Passive Aircraft Detection System & Camera Array

Alexander Alonzo, Xavier Colon, Carter Clavell, Nathanael Olson, Kai Doldan, Garrett Mosrie

Dept. of Electrical Engineering and Computer Science, University of Central Florida, Orlando, Florida, 32816-2450

***Abstract* — The P.A.D.S. (Passive Aircraft Detection System) project presents a modular and affordable dual-system surveillance device, using passive RF detection and super-resolution imaging. Designed to address radar blind zones and general passive monitoring, the system is able to detect aircraft using ambient RF signals using off the shelf components and verifies them optically using a multi-lens camera array. Super resolution algorithms reconstruct high-resolution images from low-resolution inputs. Our system has the potential to provide valuable and sensitive information while simultaneously being very affordable.**

***Index Terms* — Aircraft detection, passive radar, super-resolution, camera array, surveillance.**

I. Introduction

Surveillance of remote airspace is often limited by the prohibitive cost and infrastructure demands of active radar systems that are typically used. In areas such as the borders of a nation, remote research facilities, and underdeveloped regions, this results in "radar dead zones" referring to the gaps in detection coverage that can be exploited for illegal activity such as smuggling, unauthorized drone flight, and covert surveillance.

Our senior design project, sponsored by Roy H. Wubker Jr. of the United States Southern Command (SOUTHCOM), addresses this vulnerability with a two-step solution. First a passive RF detection unit capable of identifying the presence of an aircraft using reflected ambient signals, and a super-resolution optical imaging system that visually confirms and enhances the detected aircraft.

Unlike commercial solutions which can cost upwards of millions of dollars and make use of technology that the United States may not want to share with certain countries that ask for them, our system uses affordable, off-the-shelf components and computational techniques to achieve comparable performance, all while costing magnitudes less than typical active radar systems. This paper will present the system architecture, design decisions, results from testing, and future development roadmap as this project is meant to be a stepping stone to developing more cost-effective solutions that combine optics and electronics.

II. System Overview

Our system consists of two major subsystems; The passive RF detection unit, which detects anomalies in ambient RF signals (FM radio, GPS, etc.) using a log-periodic antenna and signal analysis and the super-resolution camera array which captures images simultaneously using multi-field images and, using AI to properly align and feature match, fuses them into a high-resolution image using back-projection super-resolution methods.

Both subsystems are powered by a power bank and are integrated with our laptops for ease of presentation. The systems currently are individualized to allow for less overall risk but has the potential to be centralized as will be discussed in the section covering future works.

III. Passive Aircraft Detection System

A. Motivation

The system we are designing is a low cost, low power remote sensing unit. Its purpose is to monitor small aircraft and radio communications in remote areas where infrastructure is lacking, particularly in regions of national security interest. The motivation behind this project is to develop a fully autonomous system that operates without the need for maintenance or user intervention once deployed. Additionally, we aim to expand the capabilities of existing technology by enabling a wide range of passive detection functions. While passive detection has seen limited implementation and research, our goal is to explore its potential and evaluate the effectiveness of such a system in real-world conditions.

B. Design Approach

The system architecture is designed around simplicity, modularity, and commodity RF component compatibility to enable quick prototyping and repeatable testing. Each subsystem was selected to make trade-offs between performance and convenience of integration, particularly in the case of spectrum analysis and RF front-end testing.

The RF front end begins with a wideband active antenna with an onboard low-noise preamplifier. To supply this preamplifier, an isolated 9V DC power source is required for the antenna. For this project, a standard 9V battery was chosen to feed power directly to the antenna. Although Aventas/WinRadio does provide another solution via DC biasing over coaxial cable—allowing the antenna to be powered directly from a common power source—we employed the isolated battery for preliminary testing. This approach obviates wiring overhead and minimizes potential points of failure during development, particularly in field installations or mobile setups where reliability is an issue and portability is important.

