Cooperative and Non-Cooperative UAS Detection

Lucas Barduson[1]
Arizona State University, Tempe, Arizona, 85281, USA

Eloisa Carrasco[2]
University of Minnesota Twin Cities, Minneapolis, Minnesota, 55455, USA

Zach Khan[3]
University of Michigan, Ann Arbor, Michigan, 48109, USA

Kelsey Krstulic[4]
University of Kansas, Lawrence, Kansas, 66045, USA

This paper describes a system that enhances airspace situational awareness by detecting and identifying Unmanned Aerial Systems (UAS). This multi-domain solution tracks both cooperative scientific flights as well as non-cooperative intrusions from "bad actors." The system supports a critical push towards safety within NASA's advanced air mobility mission. After surveying the existing technologies at Langley Research Center, the radar and visual systems were chosen for the primary and secondary detection mechanisms, respectively. These systems were tuned an upgraded to become more sensitive to UAS activity. Additionally, a Remote Identification receiver was procured and integrated into the flight surveillance system.

I. Nomenclature

AAM = Advanced Air Mobility
ADS-B = Automatic Dependent Surveillance Broadcast
CERTAIN = City Environment Range Testing for Autonomous Integrated Navigation
CFAR = Constant Flight Alarm
COTS = Commercial off the Shelf
DAA = Detect and Avoid
kNN = k-Nearest Neighbors
LaRC = Langley Research Center
MQTT = Message Queuing Telemetry Transport
R-ID = Remote Identification
RF = Radio Frequency
sUAS = Small Unmanned Aerial System
TAURAS = Traffic Awareness and Ubiquitous Real-time Airspace Surveillance
UAS = Unmanned Aerial System
YOLO = You Only Look Once

---

[1] Undergraduate Degree in Aerospace Engineering (2024)
[2] Undergraduate Degree in Electrical Engineering (2025)
[3] Graduate Degree in Electrical Engineering (2024)
[4] Undergraduate Degree in Aerospace Engineering (2025)

II. Introduction

A. Background

Parallel to the exponential growth in Unmanned Aerial Systems (UAS) usage and commercial availability, there is an increasing threat in their potential for misuse. The Federal Aviation Administration's (FAA's) attempt to counter this threat is Remote Identification (R-ID) – a law that requires all UAS to have a "digital license plate". Significant radio frequency (RF) noise in an urban environment creates a significant challenge for detecting R-ID.

This paper will dive into different methods of detecting cooperative and non-cooperative UAS along with various ways of detecting R-ID. R-ID used in the U.S. is designed to service the needs of security personnel looking to increase accountability of drone operators.[5] The solution is a broadcast identifier that can be received using commonplace Wi-Fi and Bluetooth hardware. As of March 16, 2024, 14 CFR Part 89 requires UAS over 0.55 lbs and under 55 lbs must be equipped with some sort of R-ID transmitting device—except those with special exceptions. This broadcasts the drone's location, altitude, takeoff location, and unique serial number. However, the thousands of devices that communicate on those frequencies create heavy loads of traffic – ultimately leading to a realistic broadcast range that is similar to the range of a Wi-Fi router.

B. Previous Works

Studies have been done to find the most optimal constant false alarms (CFAR) algorithm. CFAR is an algorithm that is trained on a supervised set of images so that it can build a memory and give an accurate classification **[1]**. NASAs GA-9120 doppler radar that takes advantage of such an algorithm to better classify objects against clutter. There are several different CFAR algorithms that improve up the classification quality, but they are currently not used in the GA-9120.

Working toward a common goal, the Distributed Sensing Team at Langley Research Center (LaRC) has been researching and testing ways to detect and track UAS. Their vision includes a collection of systems each sensing a specific domain to contribute to general air traffic awareness and advanced air mobility (AAM) navigation. A key element they have worked on is making a radar-camera payload that would serve as an on-board UAS detection system for detect and avoid (DAA) maneuvers. They also use similar technology for ground-based nodes, where some are dedicated for tracking single UAS and some are dedicated for groups of UAS or cluttered environments **[2]**.

C. Commercial off the Shelf (COTS) Considerations

After extensive research, it was determined that LaRC does not hold the proper technology to meet each requirement that was placed on the team. The original plan was to have a multi-layer system that starts with the radars. Once something is detected, RF and R-ID would be used to more accurately track the object the radar is detecting. Finally, cameras with a computer vision algorithm would identify the object by placing a box around the object to make classification easier. Initially, multiple companies were sought out, but this path was not chosen in favor of a more cost-effective solution taking advantage of existing infrastructure.

III. Methods and Results

D. Radar

Radars consist of a transmitter and receiver which are sensitive to specific frequencies. Emitted RF reflects off a surface and is received by the radar. This return time delta corresponds to the distance from the radar to the target. Doppler radars specifically measure the change in frequency between the emitted and measured signal; the change corresponds to the velocity of the target relative to the observer. LaRC uses two radars for primary airspace observation. The LSTAR radar is mounted on the southeast side of campus and the GA-9120 system is mounted in the northwest. LSTAR has a 50 NM diameter detection range and accurately detects general aircraft across its 360-deg azimuthal field of view. In its current configuration, the LSTAR can only detect objects larger than most small Unmanned Aerial Systems (sUAS) which is not applicable to this area of the AAM mission.

