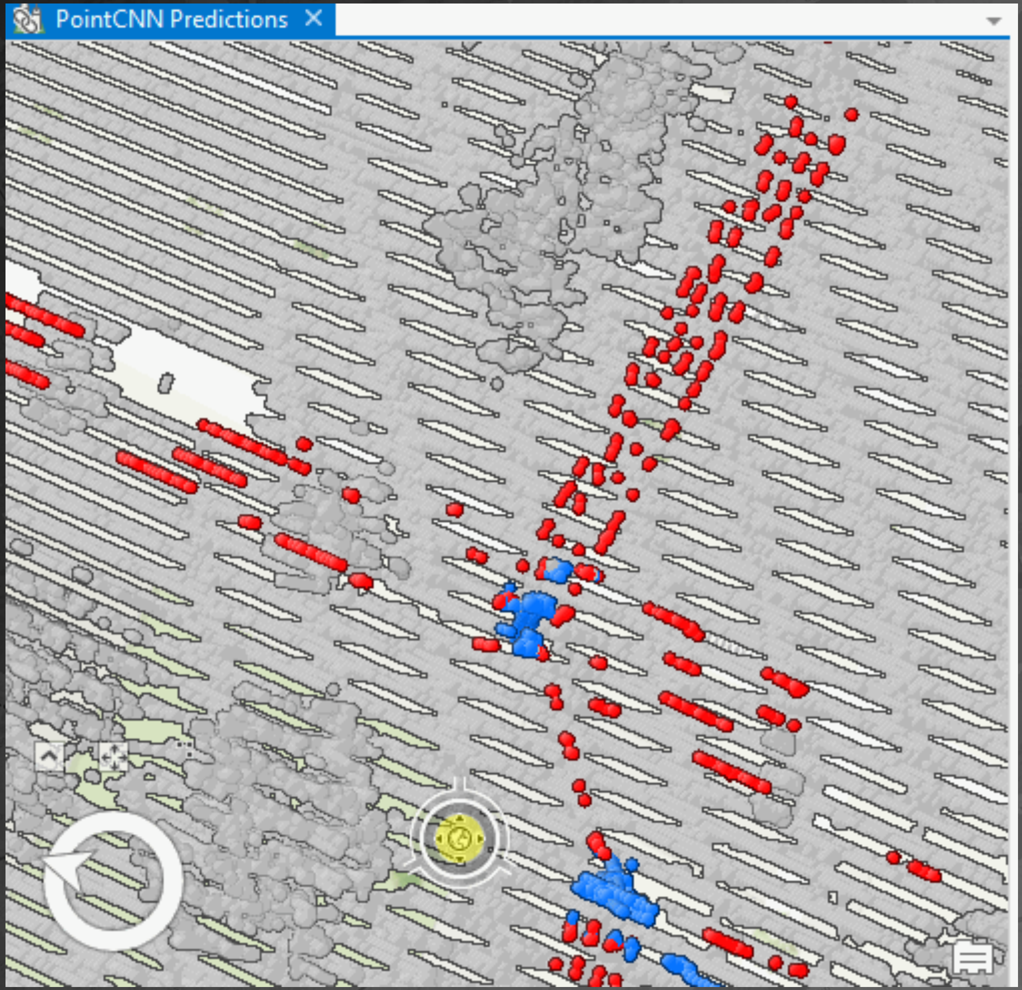

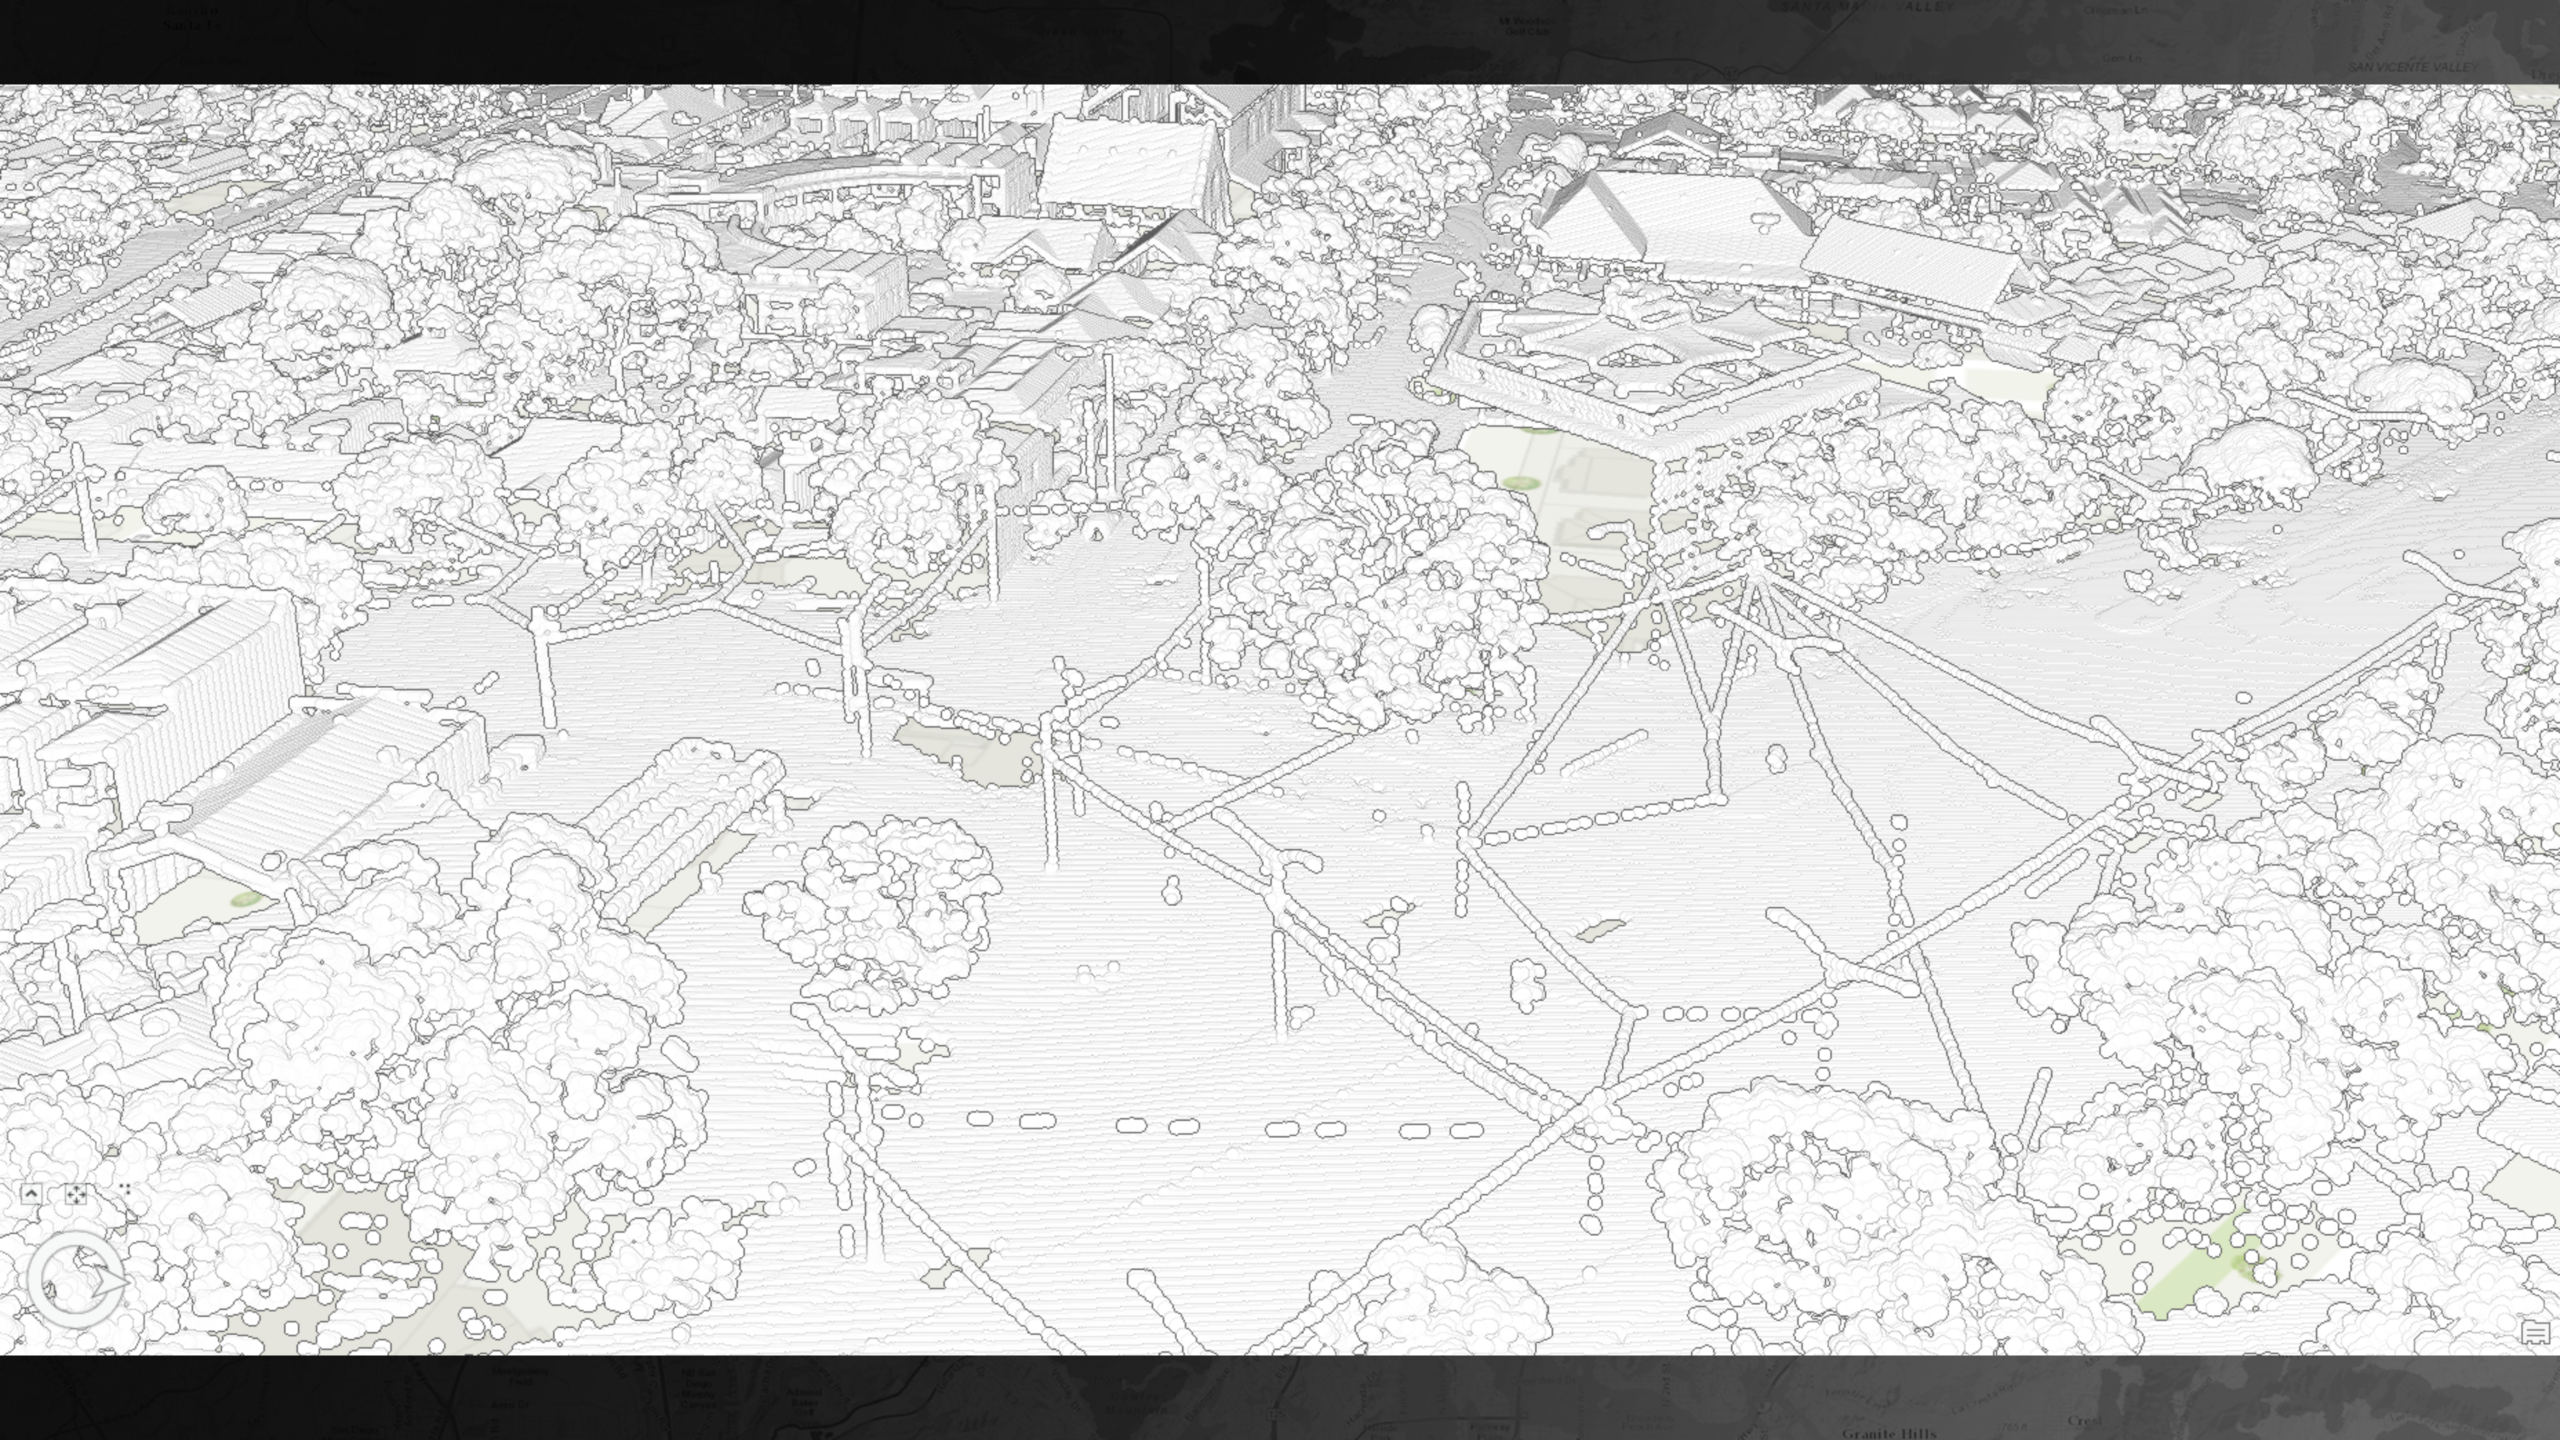
SAN VICENTE VALLEY
Granite Hills
Crest

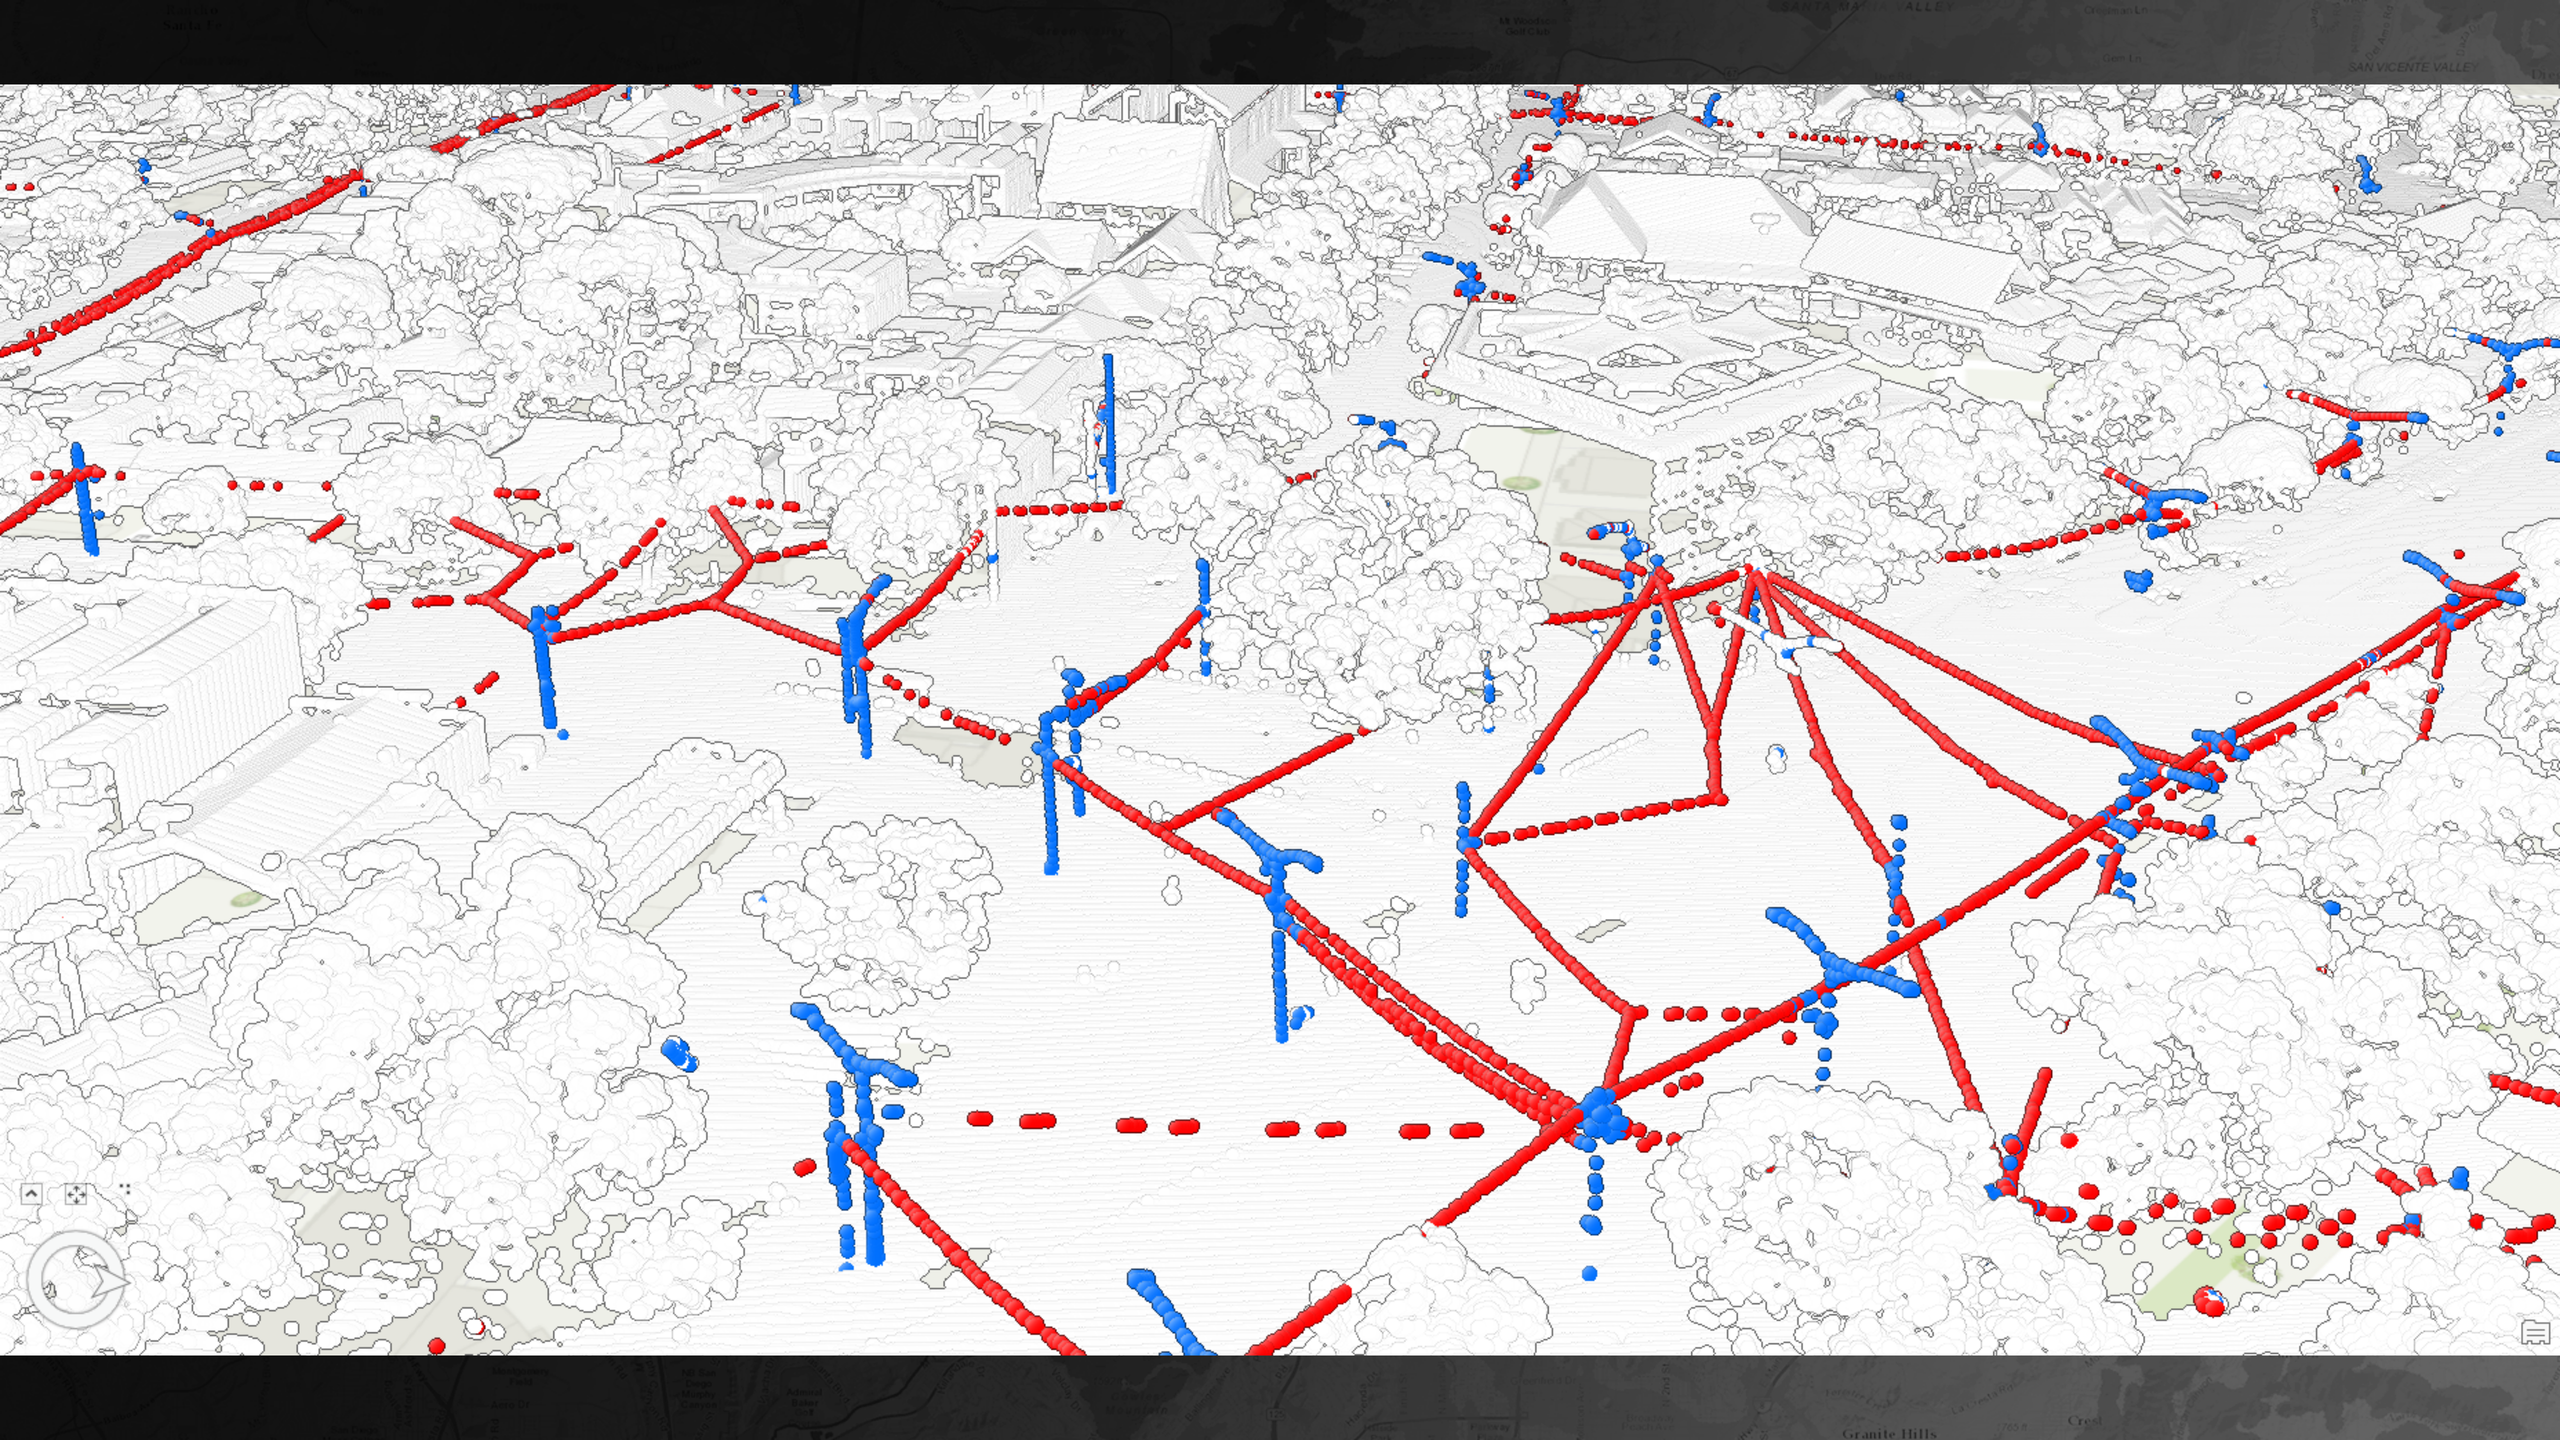

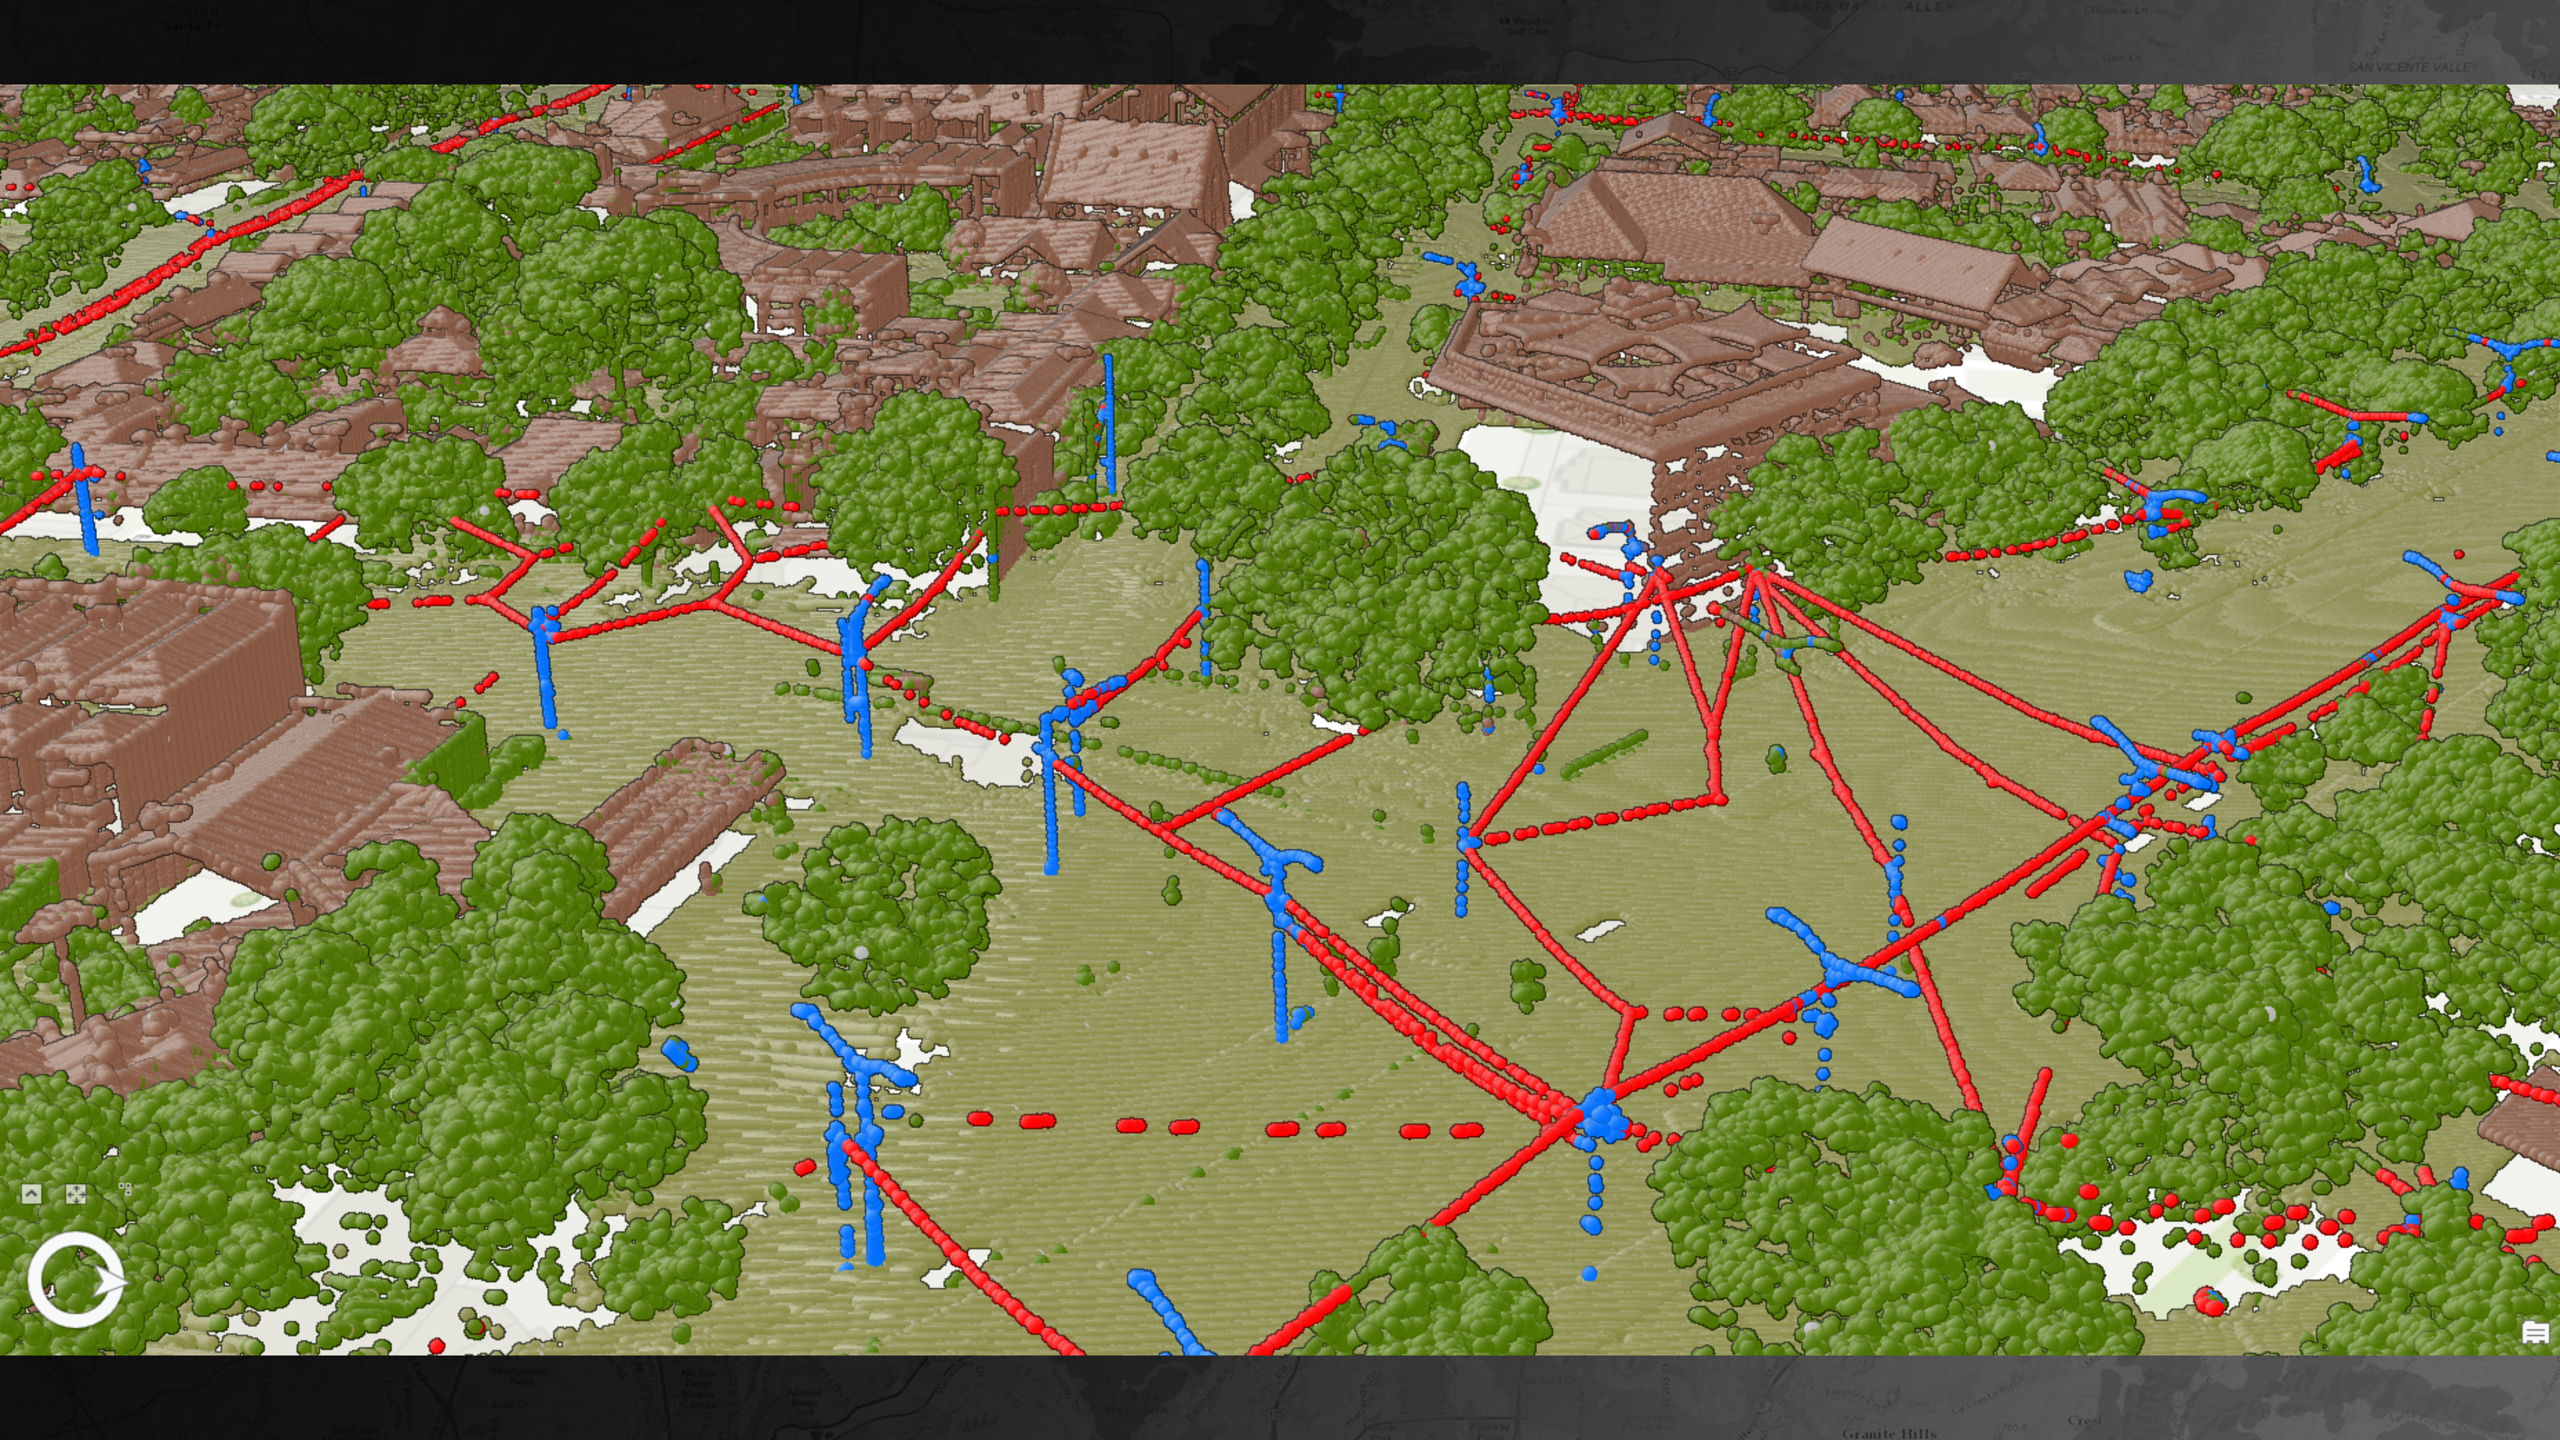