The RF signal from the antenna is received by the spectrum analyzer via a common SMA male-to-male coaxial cable. SMA connectors were selected based on their widespread use in RF systems and their inherent property of maintaining a stable 50-ohm impedance over a wide frequency range. In our testbed, a commercially available SMA cable with moderate shielding was used. As inexpensive as it was, this cable provided sufficient signal fidelity for frequency and bandwidth ranges of interest, confirming the feasibility of utilizing readily available components at the expense of no loss in system accuracy.

The spectrum analyzer is connected to the host computing platform—a Raspberry Pi 5—via USB 3.0. The interface is two-purpose: it can transfer high-bandwidth data from the analyzer to the host and supply power to the spectrum analyzer itself. This two-purpose interface significantly simplifies the system's power design and eliminates the need for a separate DC power line to the analyzer. By using USB 3.0, we are compatible with most modern computing platforms and taking full advantage of the high data rates required for real-time monitoring and logging of the spectrum.

It is powered by a typical USB C. The decision to use an embedded computing platform such as the Raspberry Pi 5 as a small-form-factor computing solution is derived from its equilibrium in processing capacity, physical dimensions, and USB peripheral compatibility. This platform is particularly well-suited for deployment in space-restricted applications such as mobile test labs or embedded RF monitoring stations.

Generally, the system design reflects a strategic focus on modularity and ease of use. By leveraging universally available RF and computing interfaces—SMA, USB 3.0, and low power—we facilitate easy assembly, operation with assurance, and rapid iteration in development and field test cycles.

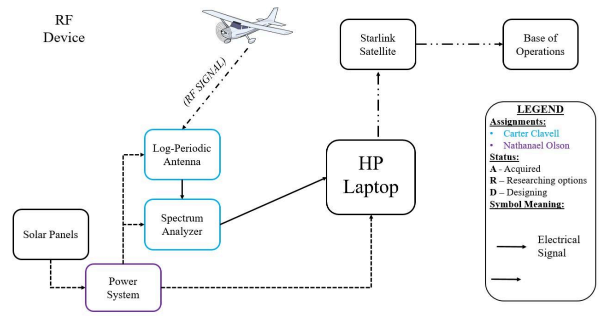

Fig. 1. Schematic diagram of RF system

C. Signal Processing Workflow

To view and interact with live RF data, we utilize Signal Hound's Spike software, which connects directly to our spectrum analyzer via USB 3.0. Spike provides a feature-rich and customizable graphical interface for displaying the RF spectrum, allowing users to see frequency power levels, bandwidth utilization, and spectral density in the time and frequency domains. It is an invaluable tool for live debugging, tuning, and capturing RF characteristics in field trials and controlled experiments.

Its combination of the wideband active antenna, SMA signal pathing, and USB-powered spectrum analyzer makes it an easy integration with Spike. The USB 3.0 connection to the Raspberry Pi 5 delivers sufficient bandwidth to stream high-resolution spectral data in real time. Spike support for Signal Hound analyzers delivers minimum latency and maximum fidelity in spectrum capture, which makes it an essential component in our signal chain.

Apart from plain visualization, our system also encompasses a machine learning-based model for detecting changes, anomalies, or unexpected behavior in the RF spectrum. The model accepts real-time spectral data captured from Spike (via exported or API-accessed data streams) and identifies deviations from normal spectral activity. The model is trained on baseline RF environments to distinguish between nominal and anomalous conditions. The procedure for applying the model to new spectral data is outlined in Fig. 2.

This cognitive layer of analysis enables the system to go beyond simple passive monitoring, with the ability of proactive identification of interference sources, unauthorized transmissions, or changes in the environment. It includes a powerful complement to traditional spectrum monitoring, with capabilities such as automatic alerting, adaptive filtering, and dynamic decision-making based on spectral trends.

By coupling the battle-proven robustness of Signal Hound's Spike platform with state-of-the-art machine learning techniques, our system achieves a trade-off of interpretability and intelligence, enabling both human and machine agents to engage the RF environment in real time.