The GA-9120 is a Doppler radar system that is used to detect sUAS and cover the blind spots of the LSTAR. Two GA-9120s are mounted to the top of the Gantry facility. The north facing radar has a pitch of 0 deg and the east facing

---

[5] Taken from interview with Andy Lacher, consultant on R-ID development.

radar is pitched up 5 deg. It has a 12.5 deg elevation field of view and a 120-degree azimuthal field of view. The software for the radar can classify the objects that were detected, but it is severely unreliable. When the GA-9120 was installed in 2022, it detected smaller objects at lower altitudes, which could be an ideal solution contributing to NASA's AAM mission. The amount of noise shown in (Fig. 1) being detected causes difficulty to distinguish real and false positives. This configuration was tuned to remove the extra noise and false positives to clearly and accurately detect UAS. The radar in its early 2024 configuration detects birds and cars but is struggling to pick up faster moving sUAS (Fig. 2).

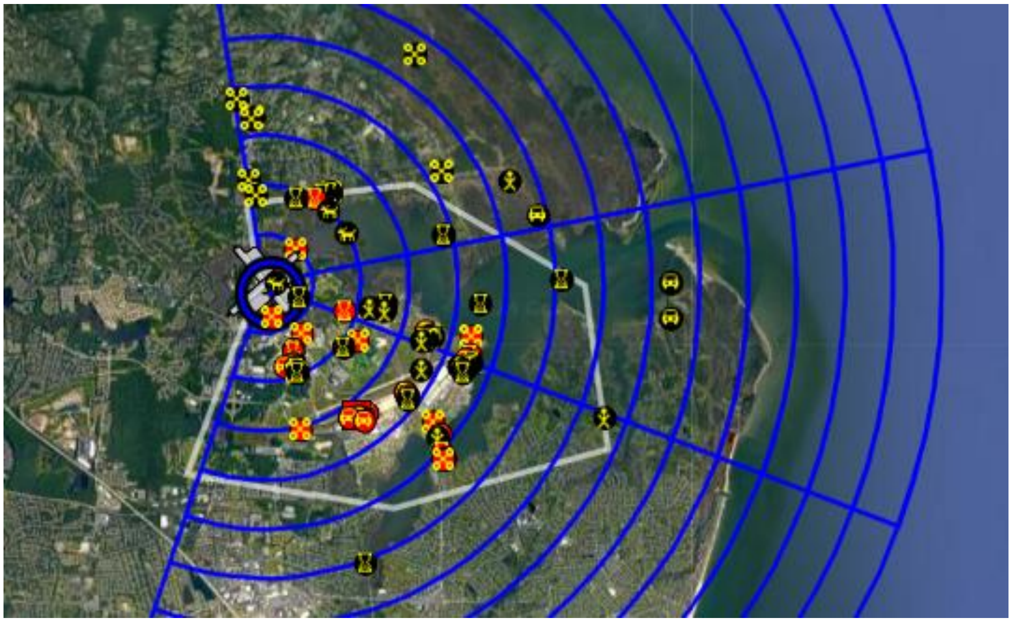

Figure 1: GA-9120 Original Configuration (2022)

Figure 2: Early 2024 GA-9120 Configuration

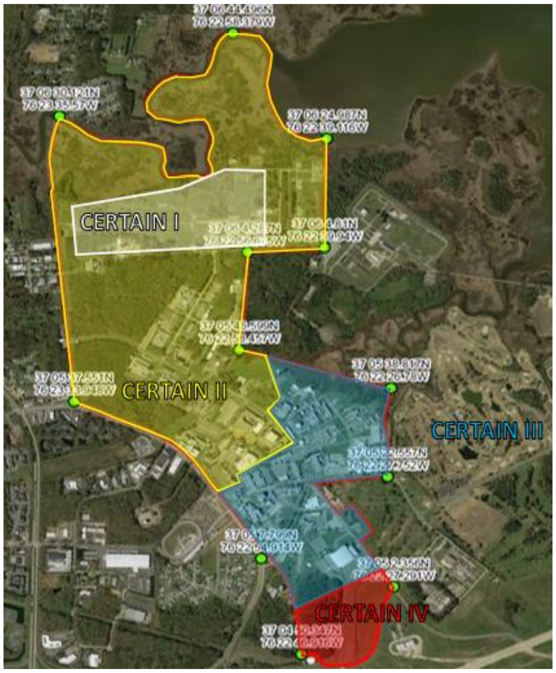

Figure 3: CERTAIN Range Locations

Configuration One:

The first configuration was developed by merging the 2022 and early 2024 configurations, resulting in a new configuration which reduces clutter (Fig. 4) while maintaining the capability to detect the UAS. At first glance, the configuration appears to do exactly what was expected. After a flight test with a Skydio, Alta 8, Alta X, and Phantom 4, it was proved that the GA-9120 could not accurately detect the UAS. The Skydio was detected at City Environment Range Testing for Autonomous Integrated Navigation (CERTAIN) range one (Fig. 3) and the Alta X was detected at CERTAIN range three (Fig. 3). These results were unexpected because only the smallest and largest UAS were detected. Another flight observed was a fixed-wing Hybrid Supervolo UAS flown at 54 kts inside the radar detection volume in a lawn mower pattern with circles at the end of the of the pattern. Theoretically, the radar should have clearly detected and tracked the continuous and repetitive movements, but the configuration failed and the UAS was not detected despite flying within the radar's detection volume (Fig. A-1).

