Improving the Detection of Explosive Hazards with LIDAR-Based Ground Plane Estimation

Andrew Buck, James M. Keller, and Mihail Popescu

University of Missouri
Columbia, Missouri, USA

SPIE Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXI

This work was supported by Army Research Office grant number 57940-EV to support the U. S. Army RDECOM CERDEC NVESD.

Outline

- LIDAR Point Clouds
- Individual Sensing Modalities
  - Forward Looking Ground Penetrating Radar
  - Synthetic Aperture Acoustic
  - LIDAR Magnitude Response
- Ground Plane Estimation
  - Target Prediction Filtering
- Sensor Fusion

LIDAR

- Laser-based range finder
- Creates massive point clouds (195 million points)
- How can LIDAR help with target detection?

Ambient Occlusion Filter

Magnitude Response

Traditional Workflow

- Used with FLGPR and SAA to produce a set of hit locations
- A size-contrast filter is applied to each beamformed image
- Resulting peaks are used as hits
- Use mean-shift to combine hits from multiple frames

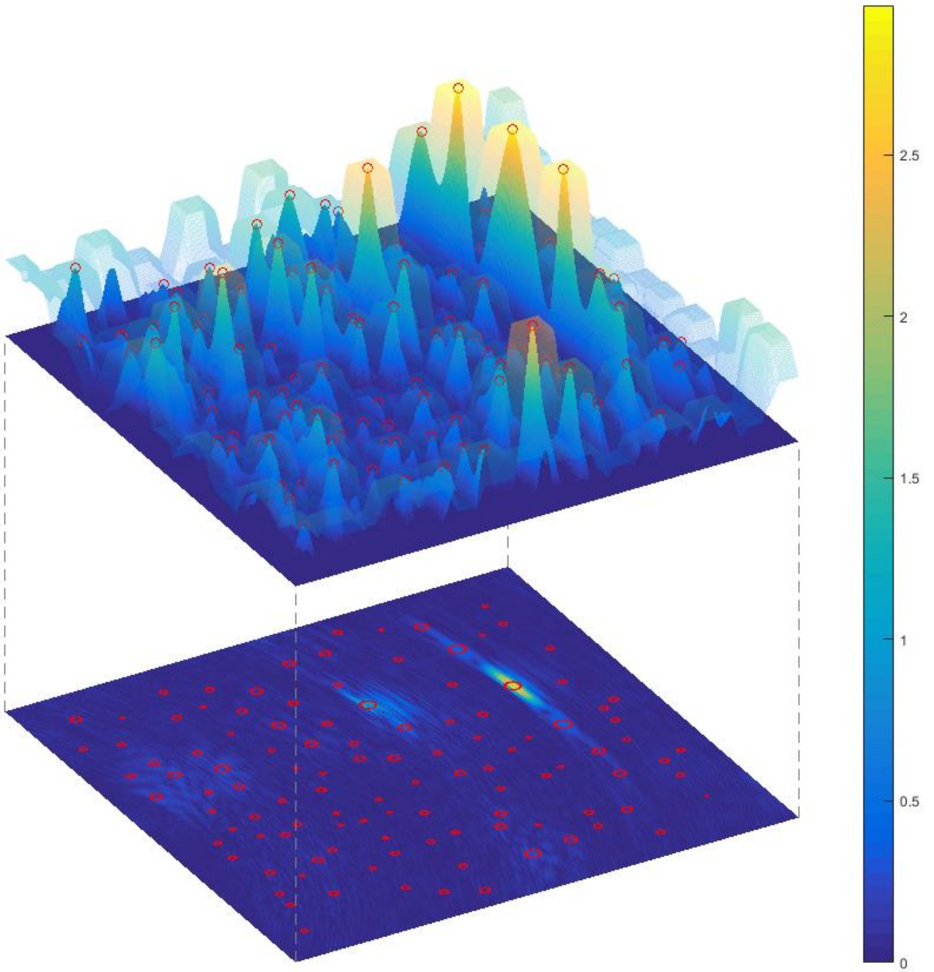

LIDAR Magnitude

- LIDAR magnitude response is much lower when reflecting off of targets
- Threshold all points above 0.2
- Use mean-shift to cluster the remaining points
- Cluster size reflects density and is used as hit confidence