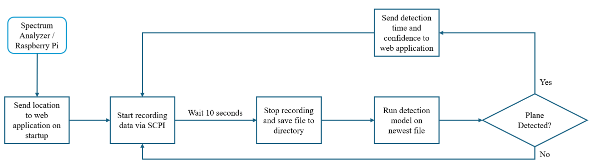

Fig. 2. Software flow diagram of detection program

D. Evaluation Metrics

The evaluation metrics for this system will focus primarily on detection accuracy specifically, the likelihood that the system correctly identifies the presence of an aircraft when one is present. In addition to accuracy, we will assess the system's effective detection range for both aircraft and radio communications. A third key metric will be detection speed, which includes the time it takes for the system to recognize and process a signal as a valid detection.

IV. Super-Resolution Optical Imaging

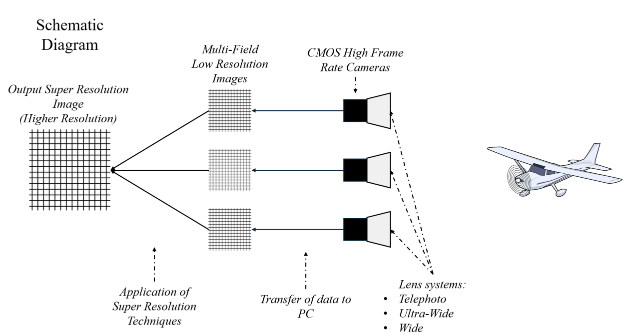

Fig. 3. Schematic diagram of camera array system and overview of the image processing pipeline

A. Motivation

The motivation for designing and including the camera array into the passive aircraft detection system was based on a suggestion by our sponsor who informed us about a project that was unsuccessful almost a decade ago that tried to leverage a similar super-resolution approach. The main appeal for our group was the leaps that super-resolution technology has made in recent years as well as the amount of optical design necessary for two optical engineering students.

Furthermore, systems that use both RF detection and optical verification that are commercially available like the TwInvis or Elvira [1,2] are far more complex than our system and cost magnitudes more. While the higher end systems do provide better performance and more valuable information, our approach achieves useful results with a budget below $250 in components and materials.

Lastly, the camera array meets our need to be stealthy. Since we are simply capturing images, it will be nearly impossible to tell that this system is even there if not completely impossible. ITAR restrictions were also met by keeping our design simple with the Arducam OV9281.

B. Camera Array Design

To best achieve super resolution imaging using a multi-camera approach, we developed a three-camera optical array that utilizes identical sensors and magnification configurations but with varying fields of view. This strategy allows us to use the slightly varying data from each camera to computationally fuse the information into an enhanced resolution image beyond the capability of any of the individual sensors. This array is demonstrated in Fig. 3.

Each of the varying fields of view is dictated by a two-lens configuration. Specifically, we employed a pair of singlet lenses, one with a positive focal point (plano-convex), and the other negative (plano-concave) to create a simple telephoto system. The specific focal lengths of these lenses were chosen to keep a constant magnification while changing the angular coverage in the systems.

The narrowest field of view system utilizes a 250mm focal length plano-convex lens paired with a -50mm plano-concave lens. The goal of this system is to capture the most fine detail as possible to build up the final image. The second, wide, system uses a 175mm/-35mm lens pair, and the third, ultra-wide, system uses a 125mm/-25mm configuration to achieve the widest field of view in the array. The purpose of all three of these systems is to provide more and more key identifying information of an aircraft such that the final image is a recognizable object.

The staggered field of views are critical for enabling high-quality image fusion in the super resolution algorithm. The overlapping information provides the necessary data to reconstruct blurry sections of the image. This optical design balances simplicity, cost-efficiency, and effectiveness. The result is a platform for super resolution imaging, most suitable for slow moving or static targets where high quality images are necessary for detection, like a far-away aircraft in the sky.