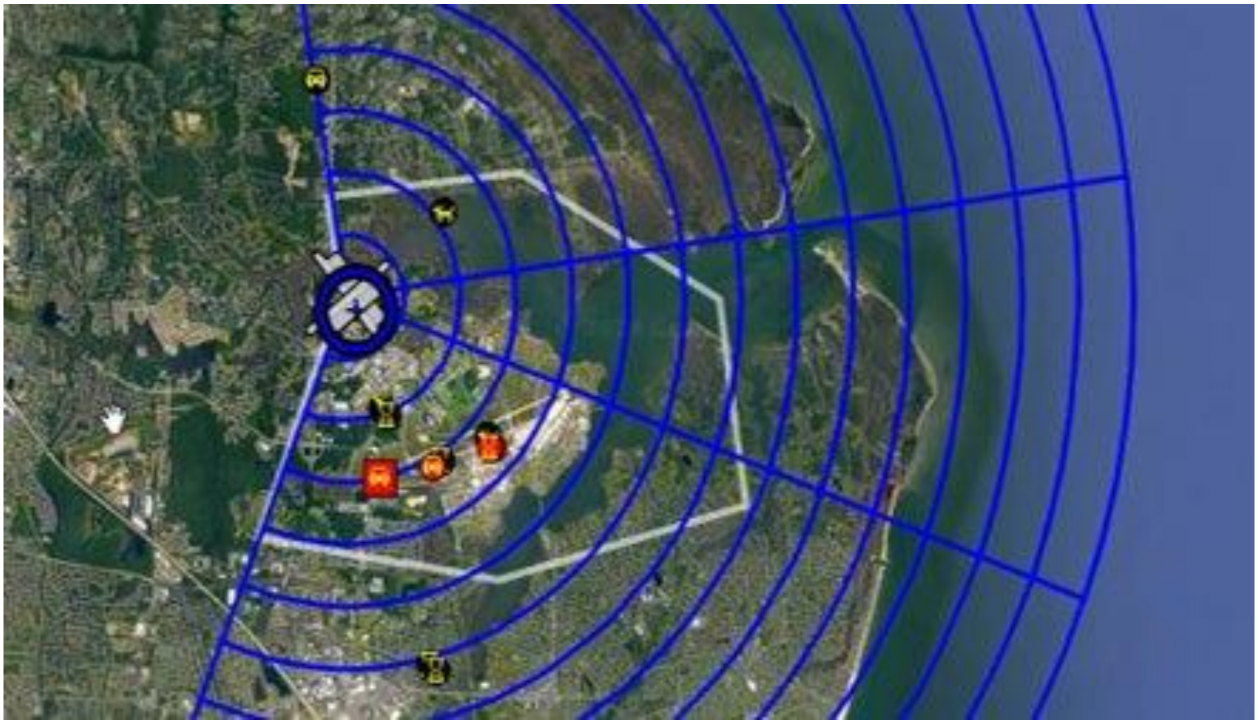

Figure 4: HORUS Configuration One

Configuration Two:

The second configuration focused on detecting lower altitude and slower moving objects with less restrictions on creating tracks. Another goal was to get rid of the classification options for cars, tractors, wildlife, and people. Configuration one on the radar displayed a choppy pattern that cut in and out for the tracks, which made it challenging to monitor the UAS effectively as it was being detected. Since this was the most important to fix, most changes to the configuration were made to the tracking algorithm parameters. During observation of scheduled UAS flights, it still detected most of the cars on the busy road outside Langley (Fig. 5), despite the radar being configured to ignore vehicles. Another thing that was noticed was the "spaghetti" cluster that formed at the northern range (Fig. 6). This type of trail is created when the radar detects one object but leaves "firework" or "spaghetti" trail, creating many false positives; it made the real objects hard to detect and classify. Overall, this configuration seemed to be worse than configuration one, but each problem had an easy fix to it.