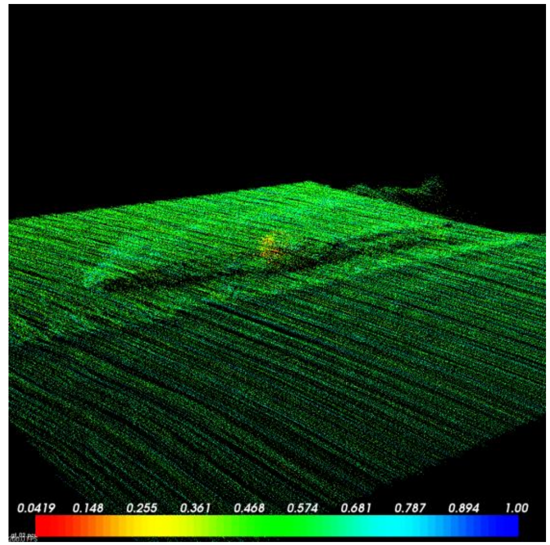

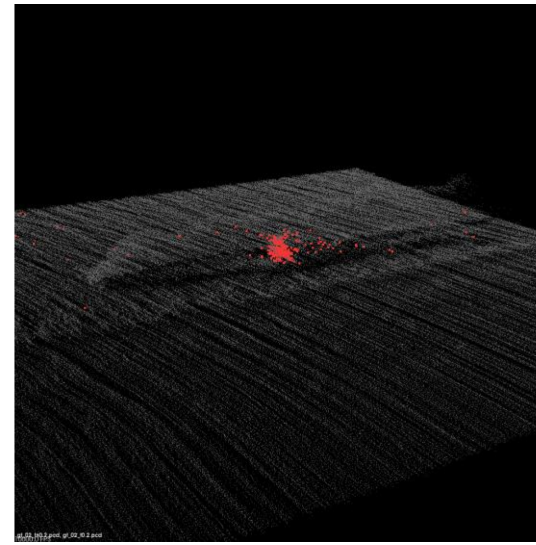

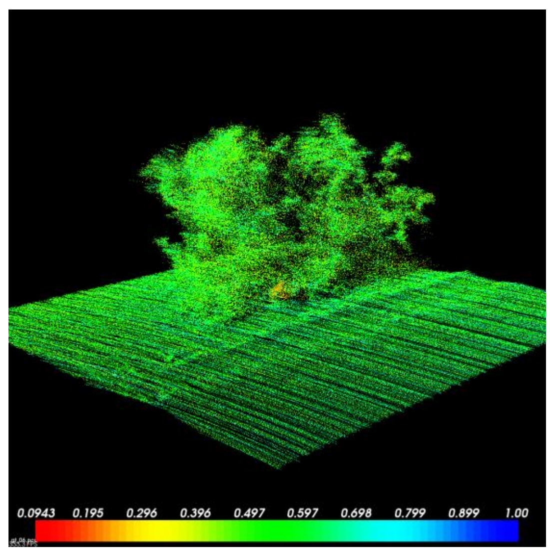

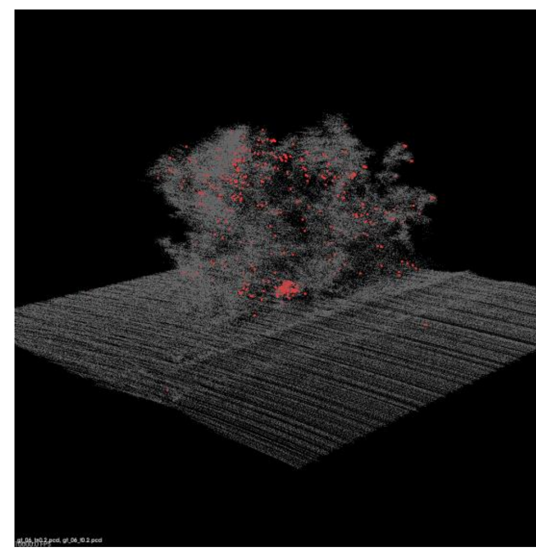

Individual Modality Scores

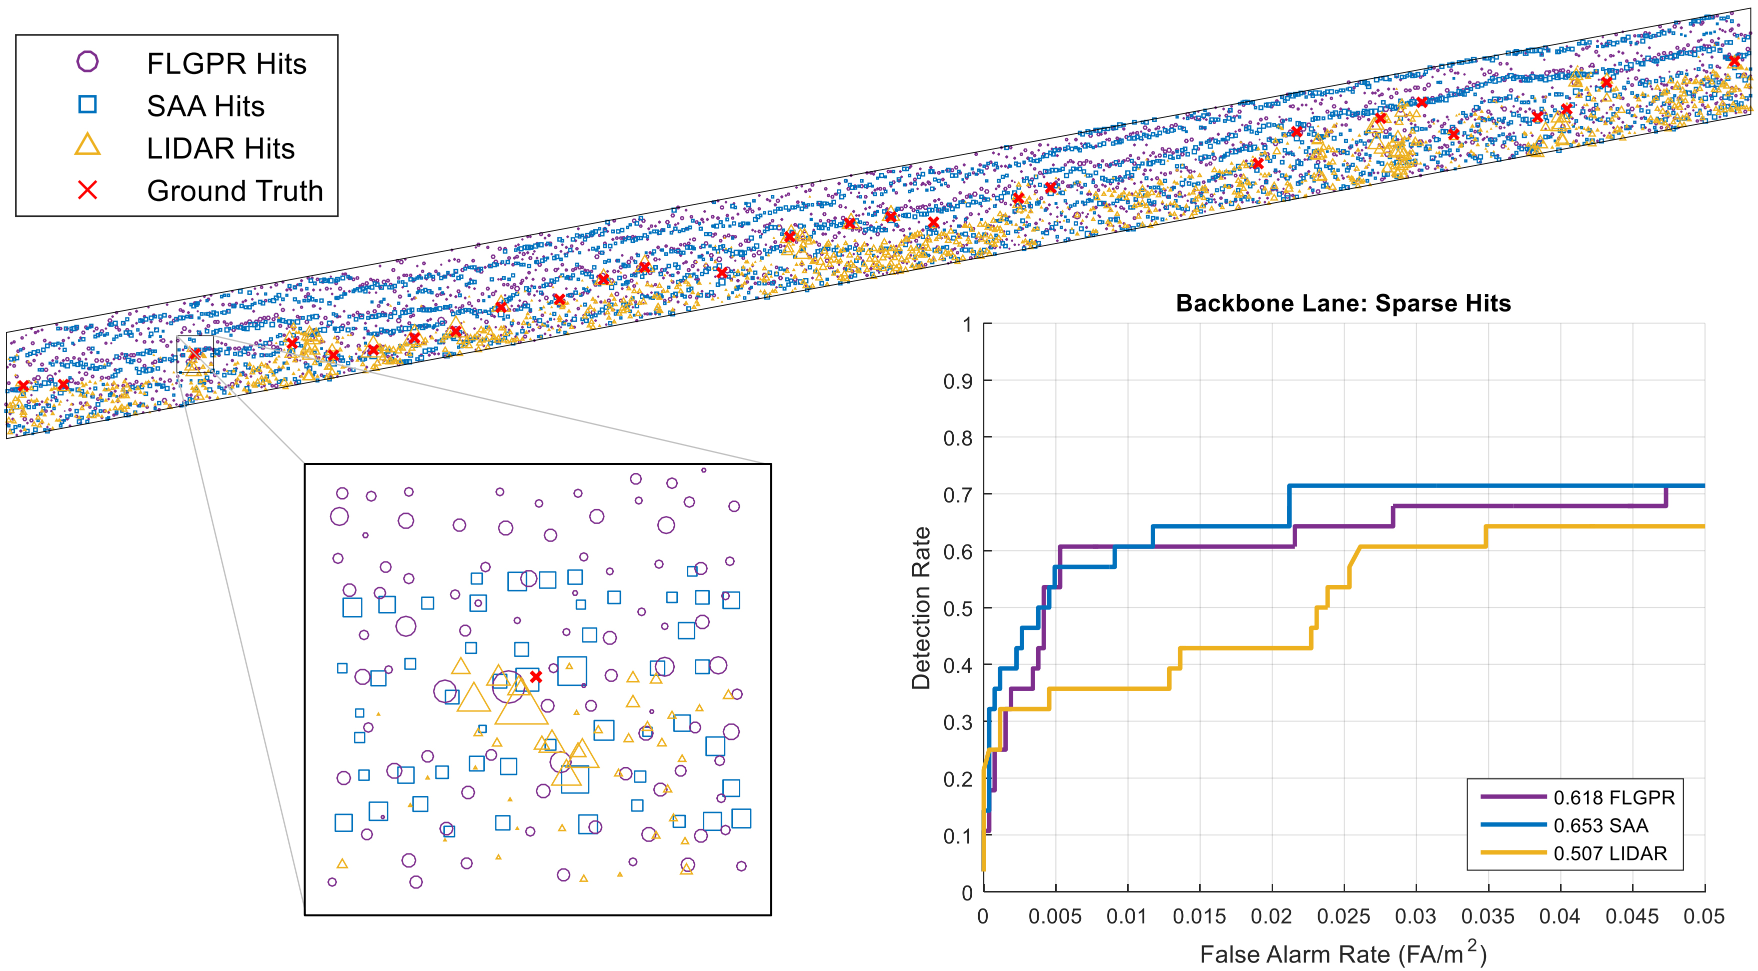

Hand-Labeled "Ground" Truth

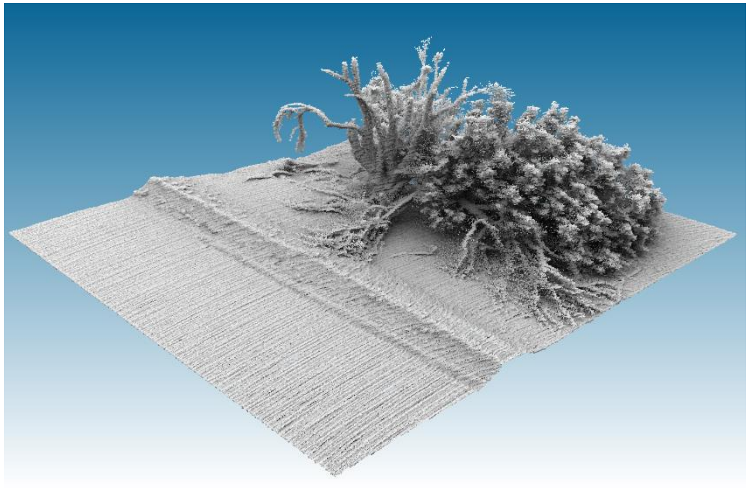

- Train a NN classifier to learn what is ground and what is not
- Hand-labeled 10m x 10m LIDAR patch with approx. 1.5 million points
- Region is divided into 10cm x 10cm grid cells