C. Image Processing Pipeline

The image processing pipeline begins with image acquisition. For our purposes, the images that are being captured should be simultaneous or as close to it as possible for super-resolution to be effective. For this, we use Python in Jupyter Notebook to detect all active cameras and then present a live feed to ensure your object is in the scene.

Next, we take those images and have to align them properly using an AI software to make sure the fields have enough matching features to be able to do super-resolution. This step is currently a work in progress but Fig. 4 is an example image of how the AI software determines matches between the different fields. The challenge here is finding the right software for our needs with the semester wrapping up soon.

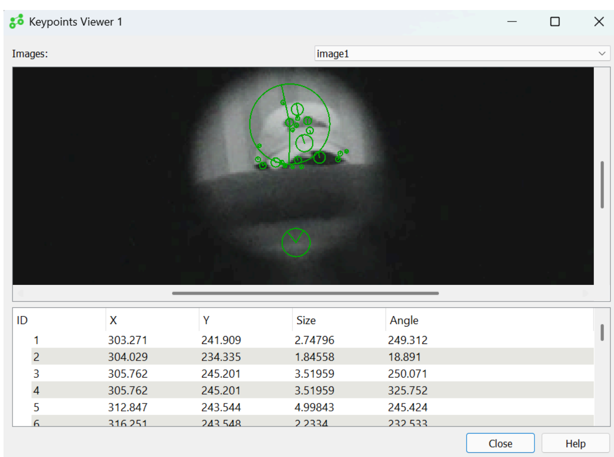

Fig. 4 AI software showing what it determines are distinguishable features

Once there are enough matching features between all 3 images, then the super-resolution algorithm can take place where we produce a singular, higher-resolution output. The super-resolution technique we use is back-projection since it works best with multi-image super-resolution. Another possible image fusion technique that could be used are those that use neural networks. Among these, ESRGAN, a relatively recent version of super-resolution, has been able to enhance images beyond what people could expect. Since ESRGAN is better, it unfortunately requires more data and the presence of a GPU, so for our purposes we will only use back-projection.

D. Atmospheric & Optical Corrections

While the current system does not offer any atmospheric or optical correction, it is important to note that the success of this system requires that there is an established line of sight, otherwise the optical system can offer no useful information. In our testing, we did notice that our cameras, while labeled monochromatic, had RGB values still associated. To work around this we decided to just convert everything to a true grayscale since without that step, it results in images that look like blue and red 3D. Atmospheric and Optical correction will be further discussed in the future work section.

E. Evaluation Metrics

The main evaluation metrics that we use for the camera system is resolution in the form WxH. While it is hard to define what successful super-resolution is, a noticeable upscale is what we want to see and is what holds the information that we find valuable. In addition, to remain consistent with our demonstrable goals, we will evaluate magnification physically and will show that image overlap must be near 100% for super-resolution to take place. Since our magnification and camera setup do not change, one evaluation of each will be sufficient.

V. Software & Backend Design

A. RF Backend

The backend of the web interface was developed using Node.js with the Express framework and MongoDB as the primary data store. It acts as the communication between the physical detection hardware and the frontend display, handling incoming aircraft signal detections and sensor location updates. RESTful API endpoints were created to support standard operations for signals and sensor data, while Socket.IO was used to enable real-time communication with the frontend.

Each time the RF detection system identifies a new signal, the data is posted to the backend and immediately emitted through a WebSocket to the frontend dashboard. Similarly, the sensor's latitude and longitude are submitted to the backend and broadcasted live to update the map view and coordinates displayed.

MongoDB was selected because of its flexibility with unstructured data and its ability to handle fast insertions of time-series information like detection events. The backend ensures data persistence, and if a user refreshes the page, all detections and the sensor status are displayed without loss.