Figure 5: Armistead Avenue North V2

Figure 6: "Spaghetti" Tracks

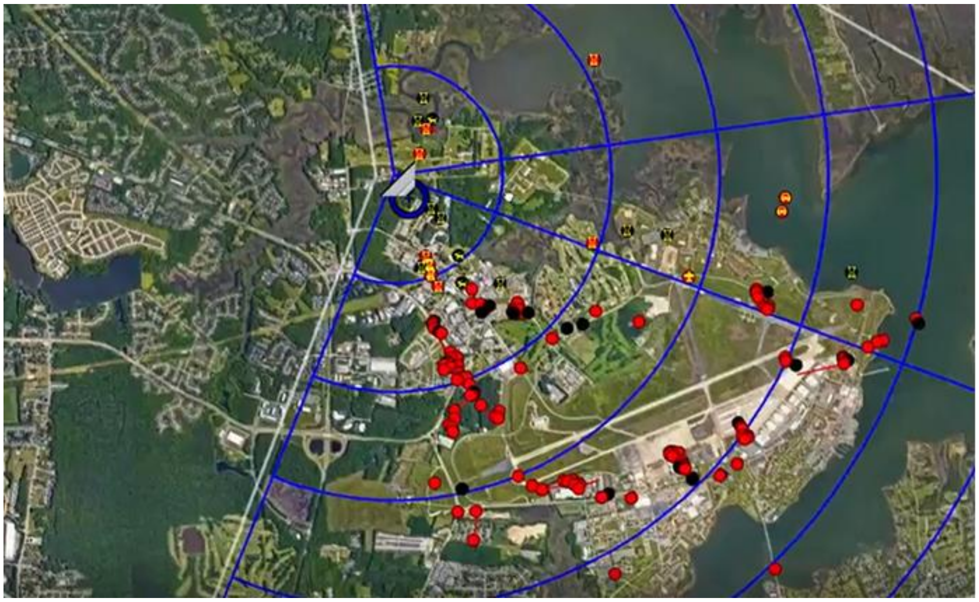

Figure 7: HORUS Configuration two

Configuration Three:

Configuration three aims to declutter the radar and improve classification accuracy in other areas by filtering out irrelevant data on Armistead Avenue and the Airforce base. To address the issue of "spaghetti" tracks and clutter, the cross-section detection size was increased to minimize the risk of detecting unwanted objects. The minimum floor altitude was also raised to specifically target UAS and reduce noise from vehicles and wildlife. However, this adjustment creates the potential drawback of possibly missing smaller UAS due to their small cross-section. It's important to note that this configuration was designed purely to control clutter and was not used to observe UAS flights, so further validation would have been needed to ensure it met the intended goals effectively.

Based on basic surveillance, the radar improved on decluttering Armistead Avenue and the Airforce base but is still struggling to classify each object as seen in figure 8 and figure 9. The figures look as if they were taken minutes or even hours apart, but they were only taken seconds apart. It is unclear why there is a stark difference, and there is not an intuitive way to fix the problem.

Figure 8: Armistead Avenue North V3

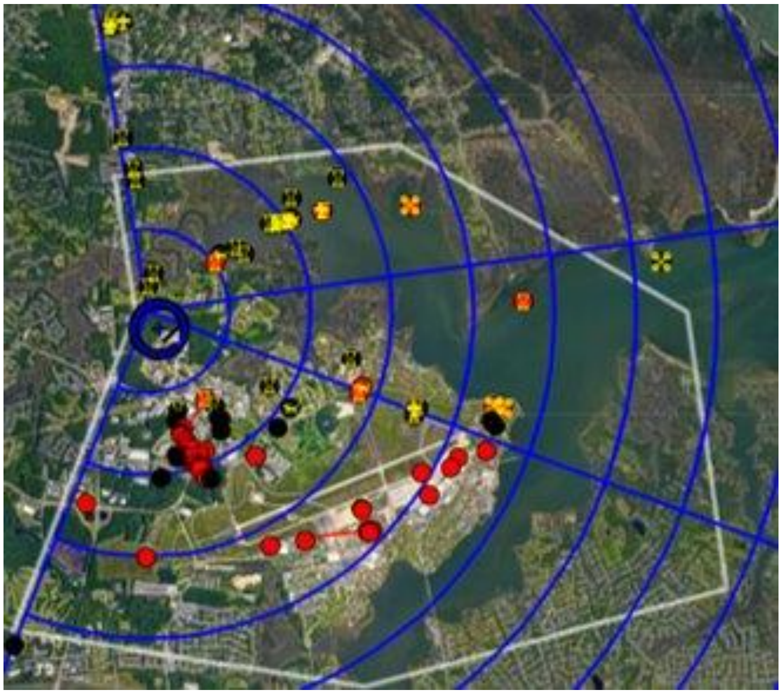

Figure 9: HORUS Configuration Three

E. RF Signals

Unless they are running entirely autonomous operations, all UAS emit RF signals. This can occur across a multitude of frequencies including Bluetooth (2.4 GHz), Wi-Fi (2.4 and 5.8 GHz), cellular (LTE and 5G), or simple radio (25-75 MHz). LaRC has five power spectrum analyzers which measure the power that is transmitted across a range of frequencies. The system's software features a 20-second live data visualizer to see what is being emitted on the spectrum.

The system was put to test by observing the strength of RF signals while UAS were flying. At 2.4 GHz, significant noise inhibited the ability to see any discrete sources or frequency hopping patterns. Radio frequencies and 5.8 GHz were not tested. LTE proved difficult to even track due to the multitude of possible frequencies. While the additional geolocation software may be been interesting to test, it seems likely that it still would have trouble finding a target at 2.4 GHz – especially one moving at fast speeds. RF tracking may, however, be a useful tool for locating the location of a remotely controlled UAS's pilot.