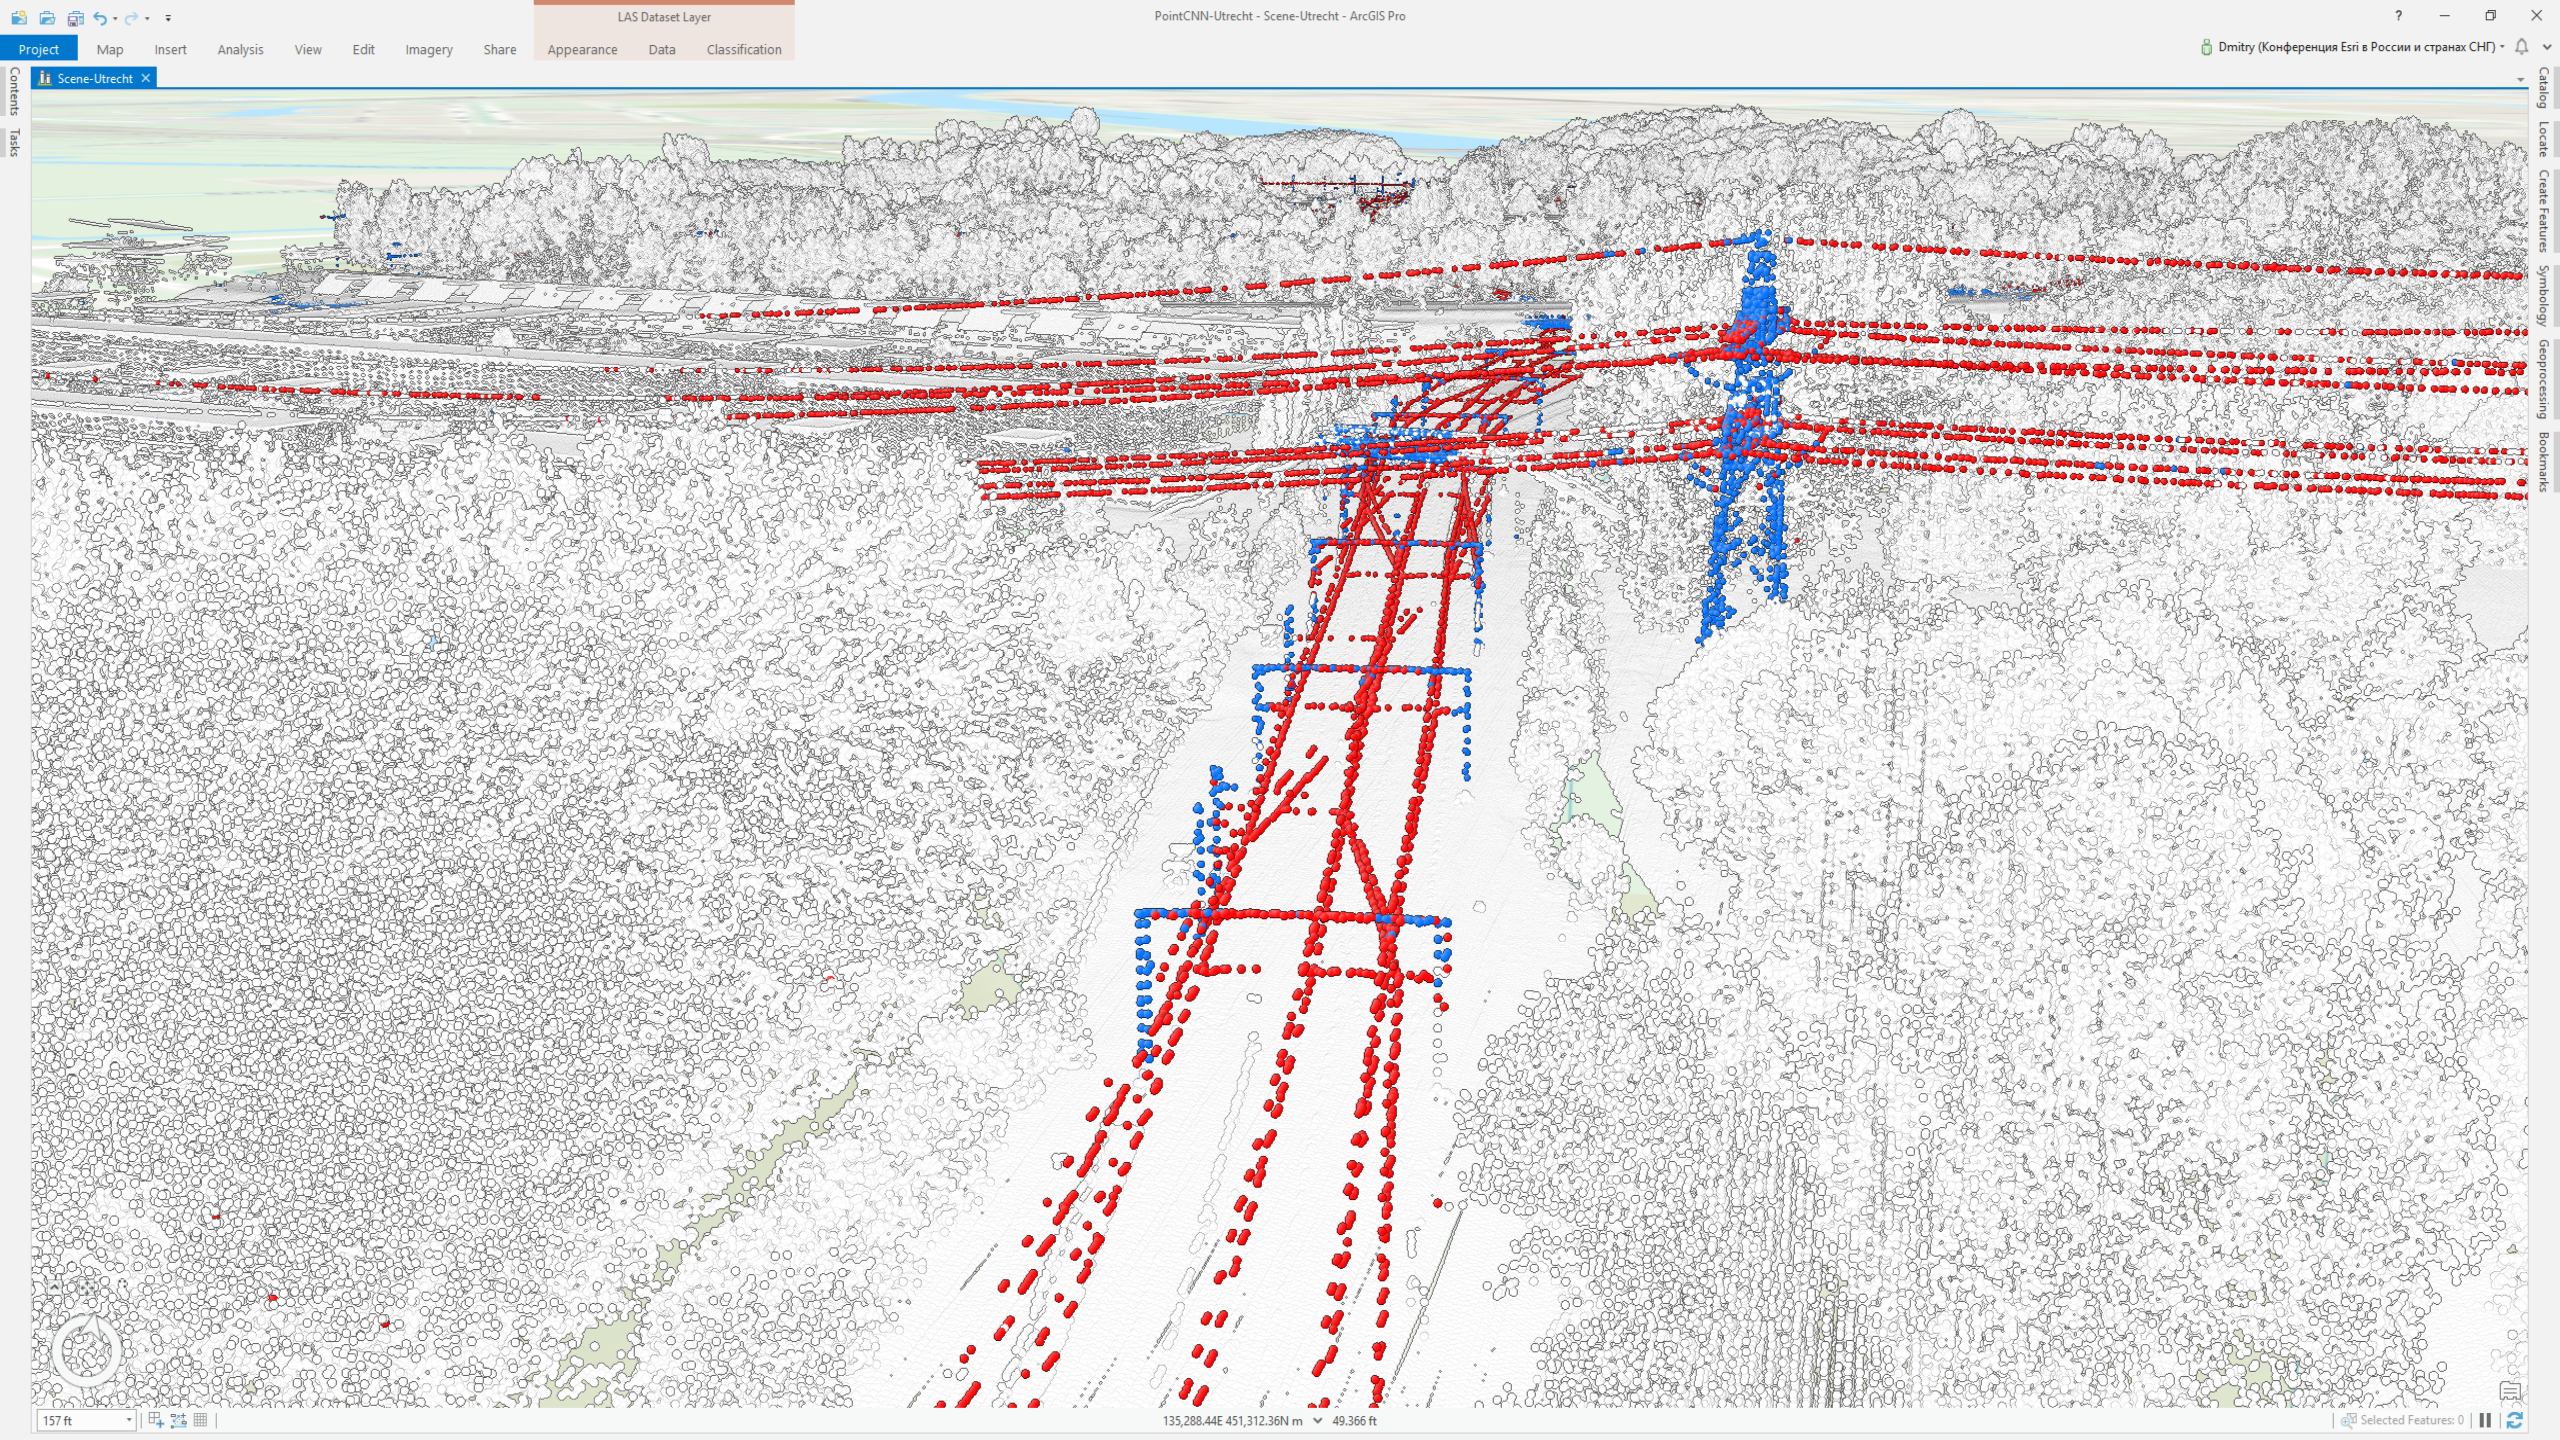
LAS Dataset Layer
PointCNN-Utrecht - Scene-Utrecht - ArcGIS Pro
Project
Map
Insert
Analysis
View
Edit
Imagery
Share
Appearance
Data
Classification
Dmitry (Конференция Esri в России и странах СНГ)
Scene-Utrecht
Contents
Tasks
Catalog
Locate
Create Features
Symbology
Geoprocessing
Bookmarks
157 ft
135,288.44E 451,312.36N m
49.366 ft
Selected Features: 0