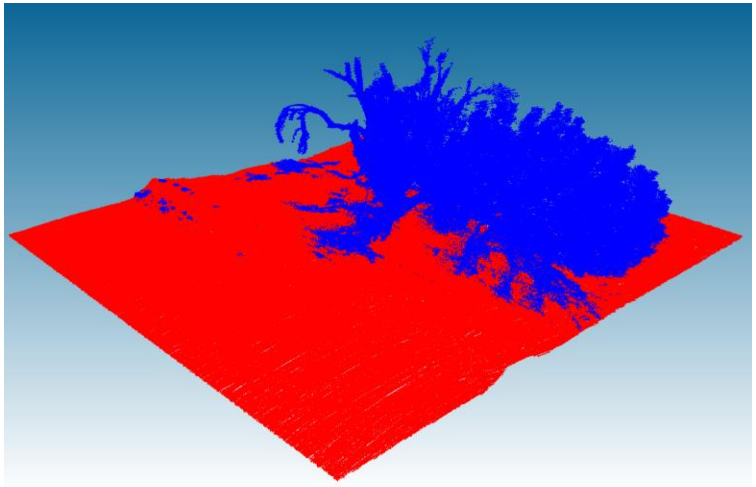

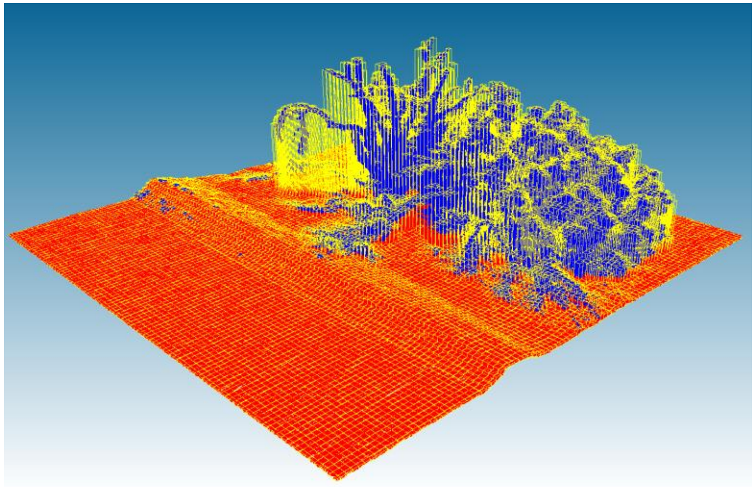

Minimum Point Raster Image

- Save the point with the lowest elevation in each grid column

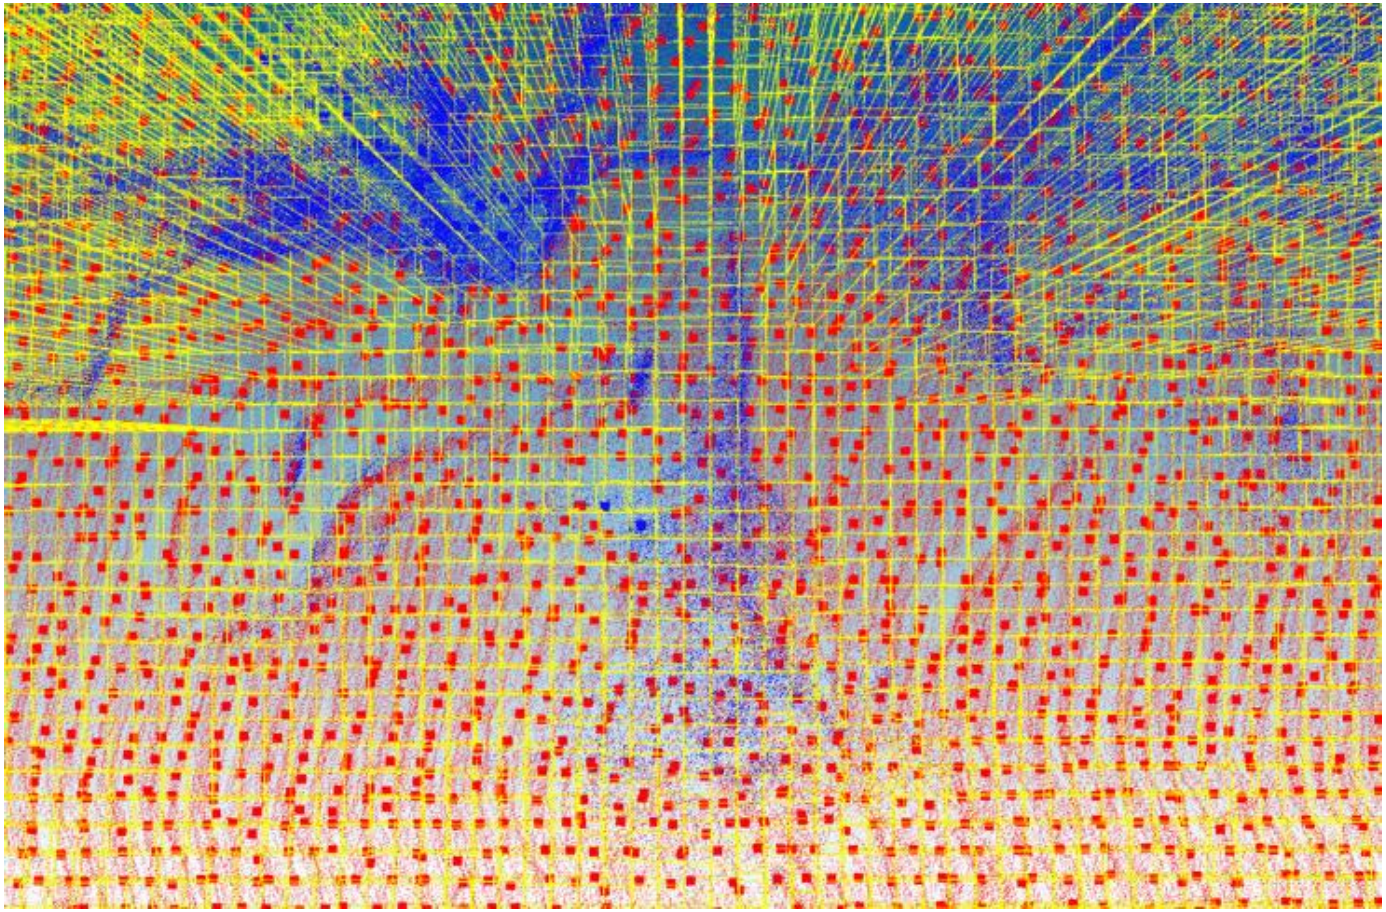

Feature Extraction

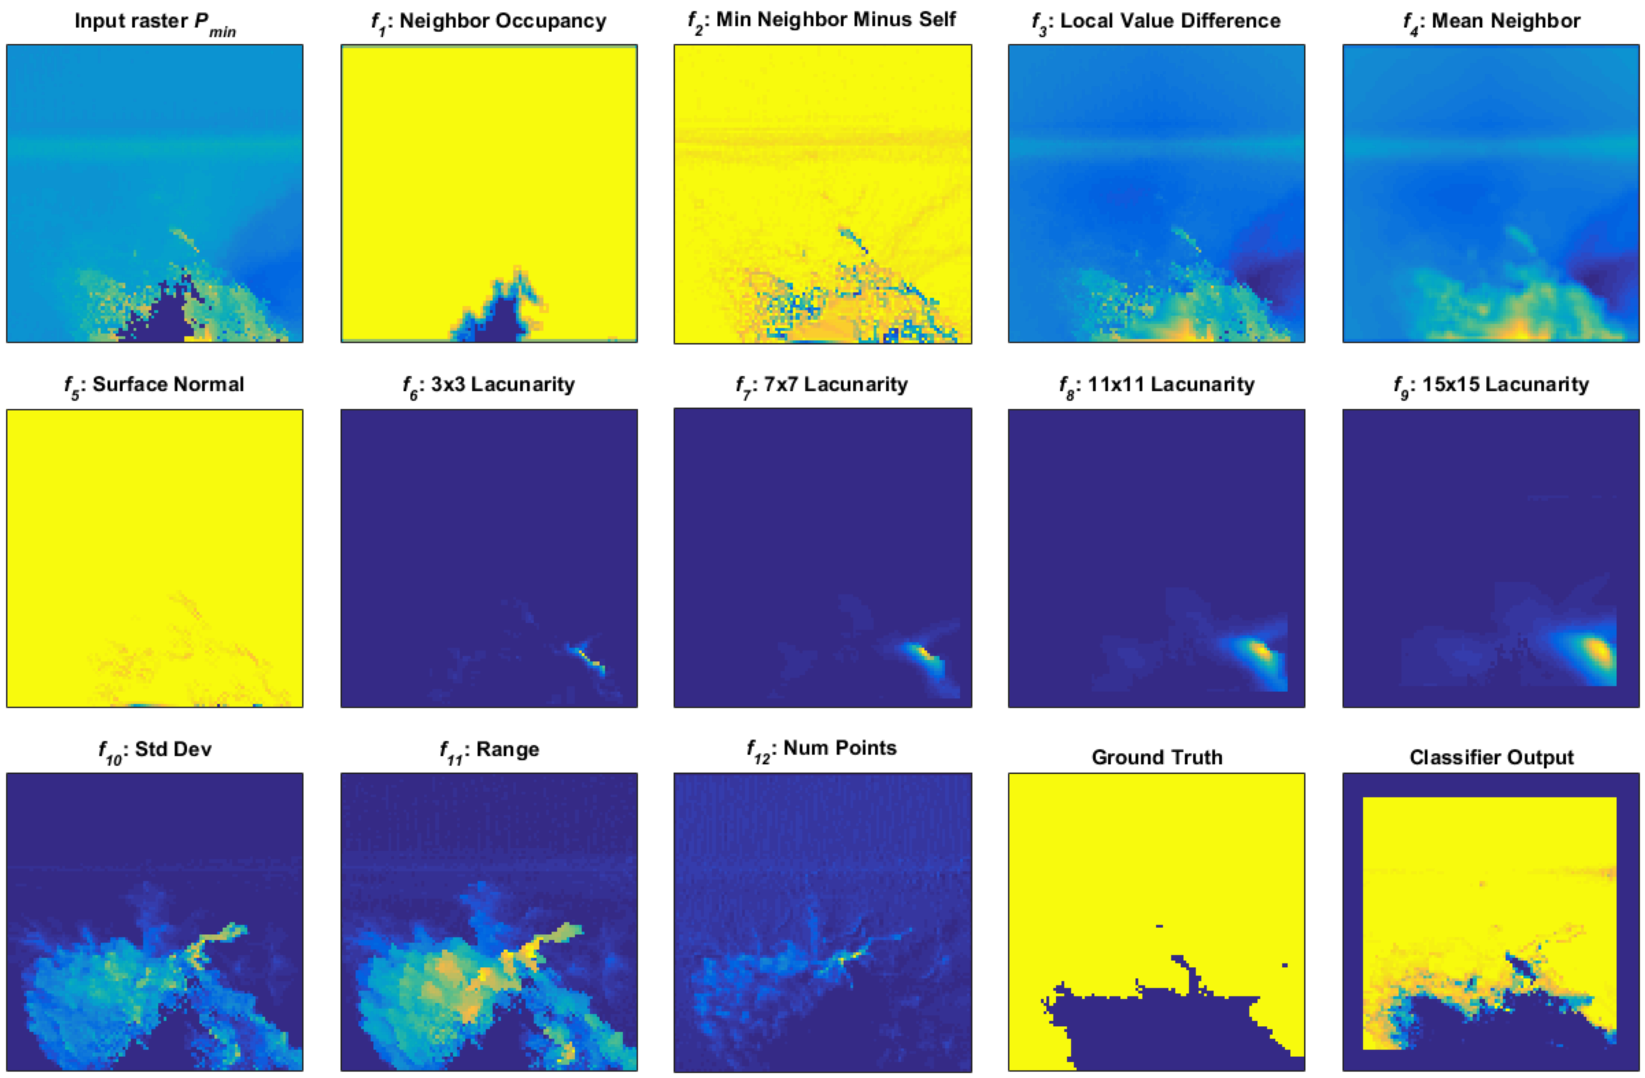

Ground Plane Interpolation

- Threshold and erode high-confidence ground locations from the classifier output
- Use this mask to initialize the ground plane model from the minimum point raster image
- Fill in remaining locations with linear interpolation