A much simpler way of tracking UAS with RF signals is with R-ID. This provides all necessary information for the safety and security of the Langley Research Center airspace; however, it is still not a perfect solution in practice. The Wi-Fi and Bluetooth frequency bands can be very busy in populated areas and do not transmit well through buildings and densely wooded areas. Additionally, many drones still require an external device to broadcast R-ID signals, meaning bad actors can easily circumvent R-ID detectors by simply taking off the transmitter. Thus, R-ID is limited to detecting compliant UAS and providing air traffic monitoring for LaRC's UAS operations.

The solution chosen for detecting R-ID was the DroneScout ds240 from BlueMark Innovations. The ds240 is a COTS R-ID receiver equipped with 15 dBi antennas which can detect signals at all R-ID frequencies up to 15 km away **[3]**. A proof-of-concept mount for the antennae and receiver was setup at ground level (Fig. 10). Initial testing was carried out to verify that the setup worked as intended and get an idea of the range that could be achieved in the specific environment. This testing consisted of setting up the ds240 at the CERTAIN one range and flying an Alta-8 and Skydio. The tests were largely successful, with some results shown below in figure 11. The ds240 was successful in receiving the R-ID signals, automatically storing them in its archive of flights, and doing so for even the furthest flights tested, which reached as far as 3900 ft (just under 1.2 km) from the receiver.

Figure 10: ds240 Setup

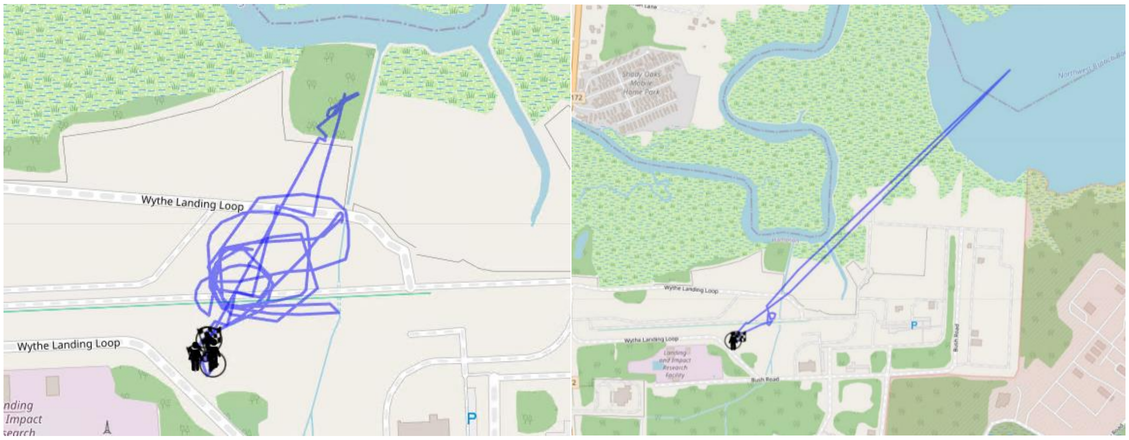

Figure 11: Archived flight paths of initial range test with Alta-8 and Skydio

The use of two different drone types had an added benefit of demonstrating that the ds240 could detect signals both from an external R-ID module (attached to the Alta-8) and a built-in module (included by default in the Skydio). The only unsuccessful aspect of this testing was that the ds240 was not displaying the live drone position on its user interface, the DroneScout Dashboard. Through repeated refreshing of the archived flight path while the drones were flying, a pseudo-real-time display was still possible despite the apparent limitations of the Dashboard software.

The second set of testing sought to remedy the lack of live data while also integrating the ds240 into the existing Traffic Awareness and Ubiquitous Real-time Airspace Surveillance (TAURAS) system currently utilized by LaRC. TAURAS is a centralized view of several safety data streams, including Automatic Dependent Surveillance Broadcast (ADS-B), radar, and now R-ID. These tests involved more flights with the ds240 placed at the same location on CERTAIN one, this time with only one flight of a Skydio. Between the initial testing and this second set of testing, TAURAS engineers developed a temporary integration by having the ds240 send data to an on-site message queuing telemetry transport (MQTT) broker connected to TAURAS. Unfortunately, this meant it was not sending data to the DroneScout Dashboard and thus the flight path archive could not actively be read. During the flight, a computer screen displaying the TAURAS and Skydio interfaces was recorded (Fig. 12).