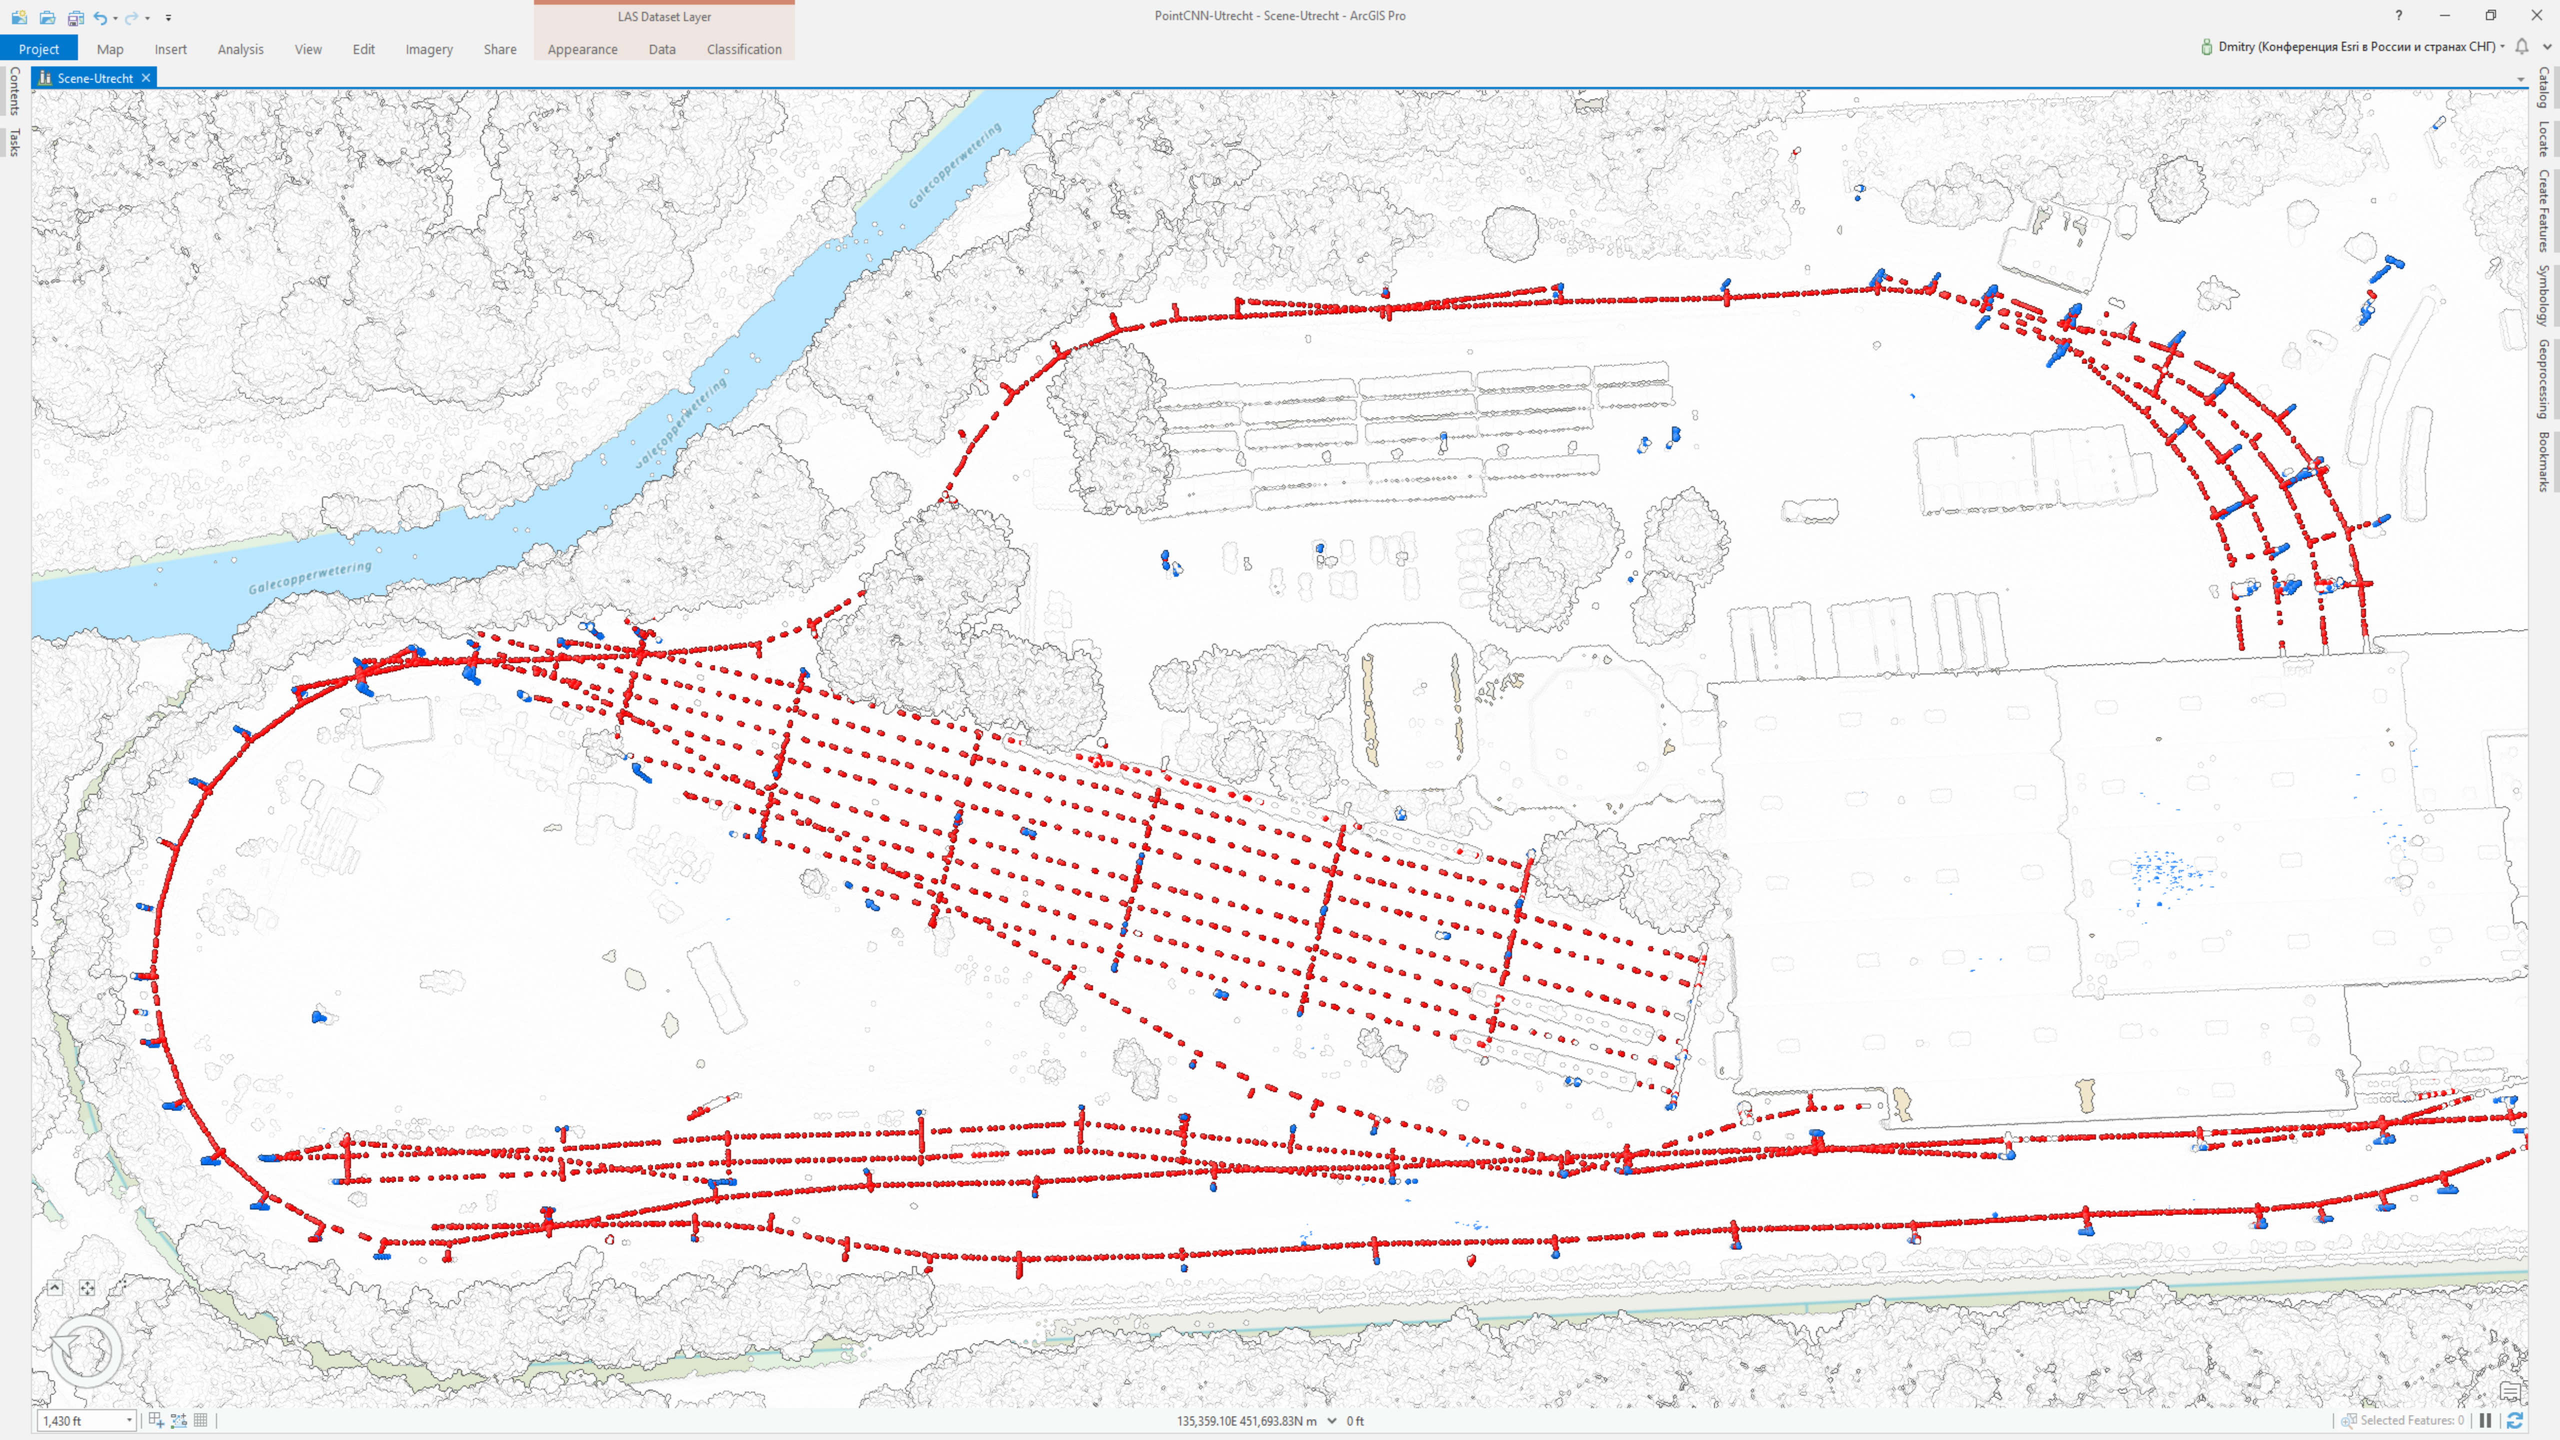
LAS Dataset Layer
PointCNN-Utrecht - Scene-Utrecht - ArcGIS Pro
Project
Map
Insert
Analysis
View
Edit
Imagery
Share
Appearance
Data
Classification
Dmitry (Конференция Esri в России и странах СНГ)
Scene-Utrecht
Contents
Tasks
Catalog
Locate
Create Features
Symbology
Geoprocessing
Bookmarks
Galecopperwetering
Galecopperwetering
Galecopperwetering
1,430 ft
135,359.10E 451,693.83N m
0 ft
Selected Features: 0

