Building an Autonomous Detect-And-Avoid System for Commercial Drones

Alejandro Galindo
VP of Research and Development
Iris Automation Inc

2021 embedded VISION summit
VIRTUAL | MAY 25-28

Iris Automation
The future in sight.

Outline

- Introduction & context
- Requirements
- Algorithmic approaches
- False positive rates
- Conclusions

Commercial and Industrial Drones

Commercial and Industrial drones have the potential to completely disrupt industries and create new ones.

- Infrastructure inspection
- Package delivery
- Agriculture
- Surveying
- Search and rescue
- 100s of other applications

Source: Censys Technologies, https://censystech.com/

The Main Challenge

Goal: Make industrial drones safe to integrate into the national airspace

Technical challenge: Detect, in real time, using a camera, crewed intruders about a kilometer+ away

The Challenge

Requirements

Main Requirements

- Detect intruder
- Track intruder
- Estimate location and velocity of intruder
- Alert/Maneuver if other aircraft is detected

Considerations

- Low Size Weight And Power
- Real time, no connectivity
- High recall rate (90%+)
- Low false positive rate (1 every 10 hours)
- No un-expected outcomes

Algorithmic Approaches

How to Detect Intruders in the Field of View

Potential approaches

- Deep learning
  - Frame-based object detection
  - Video-based object detection
- Conventional Computer Vision
  - Optical flow (dense or sparse)
  - Background subtraction

Deep Learning vs Conventional CV

Approaches

- Deep learning
  - High detection rates possible
  - High false positive rates
  - Low precision on geometry
  - Poor generalization
- Conventional Computer vision
  - Could generalize well (especially geometry)
  - Low recall rates
  - Computationally demanding (potentially)
  - 3D estimation is an ill posed problem in this setup

Deep Learning + Conventional CV

2021 embedded VISION summit

Approaches

- Conventional Computer Vision + Deep Learning
  - Can generalize well (Geometry fundamental)
  - Low false positive rates
  - High recall rates
  - Range is solved mixing geometry and appearance (DL+CV)
  - No unbounded DL with unknown outcomes
  - Explainable

End-to-End Deep Learning for Safety Systems

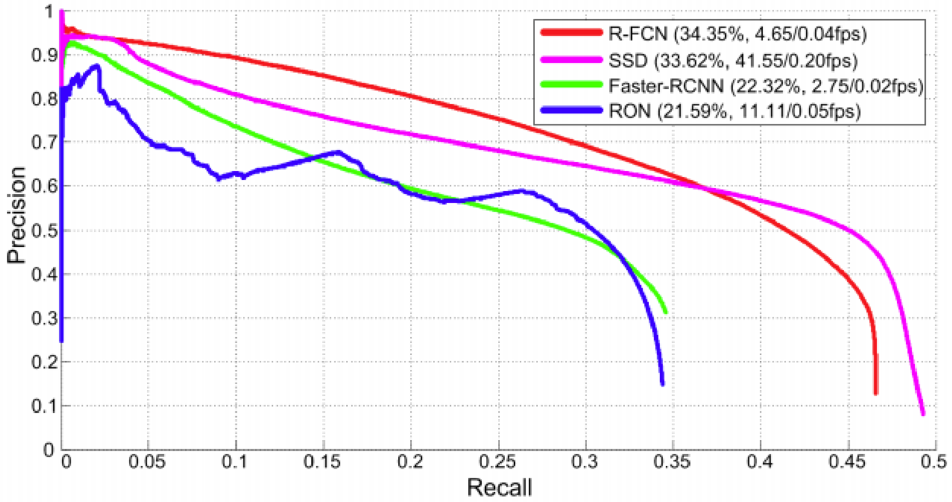

Example: Single frame object detection

1. Objects to sense are really far away (small and lower contrast)
2. Precision 0.95 → 1 false positive every 6 seconds (At 15fps)
3. Precision 0.999 → 1 false positive every 16.65 minutes
4. Unexpected actions are a possibility

Chart source: The Unmanned Aerial Vehicle Benchmark: Object Detection and Tracking - 2018, Du et al

Iris Automation
The future in sight.