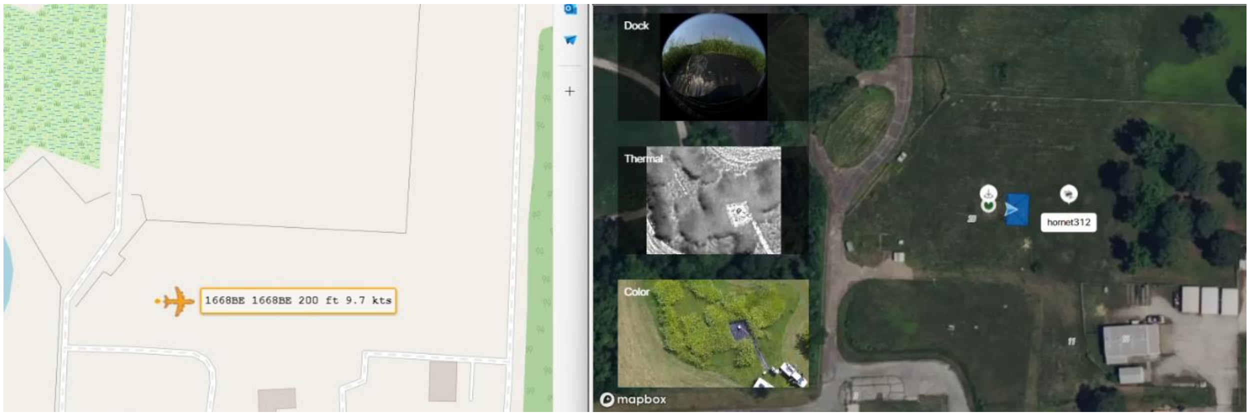

Figure 12: TAURAS on the left, Skydio telemetry software on the right

This flight was fully successful, demonstrating both that the ds240 data could be displayed in real time and that this real time display could be accomplished in TAURAS. This test also helped define more of the boundaries of the ds240's capabilities in that the Skydio could only be detected starting at approximately 100 ft of altitude. This is almost certainly due to a grove of trees about 90-100 ft tall being between the drone and the receiver. There was a brief issue with displaying the altitude in the TAURAS feed; however, this was resolved during the flight. More testing would need to be done to determine if this is a persistent problem.

An important thing to note about these two tests is that they are carried out in mostly open areas with low interference and nearly always a visual line of sight to the target. However, it was discovered during the second test that the ds240 had made two unintentional detections of recreational UAS in the area. These are believed to be true positives since they had different identifying numbers than the NASA-flown UAS. The first was during a previous day of testing where it picked up signals from the opposite side of the Northwest Branch Back River, at the end of Robert Bruce Road. This was nearly twice as far as the initial testing at 1.28 miles (just over 2km) from the receiver; the flights were recorded in the DroneScout Dashboard archive. The second detection happened as the second test session was wrapping up. A UAS was detected 2.23 miles (over 3.5km) away from the receiver. This was doubly important for demonstrating the ds240's capabilities because it was both further away than any previous testing and in a direction that had more trees and noise than any previous testing. Unfortunately, the screen recording had been stopped by the time it was discovered, but figure 13 shows its approximate location in relation to the receiver.

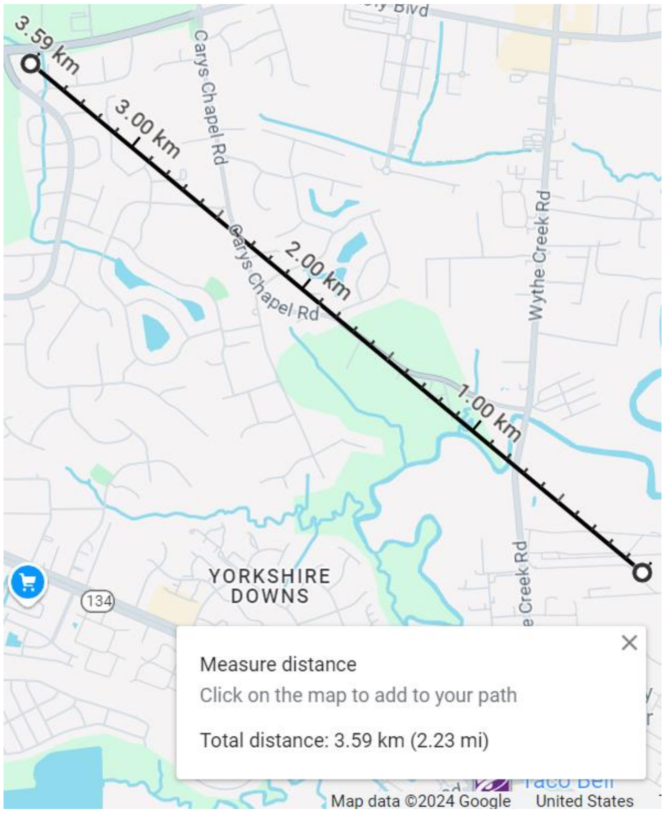

Measure distance

Click on the map to add to your path

Total distance: 3.59 km (2.23 mi)

Figure 13: Approximate location and distance of commercial drone detected by ds240

Overall, the results from utilizing R-ID have been very promising. The range, precision, and quick update rate fit the needs of the system. However, there are still future improvements to be made. The ds240 still needs to be properly integrated into TAURAS rather than using the temporary integration. A permanent antenna mount also must be created and installed wherever the ds240 will be permanently located, and the full extent of the receiver's range should be explored.

F. Visual

The existing camera system on the northern CERTAIN range consist of stationary bullet cameras that cover different areas of the range. Most of the cameras are positioned in a way that does not show much of the sky, so they are not in an ideal placement for detection or tracking of UAS during flight operations. While more capable hardware was not attainable during the project timeline, a software solution using computer vision with the existing bullet cameras was explored. The primary goal was to direct attention to the presence and location of UAS by outlining the vehicle with a box. A secondary goal was to classify the UAS based on structure, size, and model.