PointCNN:
mobile point clouds

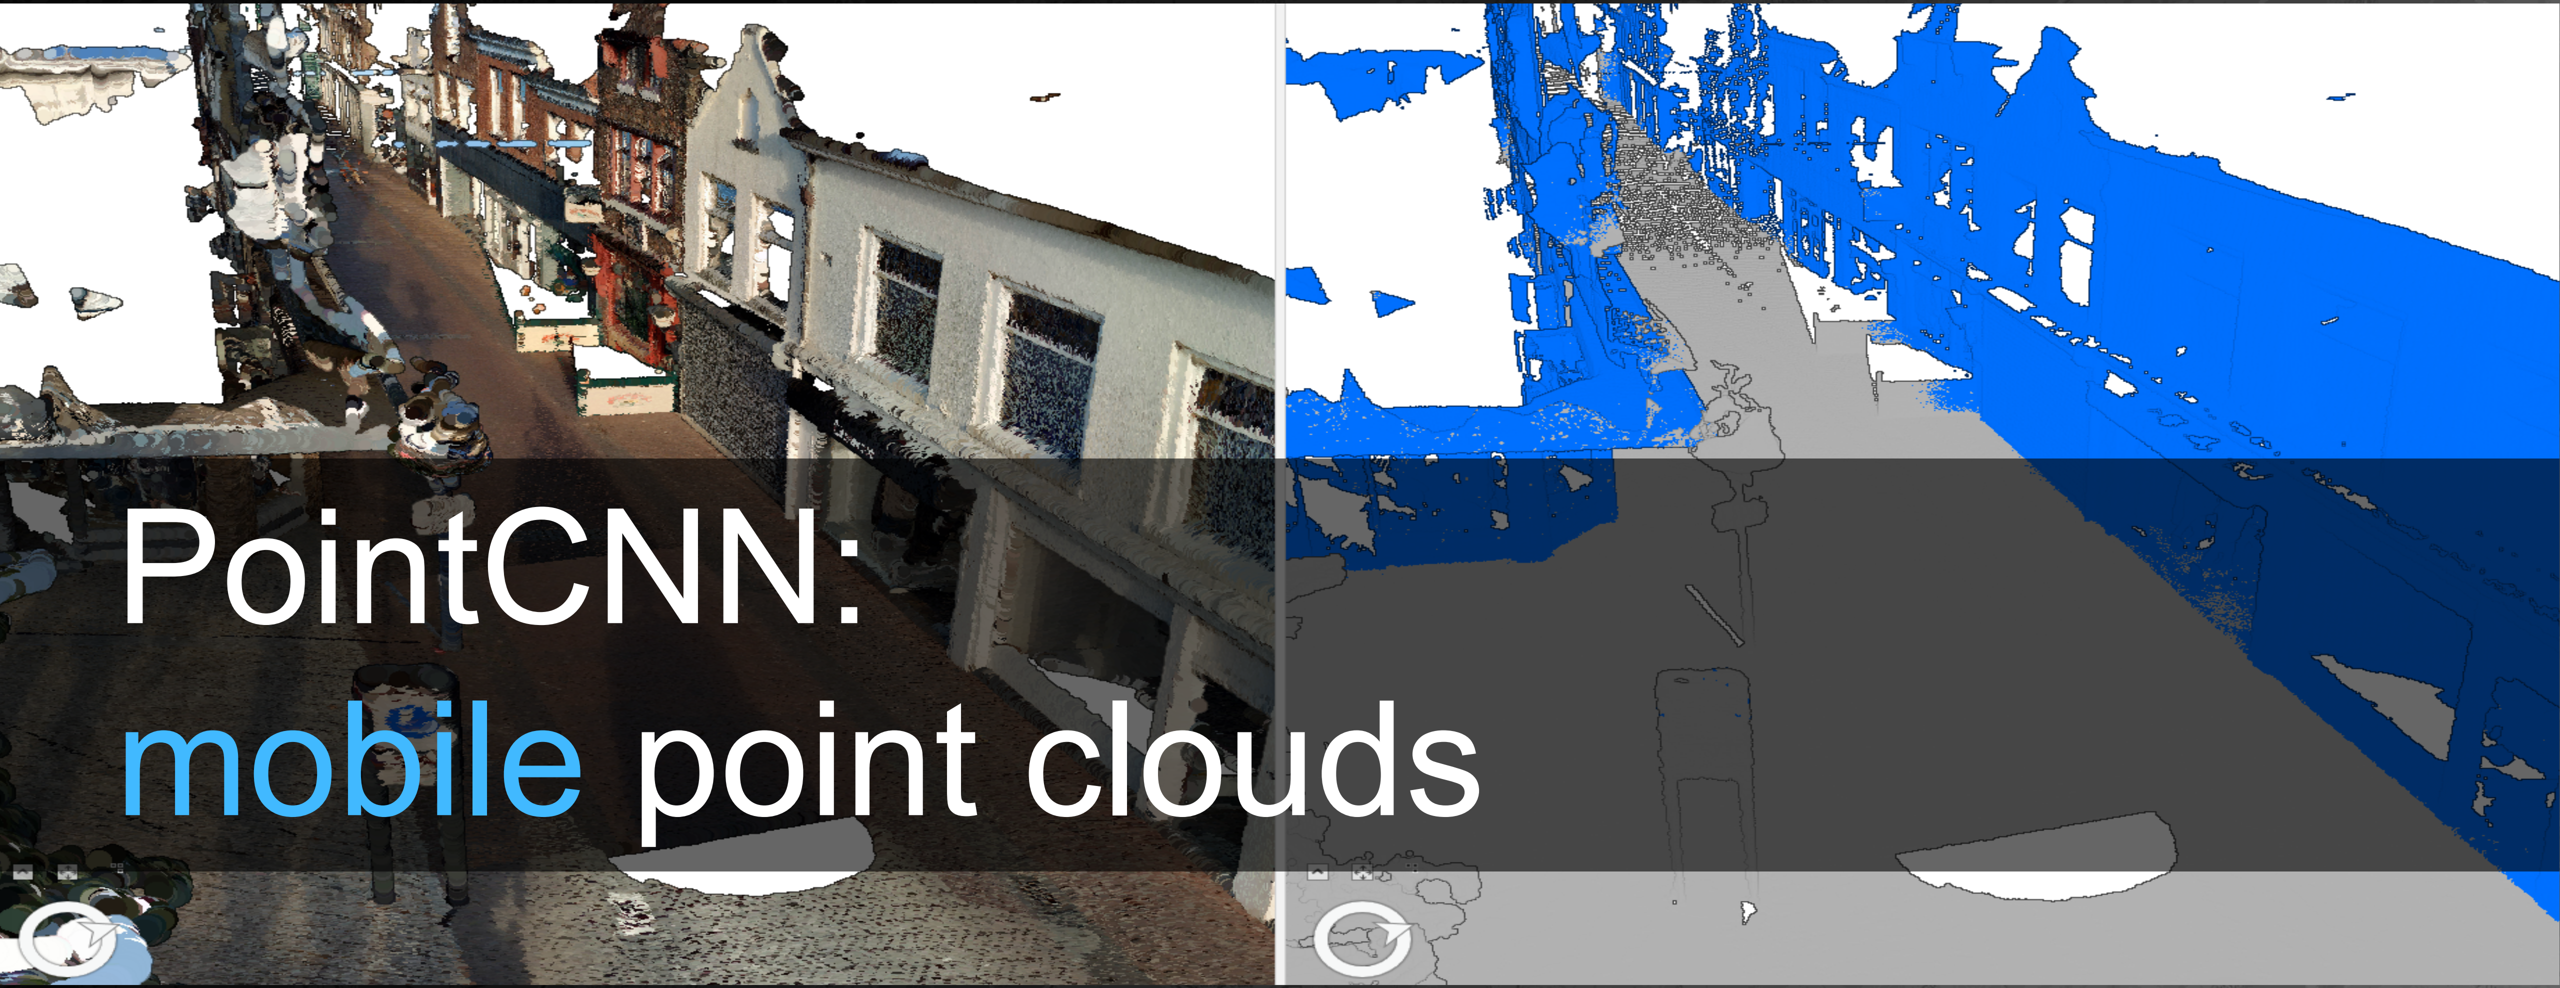

esri THE SCIENCE OF WHERE

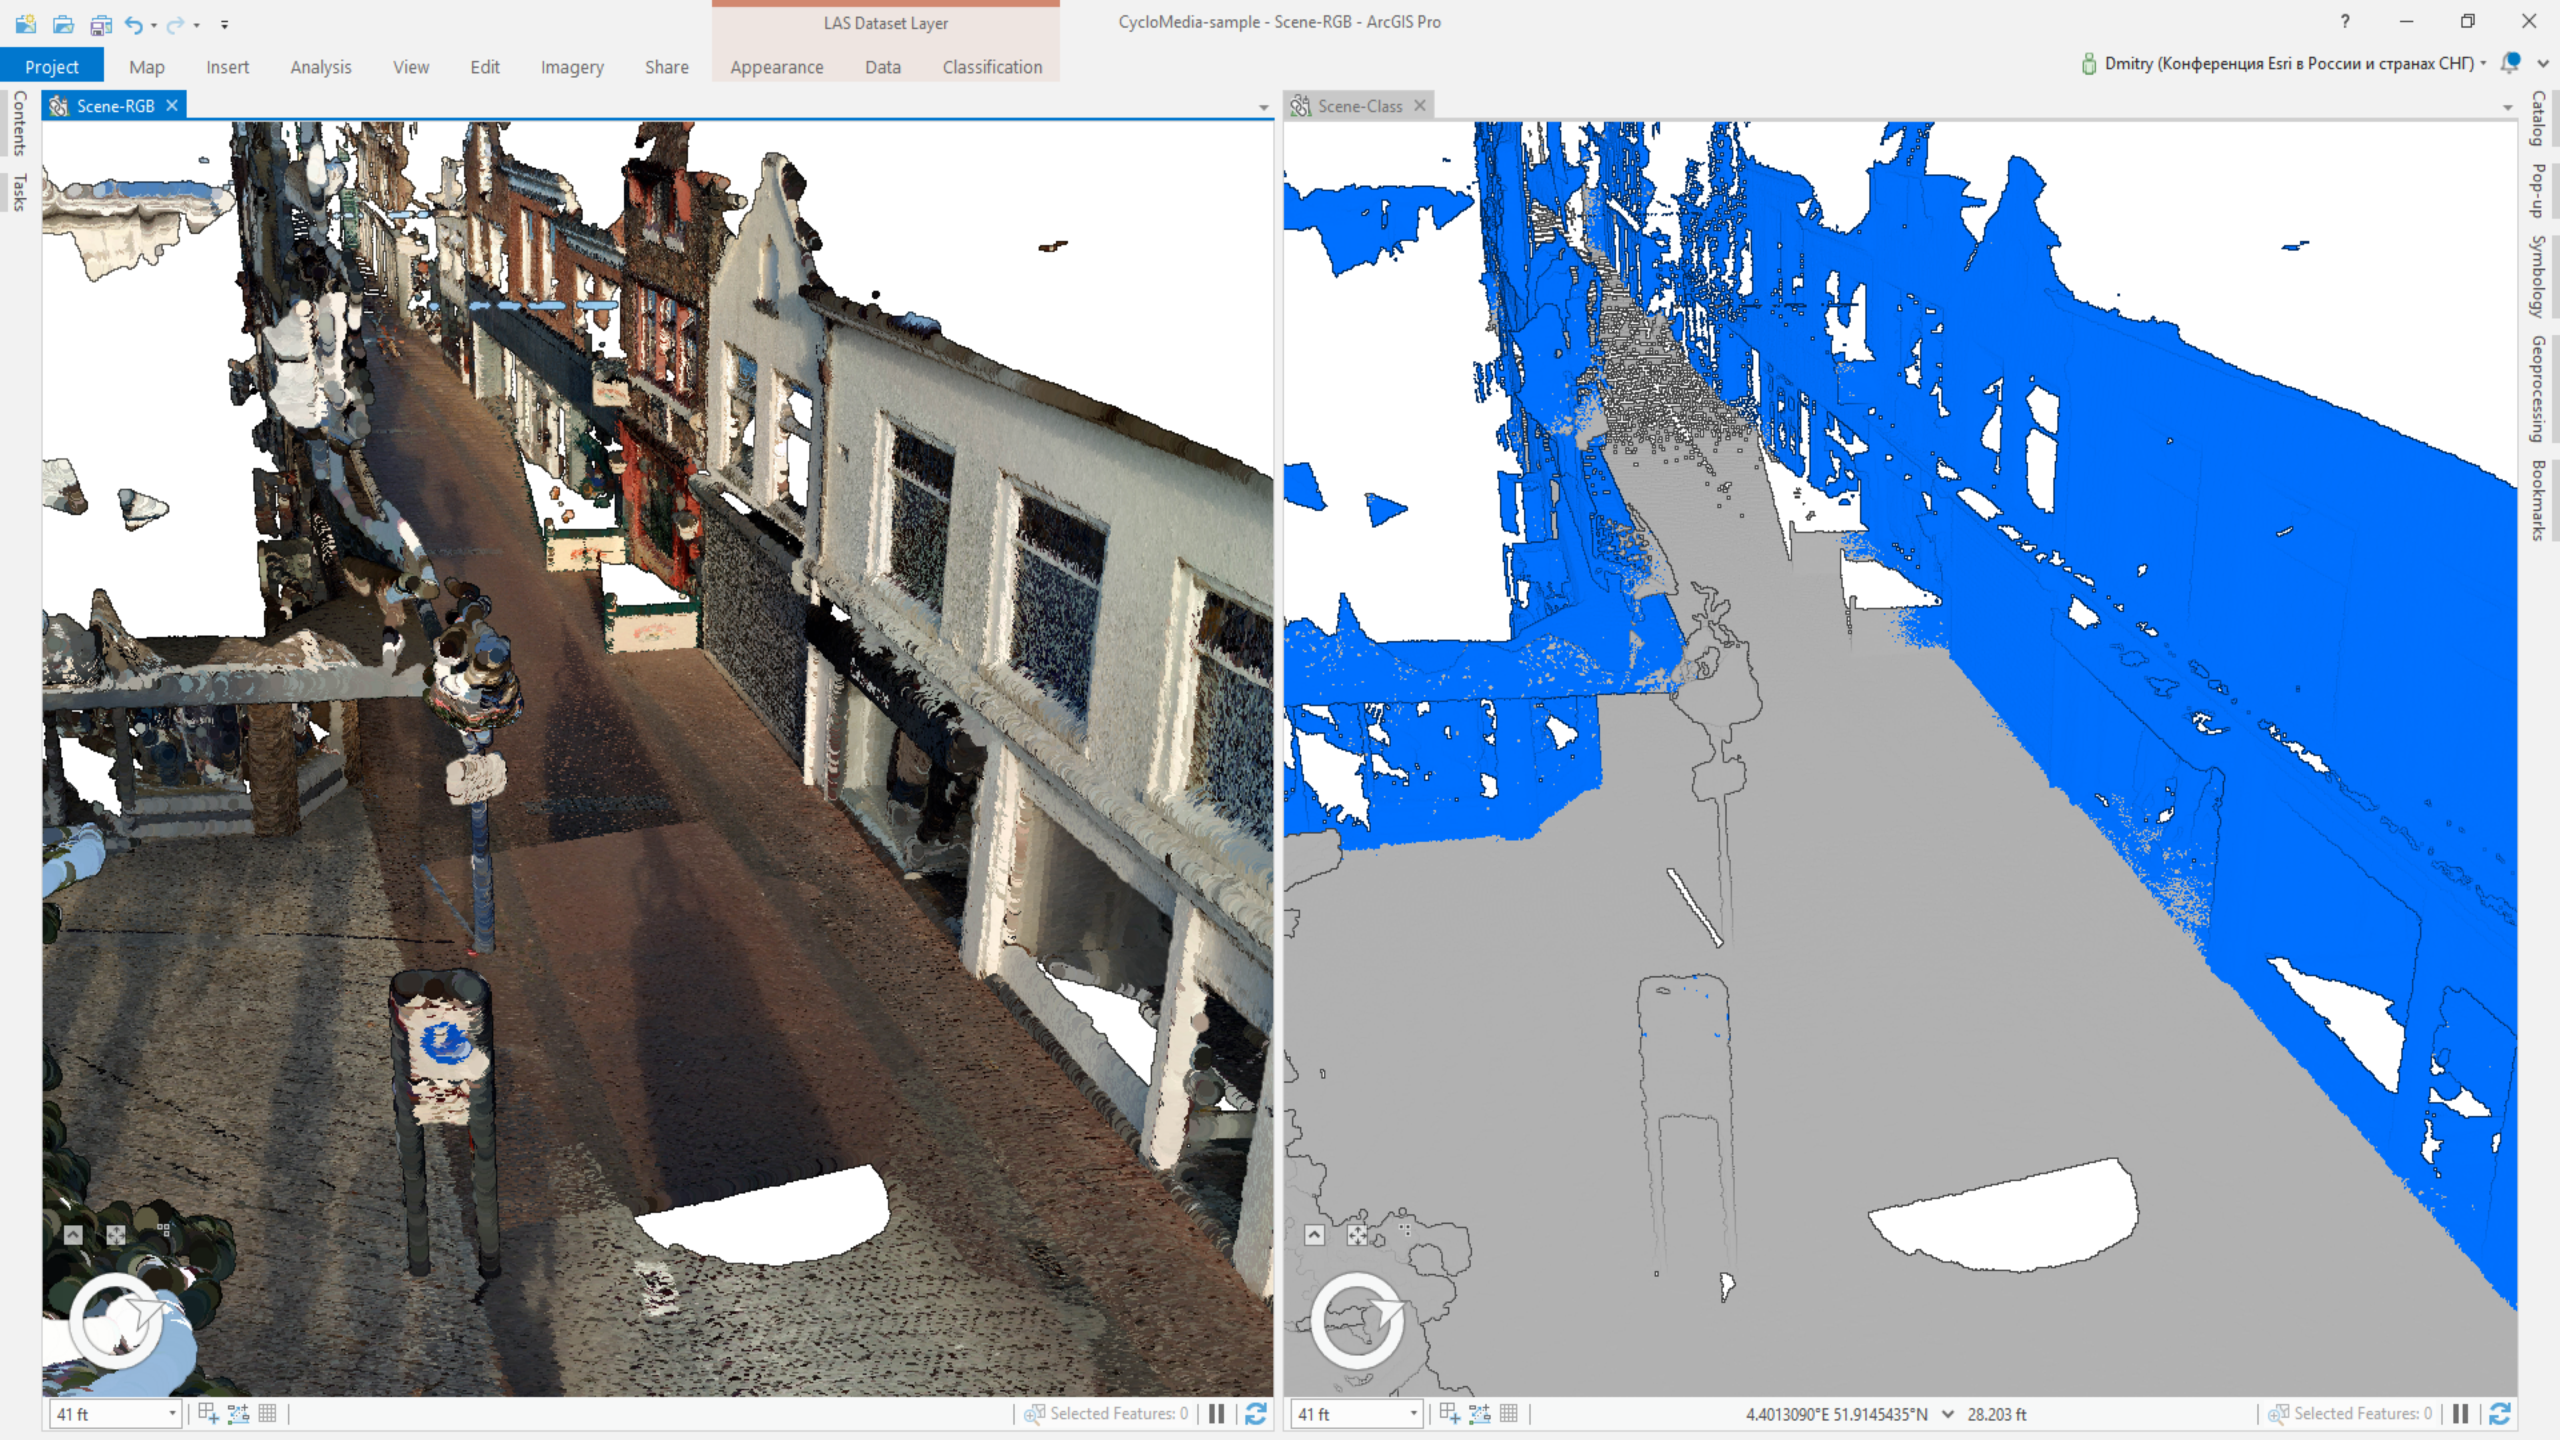
LAS Dataset Layer
CycloMedia-sample - Scene-RGB - ArcGIS Pro
Project
Map
Insert
Analysis
View
Edit
Imagery
Share
Appearance
Data
Classification
Dmitry (Конференция Esri в России и странах СНГ)
Contents
Tasks
Scene-RGB
Scene-Class
Catalog
Pop-up
Symbology
Geoprocessing
Bookmarks
41 ft
Selected Features: 0
41 ft
4.4013090°E 51.9145435°N
28.203 ft
Selected Features: 0