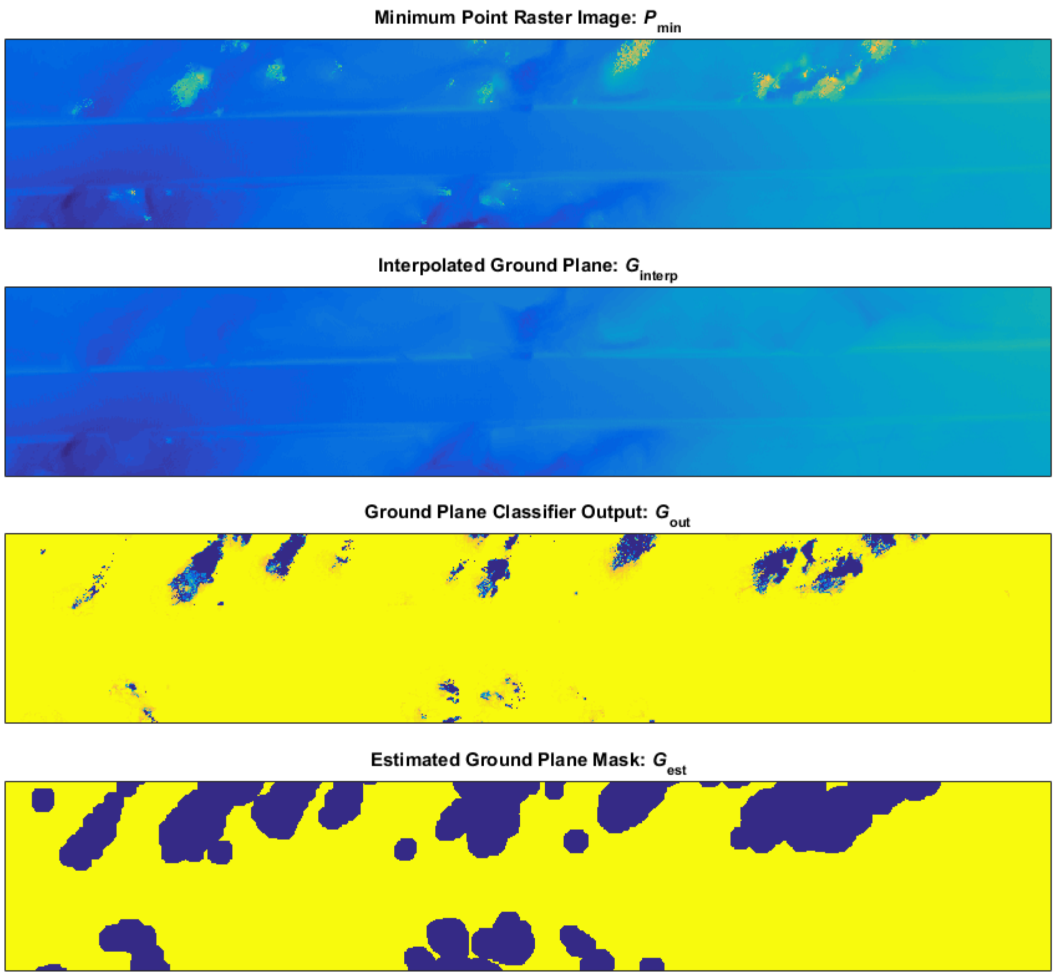

Height Range Thresholding

- Threshold the region based on the difference between the ground plane and maximum point raster image
- Remove hits that could not come from above-ground