This effort started by leveraging a software packaged called You Only Look Once (YOLO), specifically YOLOv10, because of its ability to use existing datasets to create a robust system somewhat quickly. The main issue with using YOLOv10 is that the targets are so small that they are unrecognizable by machine learning algorithms that are trained on common drone datasets which use pictures of drones from only a few feet away. This problem is worsened by the fact that a significant portion of the CERTAIN camera views are of grass and trees, which create a noisy background that makes it difficult to spot a small target. Additionally, the laptops provided to the team struggled to run such a graphically intense algorithm.

Rather than using YOLOv10, a k-Nearest Neighbor (kNN) background subtraction algorithm for object detection was investigated using OpenCV in Python, along with a Euclidean distance calculator for object tracking. This method analyzes new frames based on a history of previous frames and subtracts the ‘background’. When paired with a thresholding function, this results in a mask that is essentially a motion detector, where pixels that have changed appear white and pixels that have not changed appear black. Finding contrast on this mask allows for object detection, which allows for tracking based on Euclidean distance from the object’s center. The program then draws a box around the object, fulfilling the primary objective of the detection system as seen in figures 14 and 15.

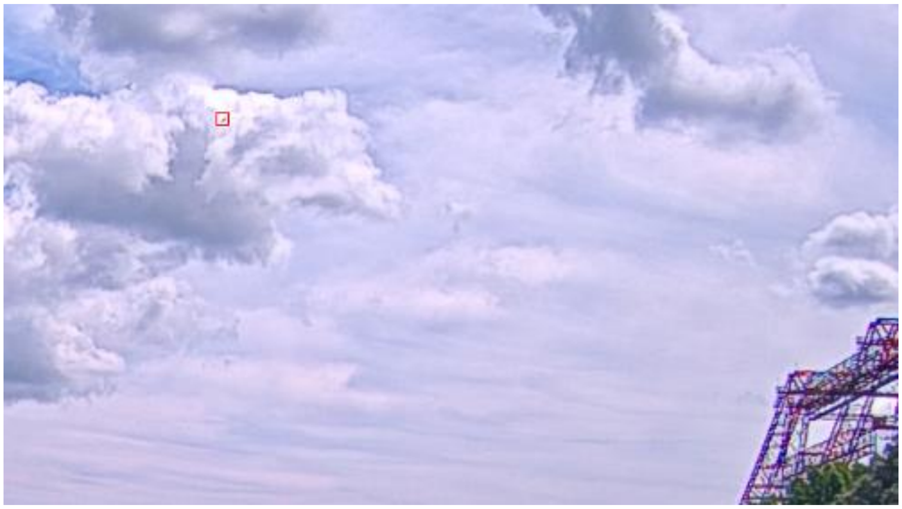

Figure 14: Bird detected by CERTAIN camera

Figure 15: Bird detected by kNN filter

This method significantly improves the problems with YOLO mentioned above. Firstly, it does not need to be trained, saving the time and effort of creating a custom dataset. Secondly, its use of frame history allows it to better adapt to the dynamic environment in the camera feed, allowing the program to detect movement as small as a single pixel. It is also less computationally expensive. While still not able to run at the full 30 frames per second of the camera feed on the team's laptops, it ran faster than YOLO, and well enough to prove the concept. Finally, the code is simple and accessible to tune for better performance.

Most of the testing and tuning was focused on CERTAIN camera 3 because of its long sightlines and diverse backgrounds to detect against. Due to the low volume of UAS flights, the OpenCV program was usually tested on birds flying in the area. This resulted in a tuning process that consisted of making a few small changes between spontaneous bird sightings.