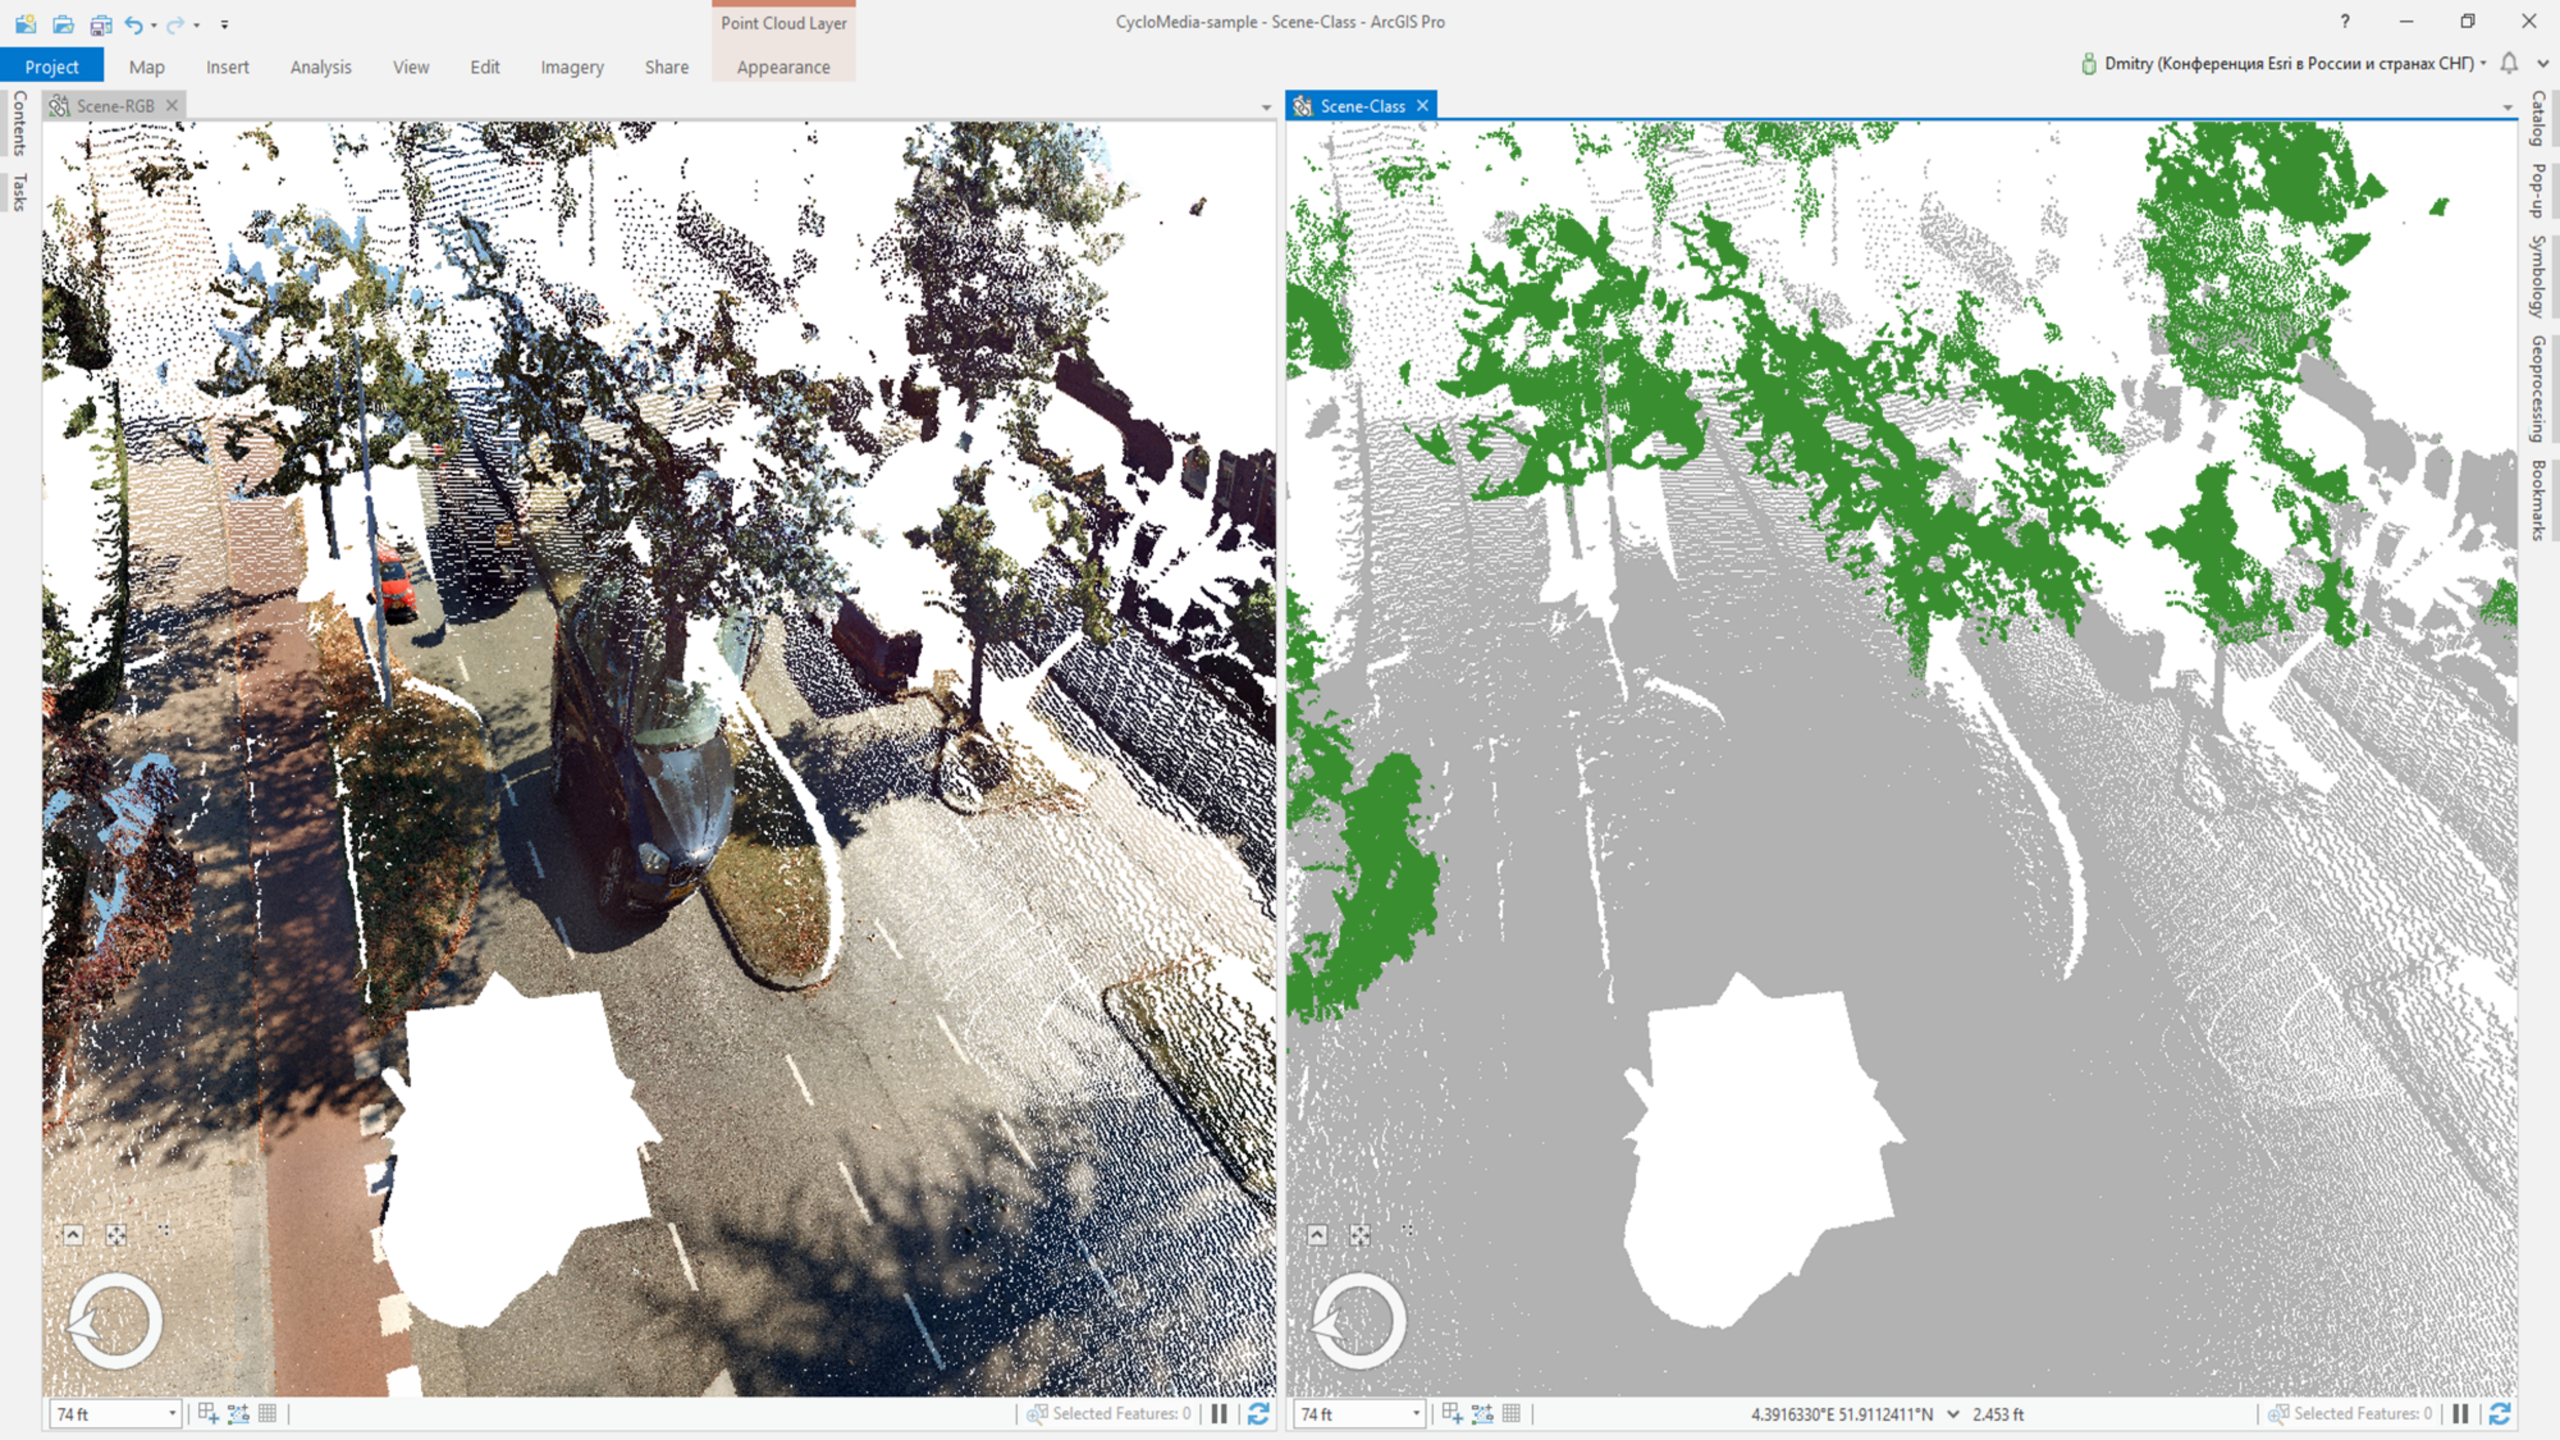
Point Cloud Layer
CycloMedia-sample - Scene-Class - ArcGIS Pro
Project
Map
Insert
Analysis
View
Edit
Imagery
Share
Appearance
Dmitry (Конференция Esri в России и странах СНГ)
Scene-RGB
Scene-Class
Contents
Tasks
Catalog
Pop-up
Symbology
Geoprocessing
Bookmarks
74 ft
Selected Features: 0
74 ft
4.3916330°E 51.9112411°N
2.453 ft
Selected Features: 0


Point Cloud Layer
CycloMedia-sample - Scene-Class - ArcGIS Pro
Project
Map
Insert
Analysis
View
Edit
Imagery
Share
Appearance
Dmitry (Конференция Esri в России и странах СНГ)
Scene-RGB
Scene-Class
Contents
Tasks
Catalog
Pop-up
Symbology
Geoprocessing
Bookmarks
64 ft
Selected Features: 0
64 ft
4.4023086°E 51.9145364°N
10.018 ft
Selected Features: 0