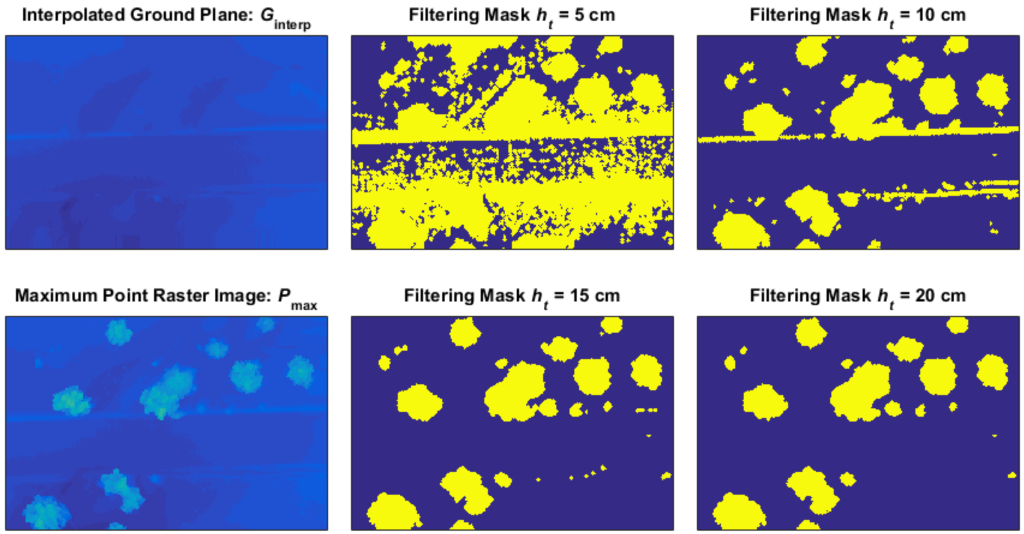

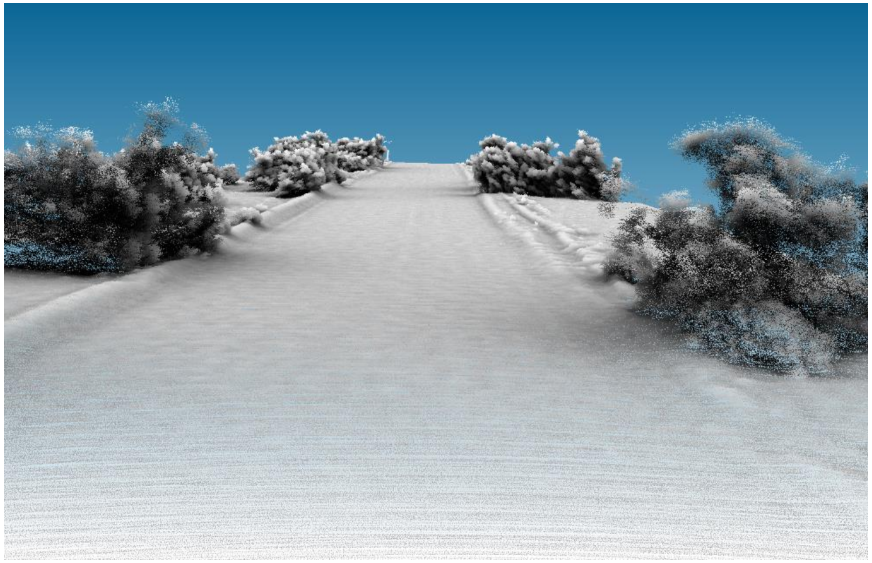

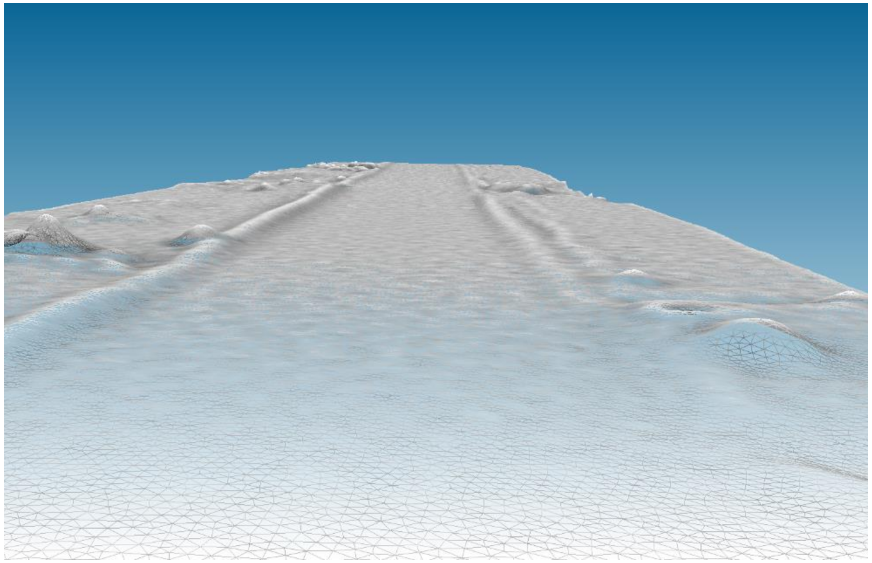

Filtering Results

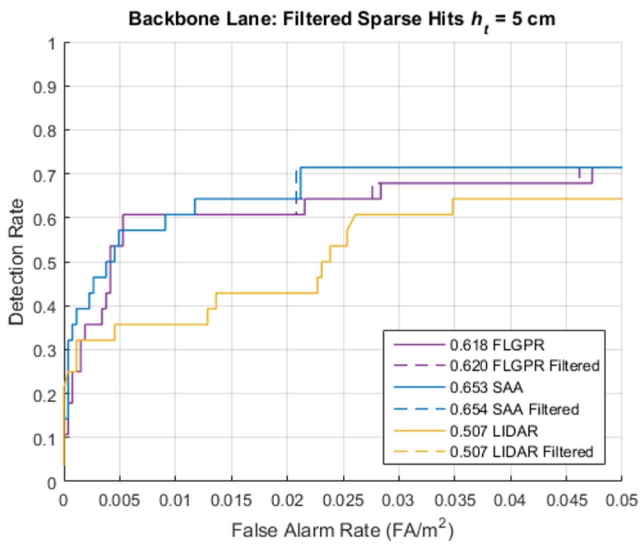

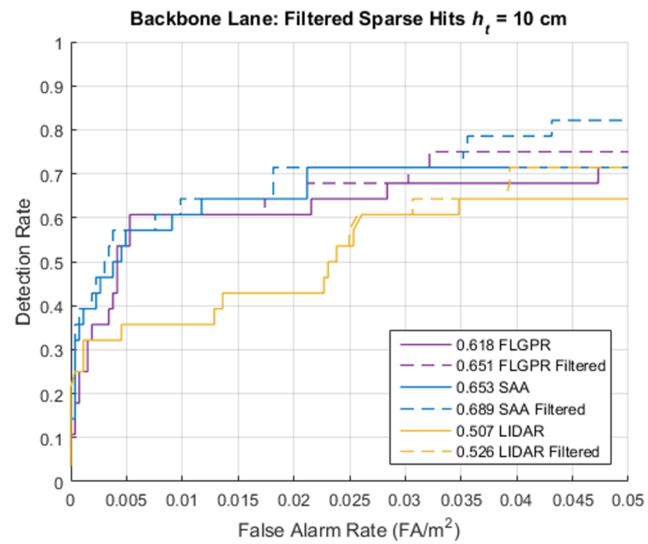

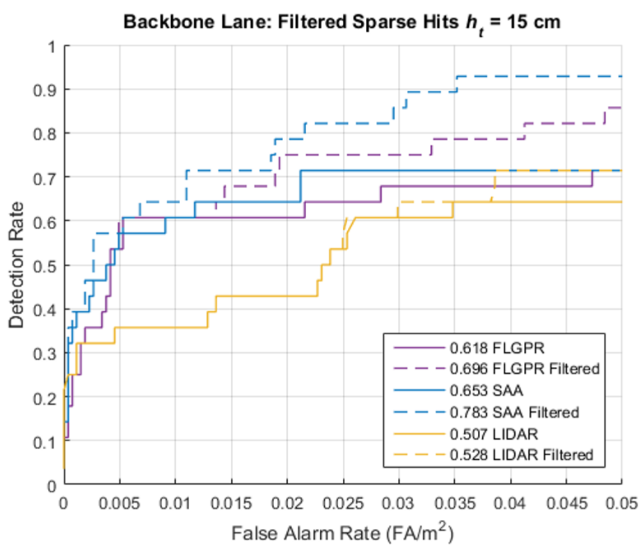

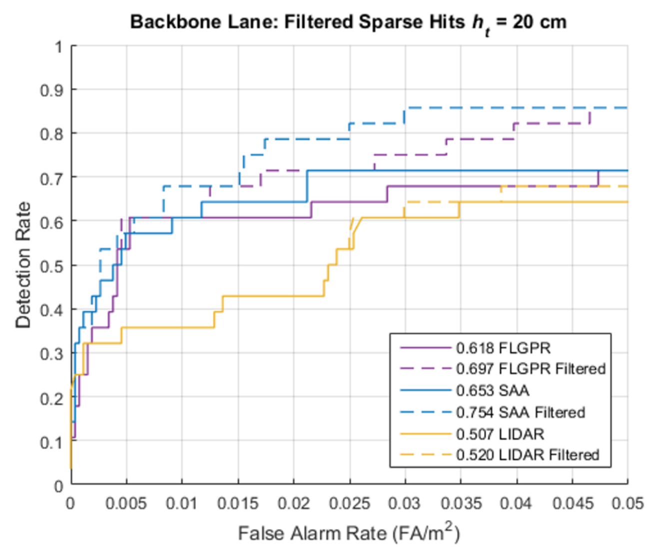

Dense Confidence Surface

- Instead of returning a sparse set of hit locations, generate a dense confidence surface from the output of the size-contrast filter
- Confidence points are aggregated like LIDAR points
- Each sensing modality produces a confidence surface raster image
- Normalize each surface by dividing by the mean