Bounded Deep Learning for Safety Systems

Cued and bounded deep learning

1. Get basic understanding of 3D geometry
2. Get remaining unknowns (e.g., classification and range estimation)
3. Estimate location and velocity of intruder
4. Optional: Define best course of action

Bounded Deep Learning for Safety Systems

Cued and bounded

1. Get basic understanding of 3D geometry
2. Get unknowns remaining (e.g., That object is moving, what is it?)
3. Estimate location and velocity of intruder
4. Optional: Define best course of action

- 1 False positive every 10+ hours
- 0.9+ Recall rate
- No unexpected outcomes

False Positive Rates

Low False Positive Rates

Avoiding “crying wolf” is fundamental to most safety systems

- Build trust in system
- Reduce induced critical conditions (many false triggers could create unwanted safety conditions)

Cueing DL Models only half the story

- Online “hard case” mining is the other half

Low False Positive Rates

Online null mining

- Data infrastructure is critical
- Proximity between data and compute is critical (while training)

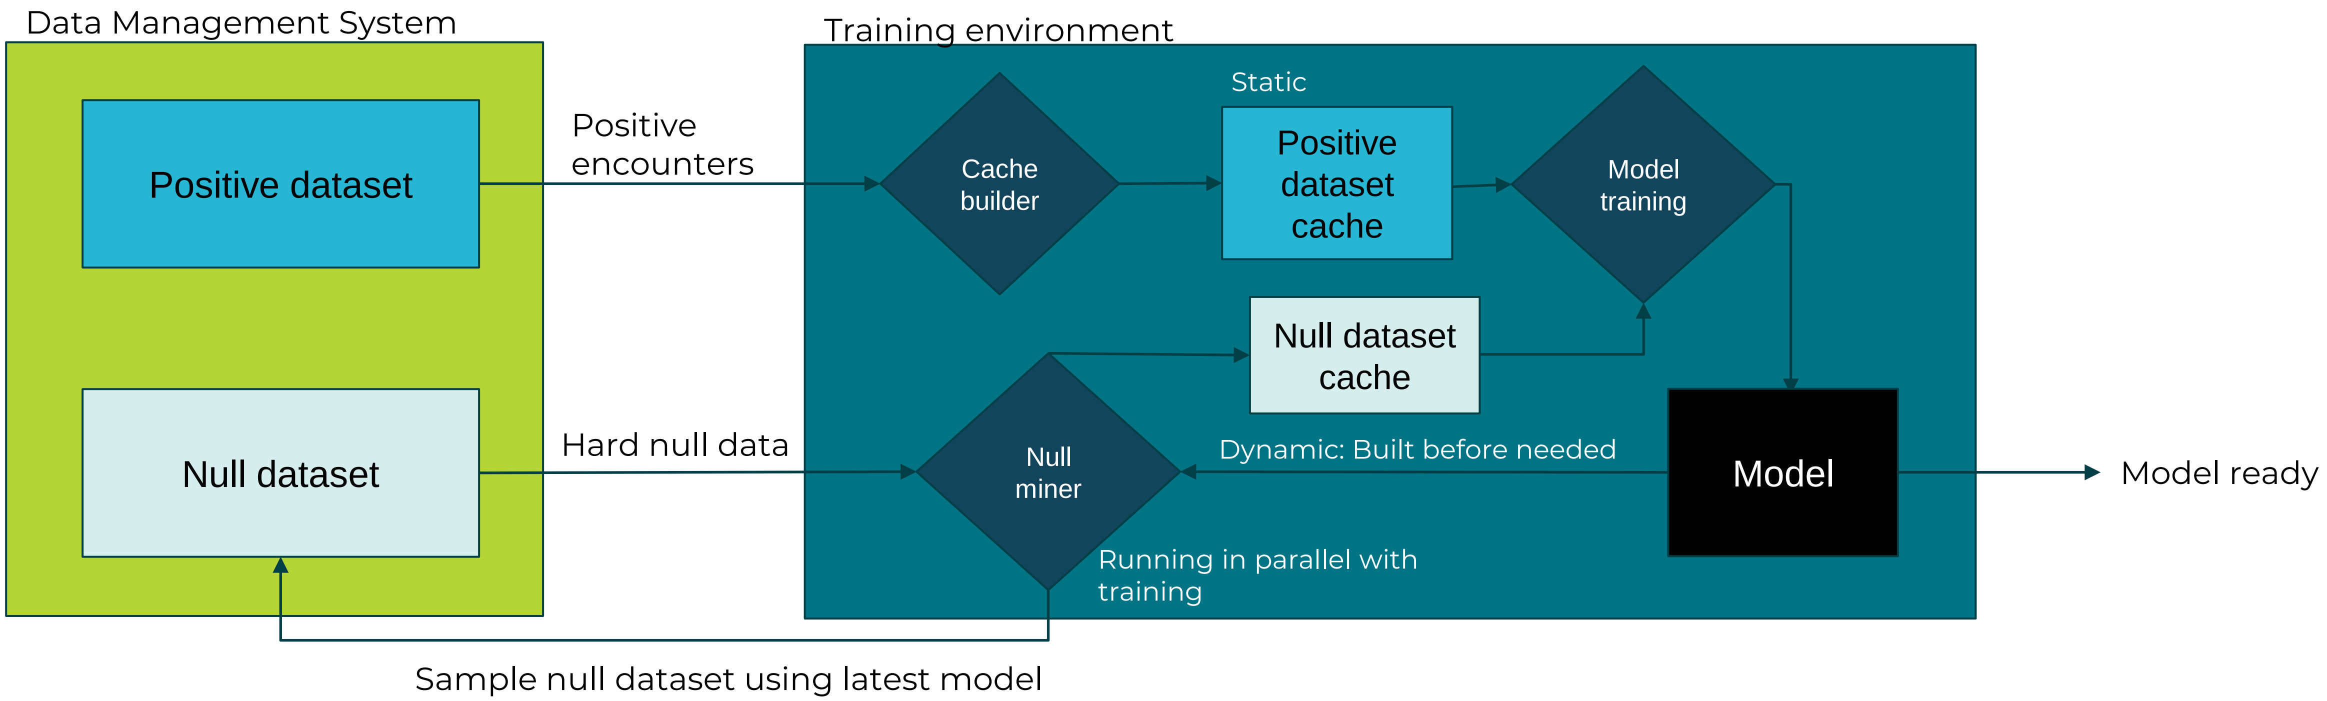

Examples

Simulating a 5-camera system on board of a drone

- Top image is a stabilized zoom of the detection
- Watch in max resolution and full screen