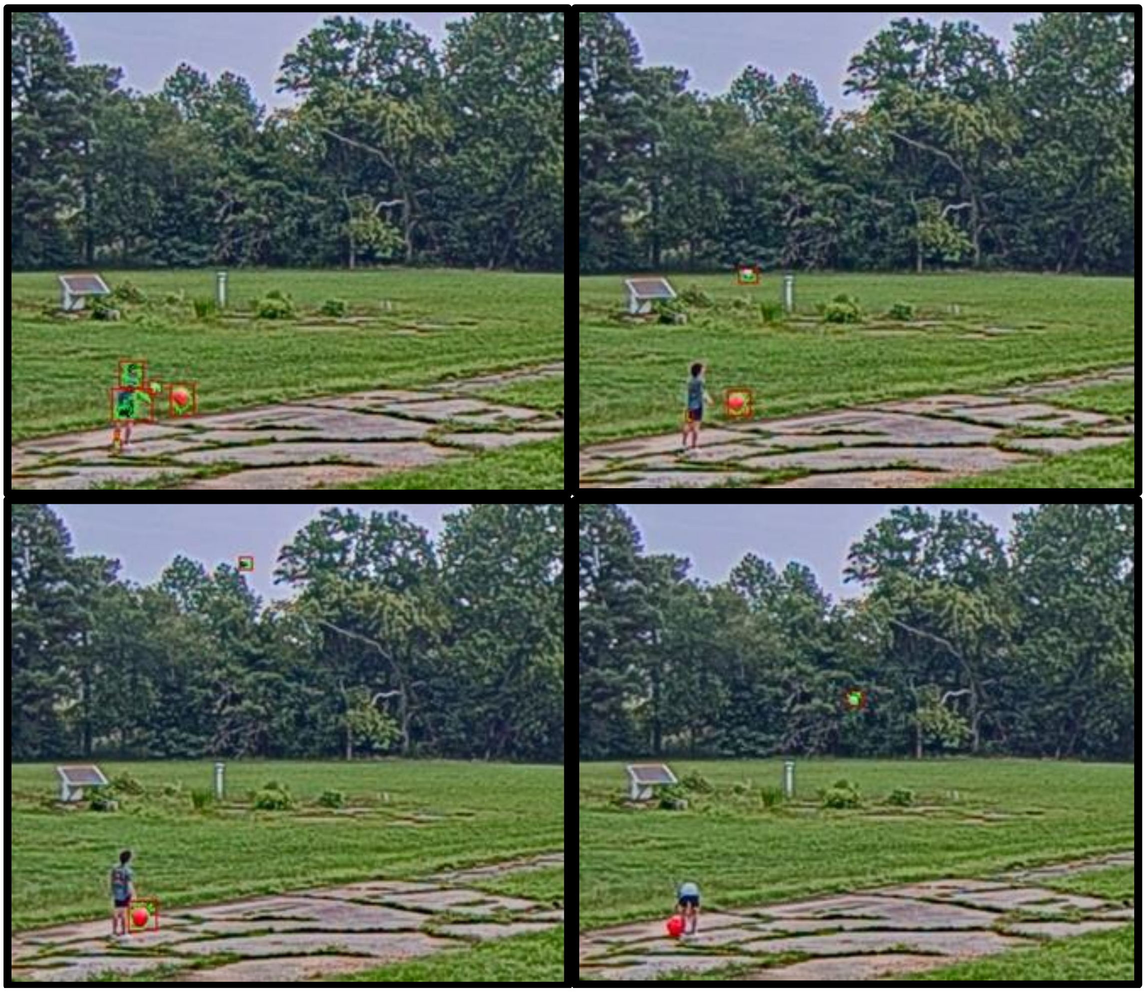

Figure 16: Real time tracking of objects approximately 110ft from camera

IV. Conclusion and Recommendations

To enhance the airspace situational awareness of LaRC and to improve the capabilities of advanced air mobility research, different modalities of detecting and classifying objects were explored. The primary instrument to accomplish this task should be the GA-9120 due its large detection range and its ability to track objects. Since the radar is prone to unclear tracks and false positives, the computer vision algorithm and R-ID receiver act as secondary tools for confirmation. R-ID can keep track of the compliant UAS, and the vision system would detect non-compliant UAS, compliant UAS, and birds.

Recommended changes to the GA-9120 configurations includes raising the confidence level for an object to be broadcasted on the map thus creating less clutter. Decreasing the minimum track quality for the radar to start classifying an object allows for a more accurate classification. The minimum altitude for detection can be increased to lessen the likelihood of detecting cars and wildlife that are not relevant. However, this creates a dead zone where any UAS that are flying low to ground will go undetected.

To improve the computer vision system, working on the software should take priority as it can be used on any camera that is accessible. Tuning the tracking threshold within the software would improve its accuracy. This is an important factor considering the differing sizes of UAS and the speeds at which they can move. Positioning and adding cameras are also a recommendation, but further investigation should be done as to what types of cameras and where they would be placed.

Appendix

Appendix A:

Figure A-1: GA-9120 Detection Volume

Acknowledgments

The authors would like to extend their gratitude to Christopher Newport University and NASA Langley Research Center for funding the members of this team and other members in the NASA Aeronautics Academy.

We would also like to thank Matthew Coldsnow, Matthew Underwood, Paul Krasa, Kyle Smalling, and Patrick Quach for being at our weekly advisory meetings and providing helpful input.

Additionally, we thank James Nicolson, Skyler Hudson, Patrick Hill, Bryan Petty, Paul Sugden and Benjamin Jenkins for not only volunteering their time to HORUS but going out of their way to push for the success of this project.

Thank you to Kelly Au and Steven Williams for editing and revising this paper.

Finally, we are grateful to Elizabeth Ward giving us this chance to start our professional career at NASA. We thank you for all of the time and effort you have put into the Academy.

References

[1] Wang, C., Tian, J., Cao, J., and Wang, X., "Deep learning-based UAV detection in pulse-doppler radar," IEEE Transactions on Geoscience and Remote Sensing, vol. 60, 2022, pp. 1–12.
doi: 10.1109/TGRS.2021.3104907

[2] Chester Dolph, Thomas Lombaerts, Corey A. Ippolito, Vahram Stepanyan, Evan Kawamura, Keerthana Kannan, George Szatkowski, Todd Ferante, Christopher Morris, Federica Vitiello, Flavia Causa, Roberto Opromolla and Giancarmine Fasano. "Distributed Sensor Fusion of Ground and Air Nodes using Vision and Radar Modalities for Tracking Multirotor Small Uncrewed Air Systems and Birds," AIAA 2024-1781. *AIAA SCITECH 2024 Forum*. January 2024.

[3] DroneScout ds240 receiver. BlueMark; 2024 [Available from: https://dronescout.co/product/dronescout-ds240-receiver/?v=7516fd43adaa]