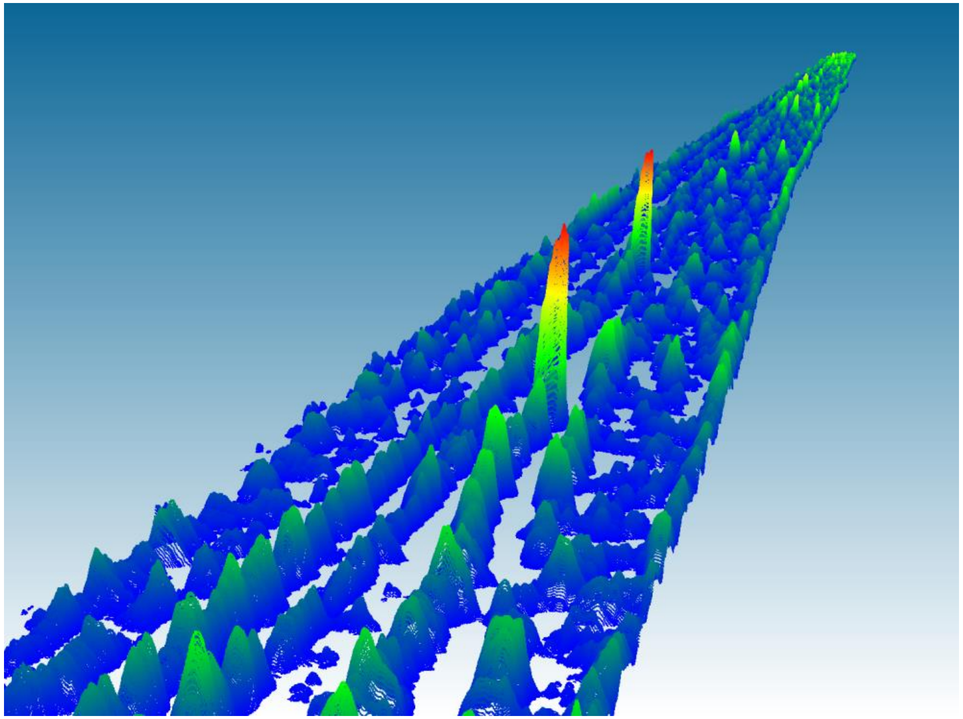

Sensor Fusion

- Surfaces are fused with a weighted average
- Non-maximum suppression is used to generate a hit list for scoring