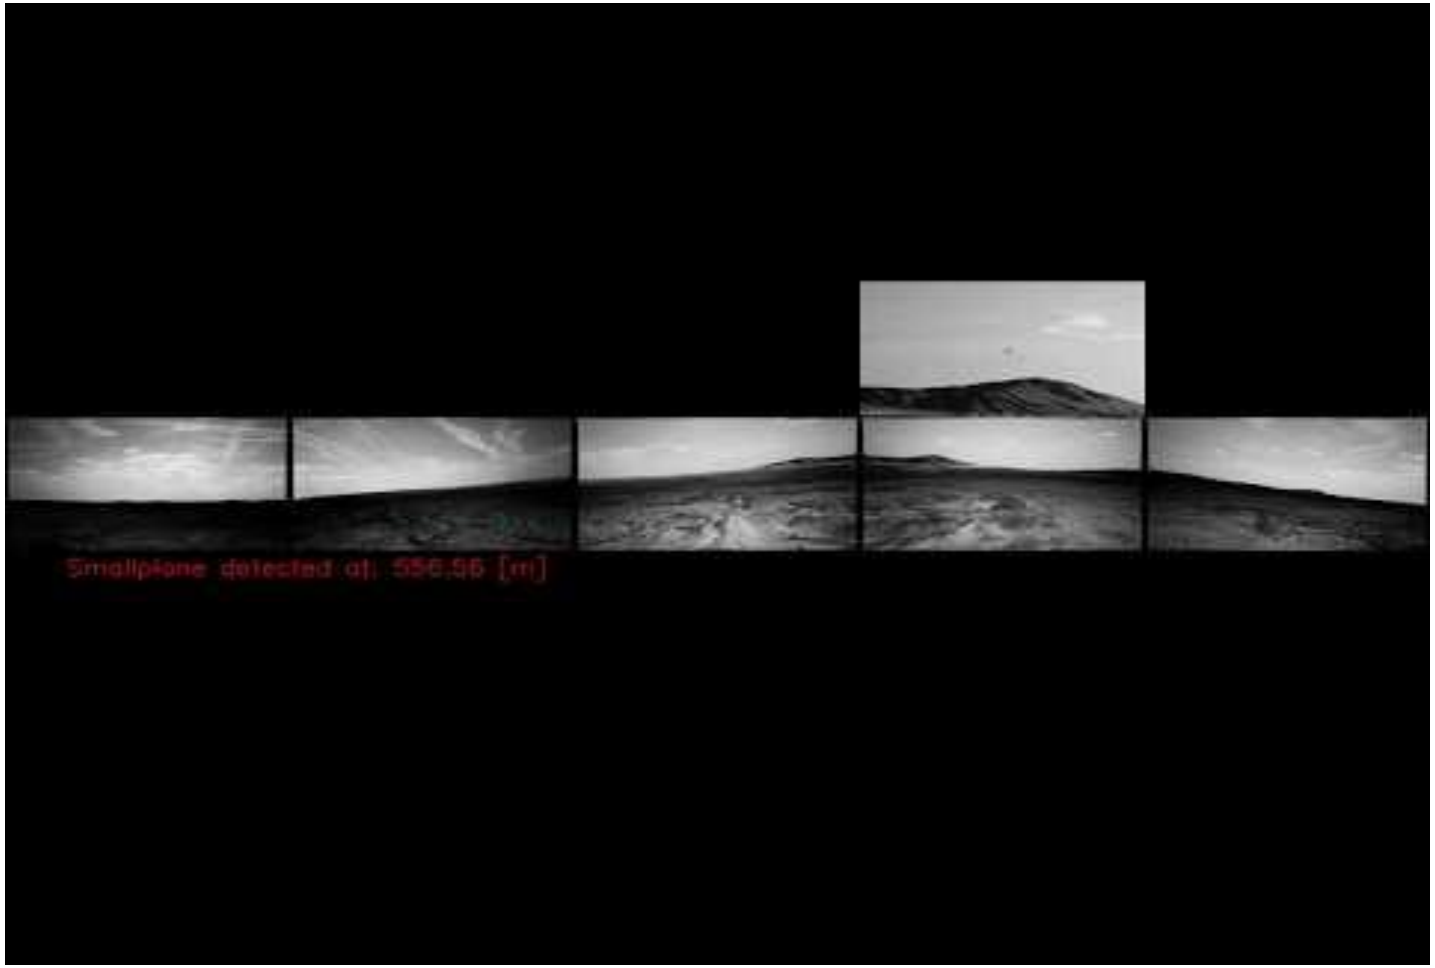

Iris Automation
The future in sight.

Conclusions & Learnings

How to evaluate safety systems?

- Evaluate on hundreds of thousands of examples (if not millions)
  - The number of examples might depend on the likelihood of an event
- Consider using synthetic data when possible
  - Many cases you can’t recreate in real life (safety concerns)
  - Many scenarios are just too expensive in real life
  - Domain gap can be bridged from both sides

Conclusions and Learnings

In the context of safety systems

- Consider your performance requirements and state of the art
  - Deep learning might help
  - Deep learning might be a good prototype

- Consider cue and bounding deep learning solutions
  - Focus compute on what DL is really good for
  - Cue to reduce the rate of false positives
  - Avoid unexpected outcomes

Conclusions and Learnings

Infrastructure and techniques

- Implement a data management system
- Hard cases mining is needed to deliver on solution beyond prototype

Resources

- Website (We're hiring!)
  https://www.irisonboard.com

- Excellent Multi-View geometry book
  https://www.amazon.com/Multiple-View-Geometry-Computer-Vision/dp/0521540518

- Hard case mining
  - Class Rectification Hard Mining for Imbalanced Deep Learning - Dong et al - ICCV 2017
  - Triplet Loss and Online Triplet Mining in TensorFlow