PoC: PointCNN in Mobile Point Clouds

CycloMedia's "CycloRama" sample point cloud with synced 360-imagery

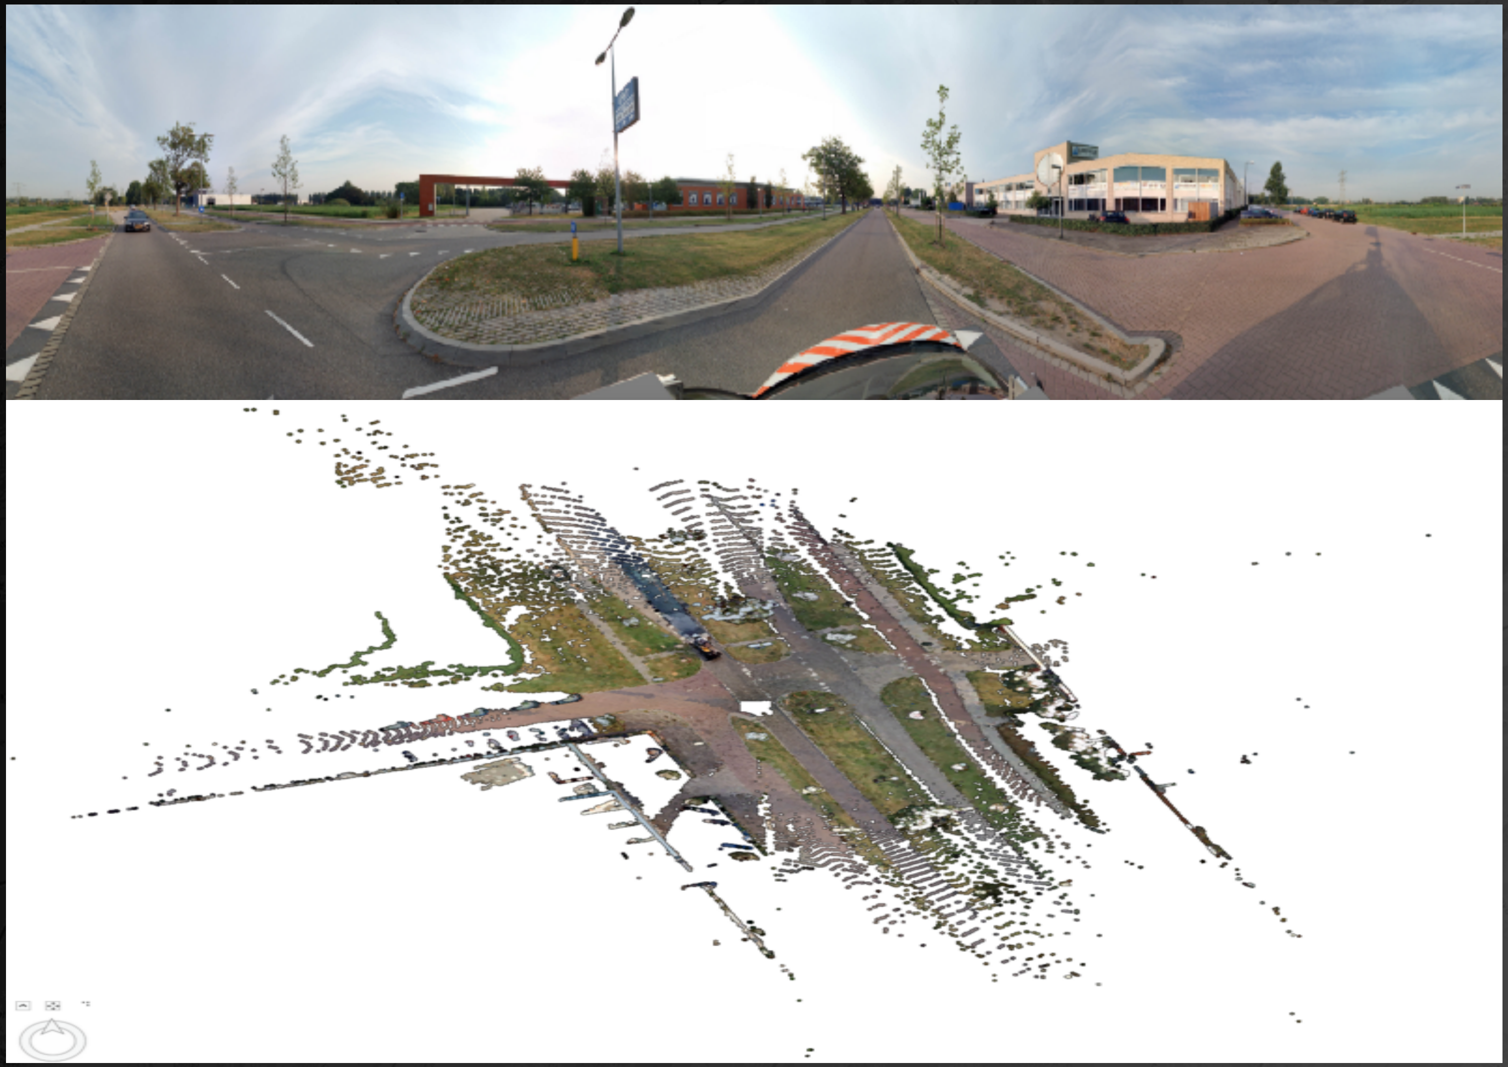

PoC: PointCNN in Mobile Point Clouds

After ~36 hours of training on GV100:

```
@E150K
                OTHER        BUILDING
Precision:[0.99126286 0.96047395],
Recall:   [0.9923109  0.95528204],
F1-score: [0.9917866  0.95787096]
```

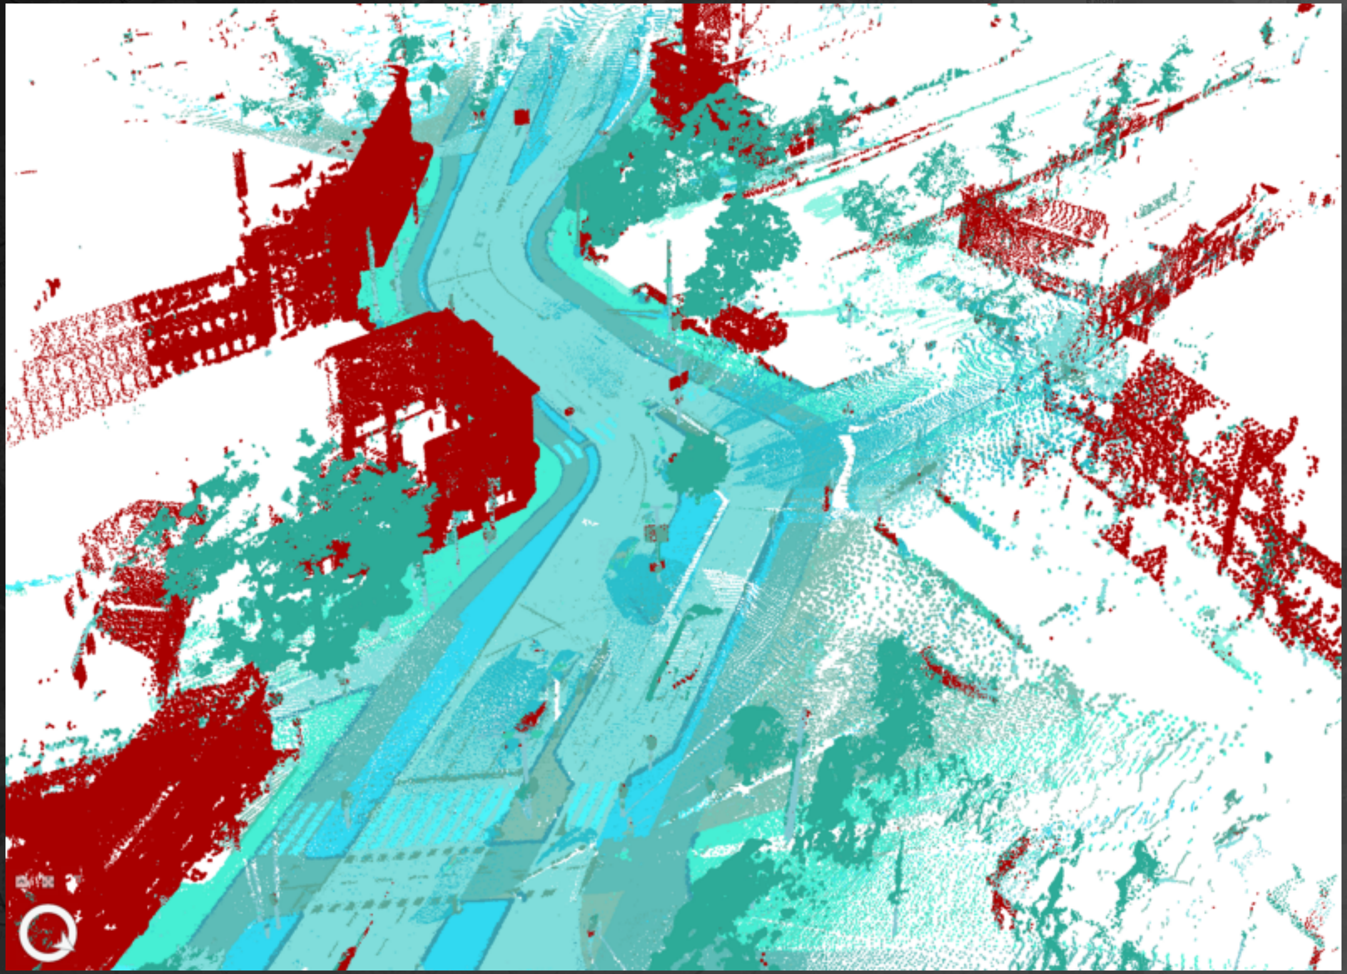

PoC: PointCNN in Mobile Point Clouds

After ~24 hours of training on GV100:

```
@E100K
                OTHER        TREE
Precision:[0.98007521 0.94283636],
Recall:    [0.98718658 0.91327329],
F1-score:  [0.98361804 0.92781940]
```

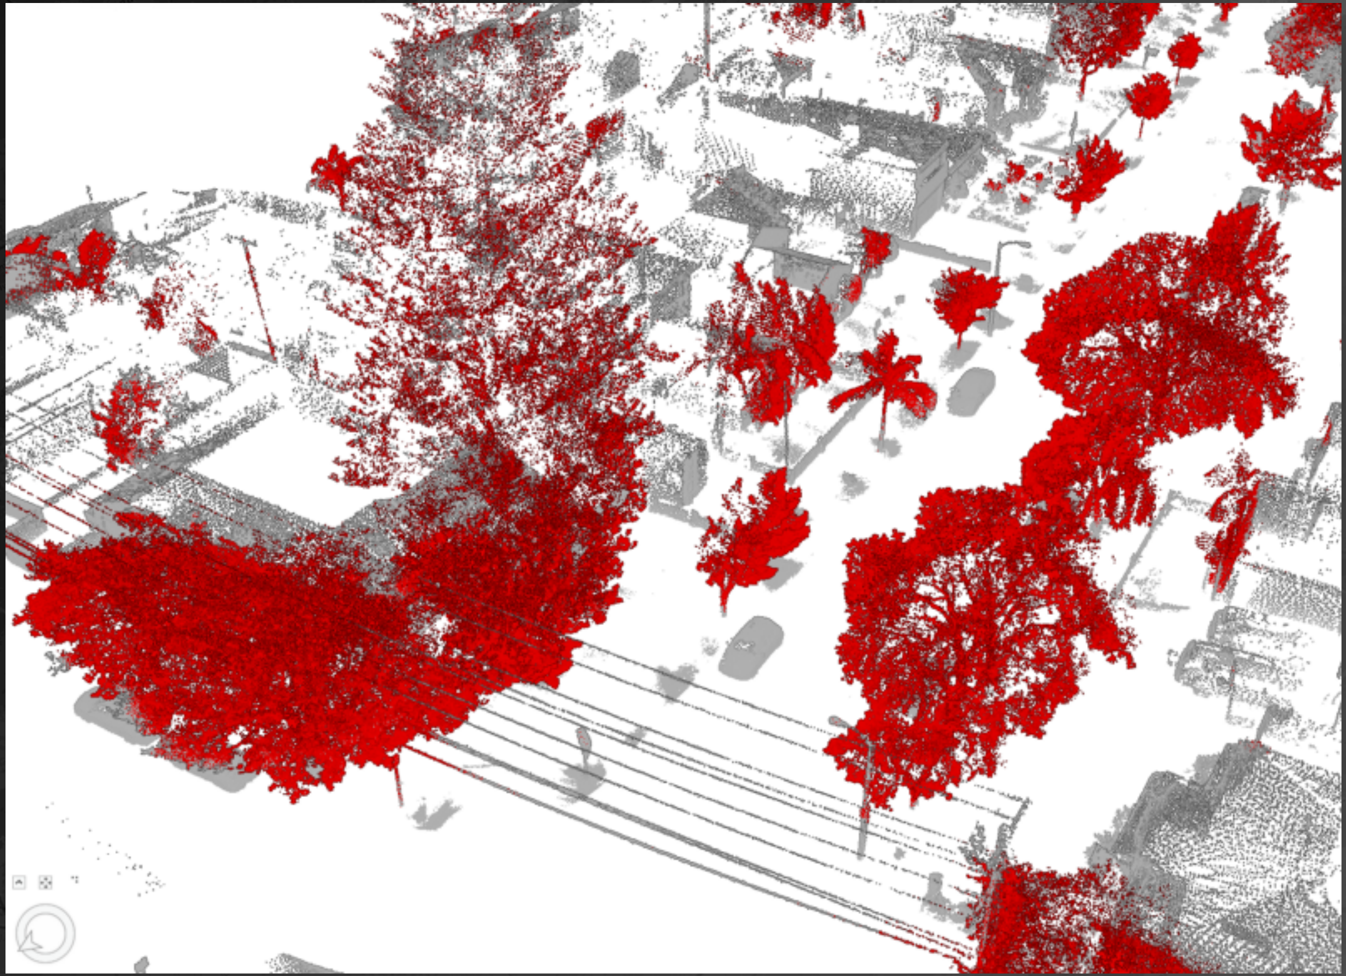

Object extraction from labeled point cloud: Trees

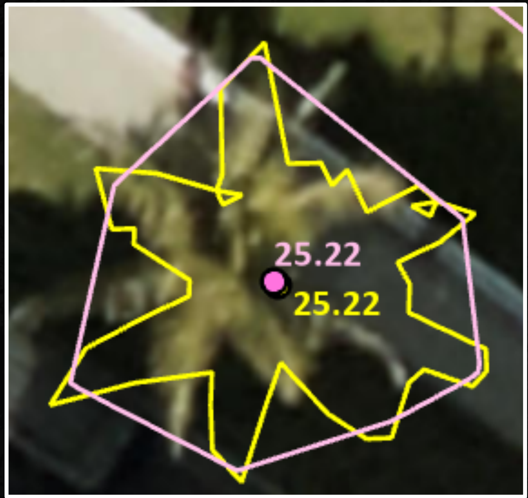

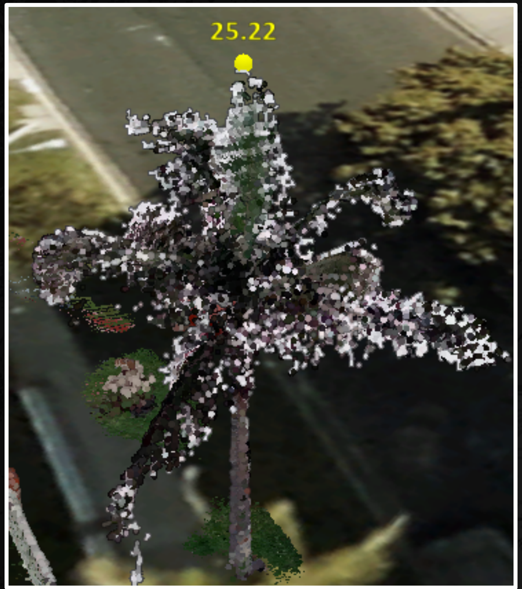

PoC: PointCNN in Mobile Point Clouds

- Traffic Lights are harder:
  - Small number of samples
    ~580,000 points (0.032%)
  - Different types, attachment options
  - Located at most noisy intersections

PoC: PointCNN in Mobile Point Clouds

After ~72 hours of training on GV100:

```
@E300K
              OTHER      TRAFFIC LIGHT
Precision:[0.99974482 0.51987179],
Recall:    [0.99939181 0.72079467],
F1-score: [0.99956828 0.60406375]


Extracted Objects, after DBScan:
DBSCAN 30pts/0.5m
Precision = .655
Recall = .704
F1-Score = .679
```