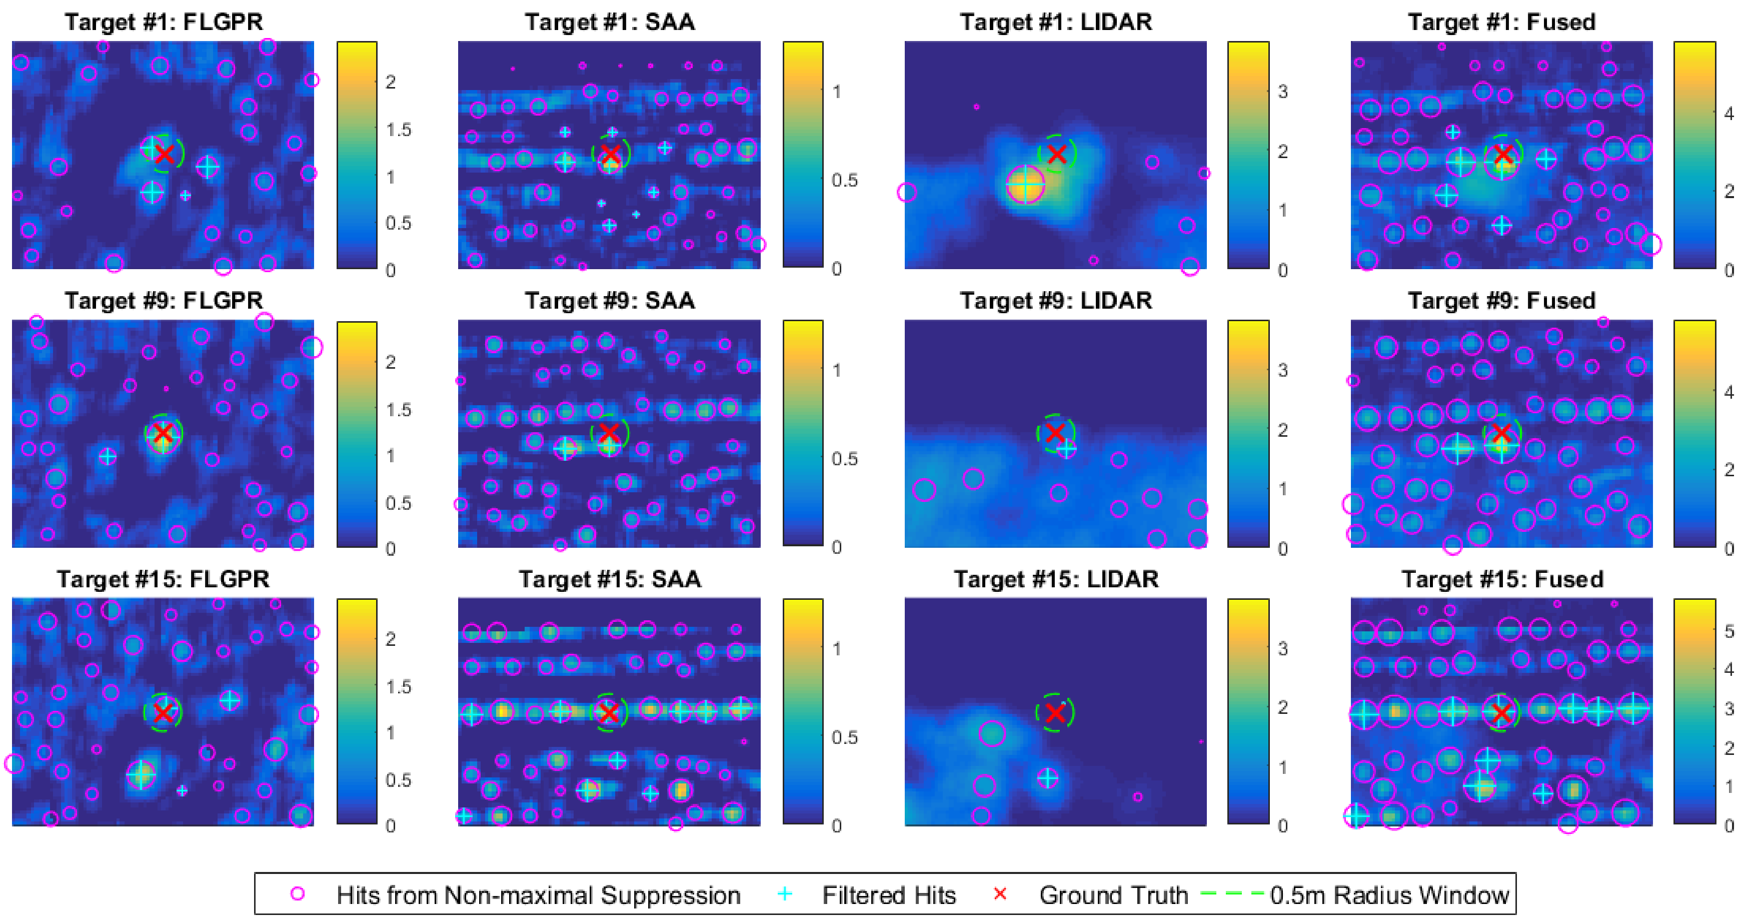

Sensor Fusion Results

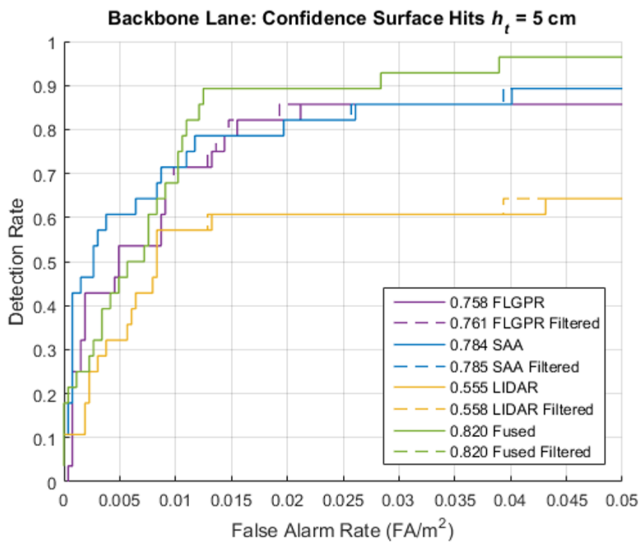

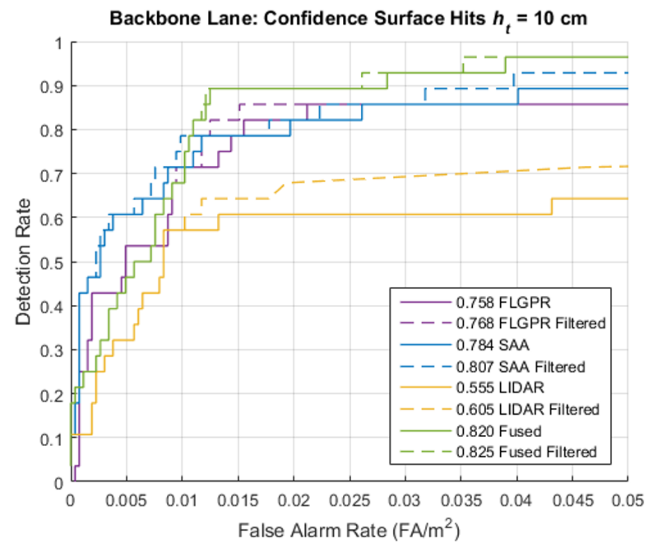

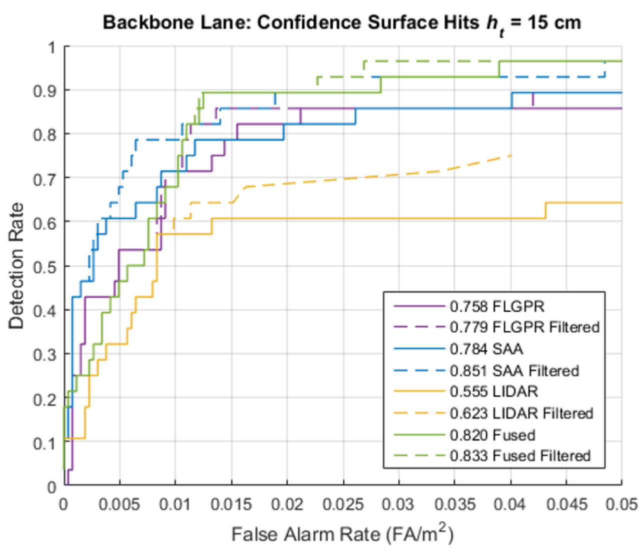

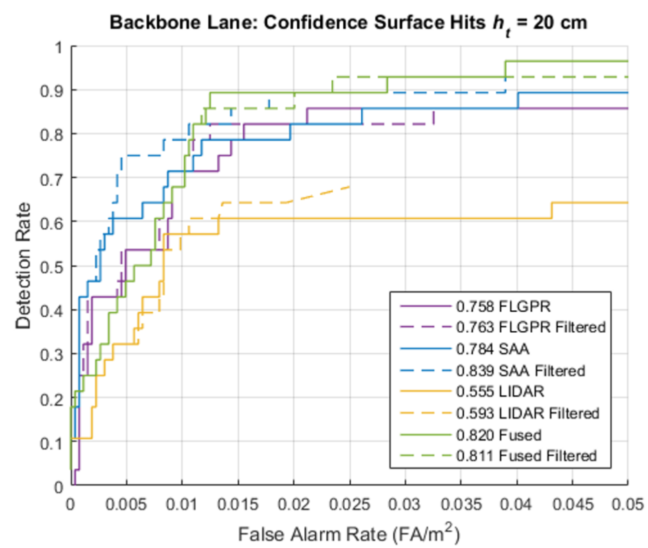

Conclusion

- LIDAR can augment existing workflows
  - Independent sensing modality
  - Improved scene understanding
- Ground plane estimation
  - Filter prediction locations based on environment
- Sensor Fusion
  - Create a confidence surface for each modality
  - Generate predictions from the combined surface