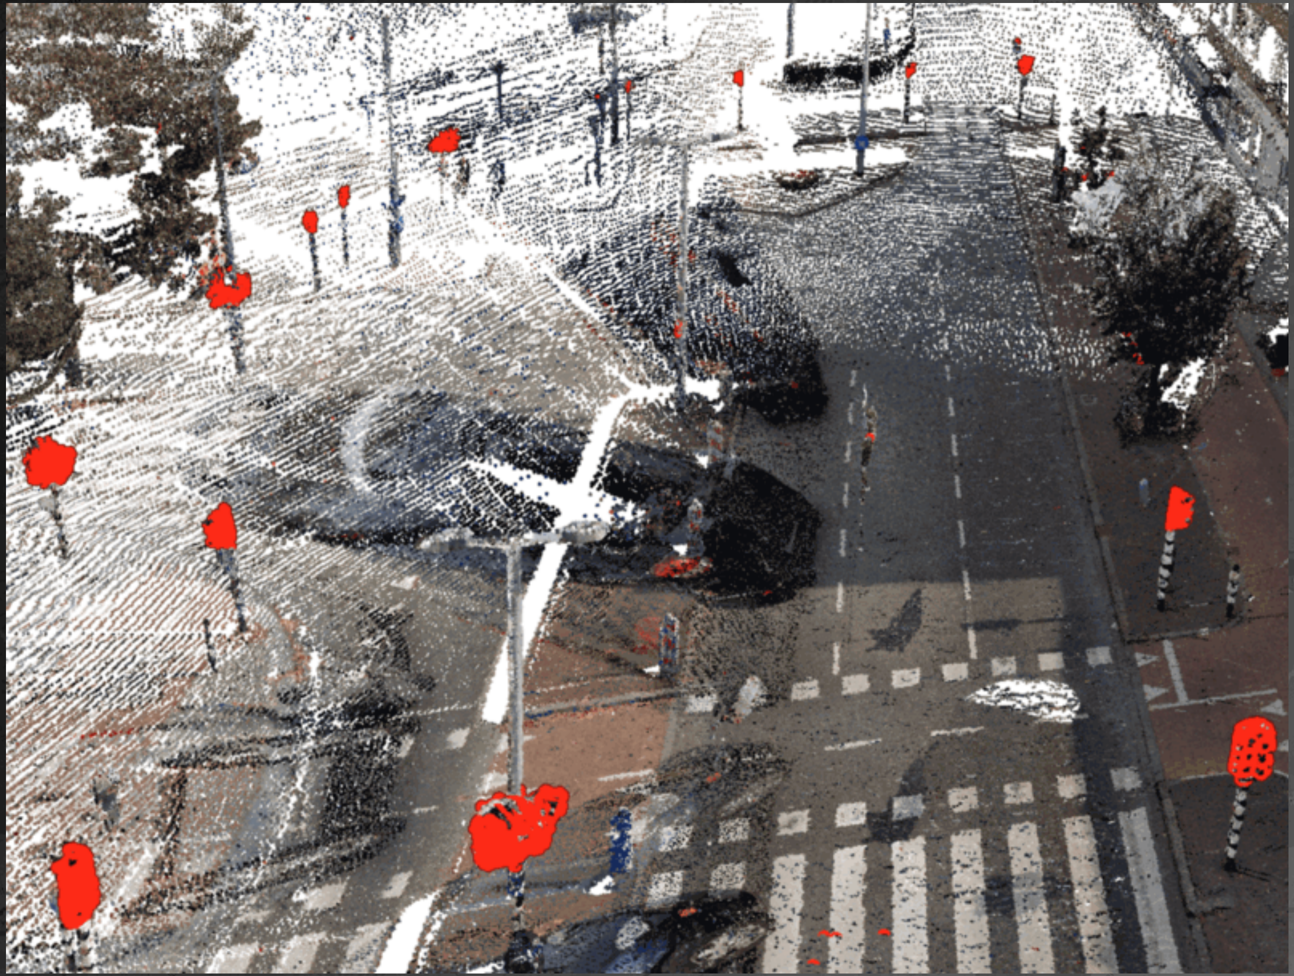

Objects extracted from labeled point cloud:

Traffic lights

- #1 – Traffic light
- #2 – Noise filtered out by DBScan

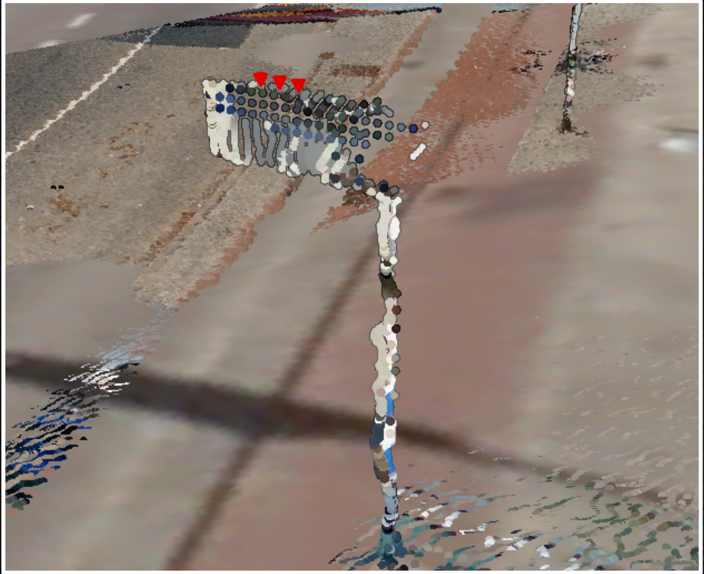

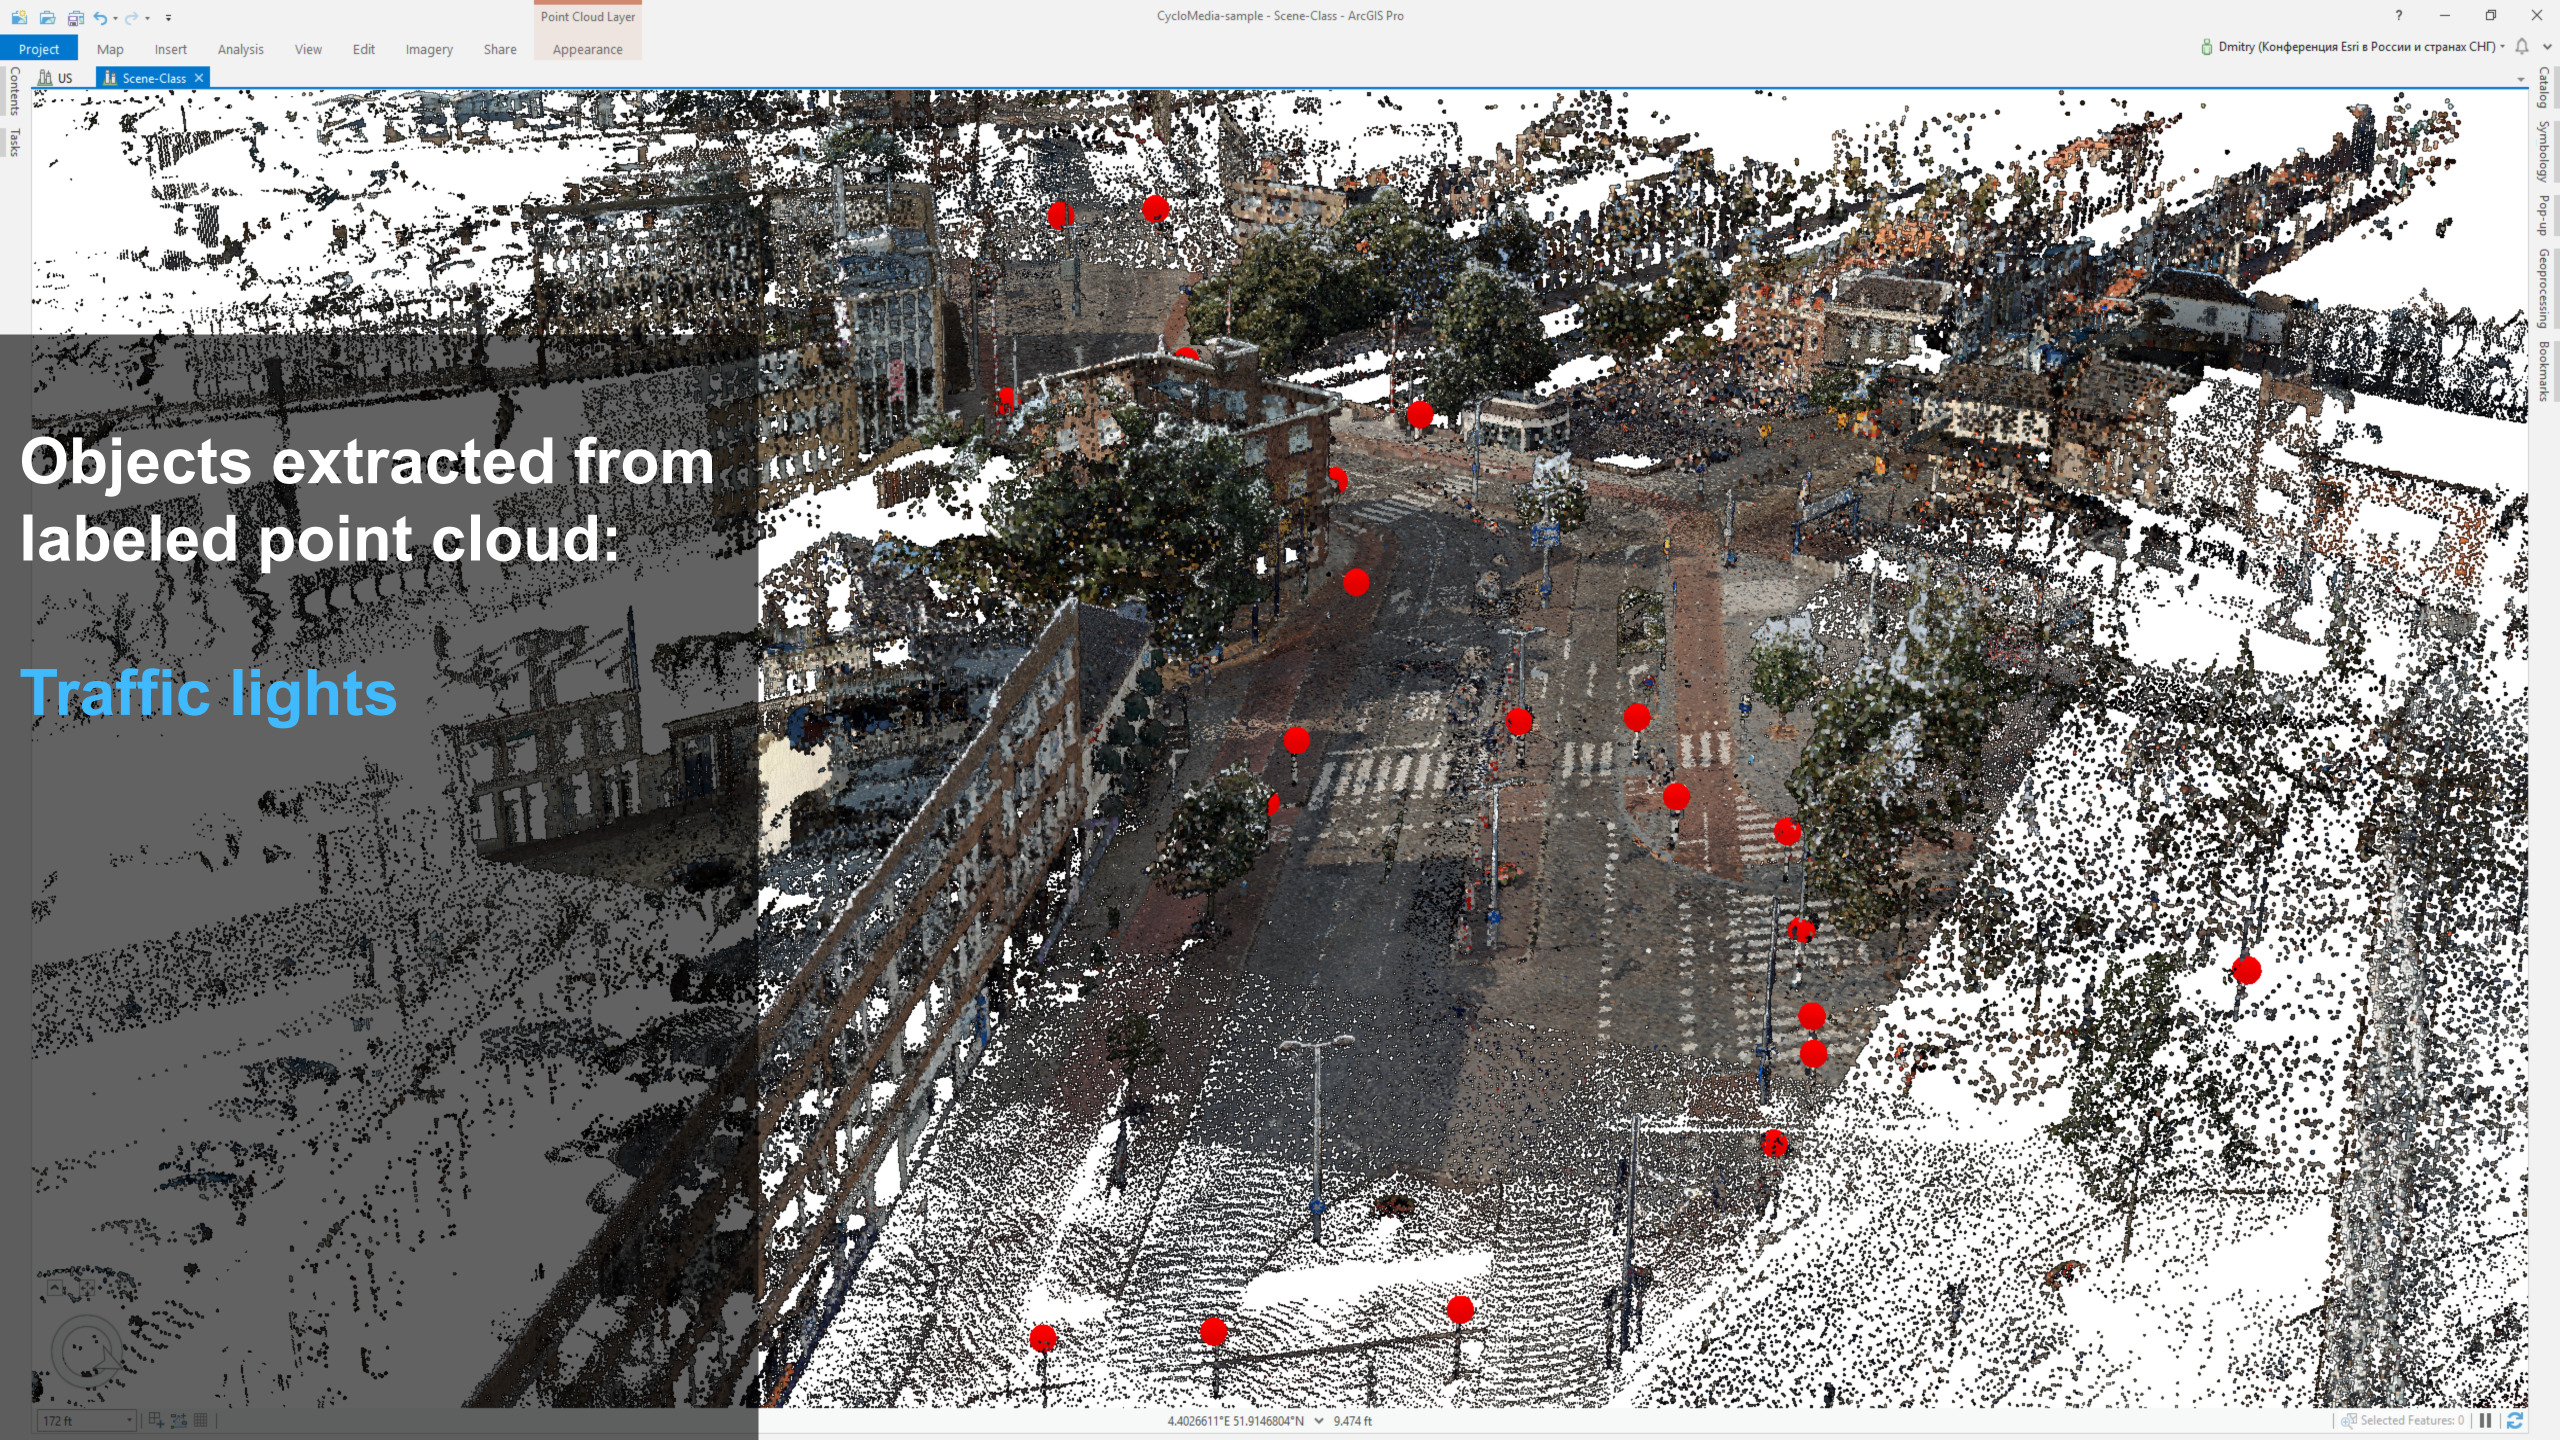
CycloMedia-sample - Scene-Class - ArcGIS Pro
Point Cloud Layer
Project
Map
Insert
Analysis
View
Edit
Imagery
Share
Appearance
Dmitry (Конференция Esri в России и странах СНГ)
US
Scene-Class
Objects extracted from labeled point cloud:
Traffic lights
172 ft
4.4026611°E 51.9146804°N
9.474 ft
Selected Features: 0

Want to know more?

medium.com/geoai

- Reconstructing 3D buildings from aerial LiDAR with AI
- 3D cities: Deep Learning in three-dimensional space
- PointCNN: replacing 50,000 man hours with AI
- Object extraction from Mobile LiDAR point clouds with Machine Learning
- …and much more

Q&A

AMLD 2020
Dmitry Kudinov, dkudinov@esri.com
Sr. Principal Data Scientist, Esri