B. Camera System Backend

The camera system backend consisted of Zemax for lens system design and testing, MATLAB for image processing, Python for simultaneous image capture, and an AI software for field alignment post image capture. While there may be a more concise way to execute what the system needs to do, image processing is central to the optical engineering part of the camera array so it was taken on by the student.

In addition to the image processing and image capture, as a way to tie the two sides of the project closer together, we developed an object detection CNN where we can submit the images of the object we captured and identify it.

C. Frontend

To support the passive RF detection, we developed a lightweight and responsive web interface that allows users to view live aircraft detection data in real time. The Frontend is built using React and is designed around a clean three-column layout. The center columns serve as the primary dashboard for displaying detected aircraft signals, complete with timestamps, confidence levels, and dynamic updates using Socket.IO for real-time interaction. The user flow for the application is outlined in Fig. 5.

The left column provides information about the physical location of the RF sensor, including a live Google Maps embed that automatically updates based on the coordinates transmitted by the sensor unit. Coordinates are formatted with directional indicators ("N" for north, "W" for west, etc) for clarity and professionalism. A color-coded legend is also included to help users quickly interpret signal confidence levels.

The third column is reserved for integration with a flight comparison system that will visually differentiate between known public aircraft and potentially unidentified signals. Here, we will indicate whether what we have detected exists on any publically available flight record.

All frontend components are styled for readability, system awareness, and constant real-time updates, making the interface suitable for research and deployment use.

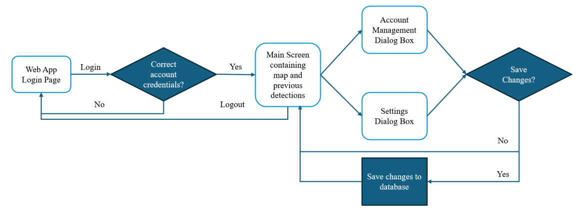

Fig. 5. User flow diagram of RF web application

VI. System Testing & Results

A. RF System Results

Our testing at MCO yielded impressive results in both machine learning performance and detection accuracy. We evaluated the system based on key engineering requirements critical to its practical use: detection time, range, and accuracy. The results showed strong performance across all three metrics, demonstrating the effectiveness of the machine learning algorithm.

Initially, the model performed poorly due to a lack of training data. However, after several months of data collection, we were able to train a model that could reliably detect aircraft at ranges from as close as 0.25 miles to as far as 4 miles. The dataset included approximately 200 data points per category aircraft present and aircraft absent which provided a solid foundation for experimental testing.

During training, we found that maintaining a balanced dataset was crucial. An imbalance would cause the model to either falsely detect aircraft in every sample or miss them entirely with high confidence. Our most successful model at MCO used around 160 data points with aircraft and 200 without, which resulted in the most consistent performance.

Interestingly, we observed that the model performed best at detecting aircraft farther away, assuming signal attenuation remained manageable. At those ranges, signal scattering became more pronounced and distinguishable, making it easier for the model to identify aircraft-related patterns. This setup achieved detection accuracies as high as 75%.

B. Camera Array Tests

For lens system testing, lenses were computationally chosen primarily from the equations for magnification, effective focal length, and field of view.

$$M = \frac{f_{positive}}{f_{negative}} \quad (1)$$

Lens system magnification

$$\frac{1}{f_{eff}} = \frac{1}{f_1} + \frac{1}{f_2} - \frac{d}{f_1 * f_2} \quad (2)$$

Lens system effective focal length

$$FOV = 2tan^{-1}\left(\frac{w}{2*EFL}\right) \quad (3)$$

Lens system field of view. w is the sensor width

These equations are shown in equations 1, 2, and 3. It was important for each lens system to have a constant magnification of 5 times so that the image overlap requirement of super resolution imaging was met. Additionally, the field of view needed to slightly vary for each lens system. This is so, in the super resolution algorithm, each image provides similar core information with a small amount of additional peripheral details.

After these equations proved the lens system we chose would work theoretically, we then had to test in simulations before ordering the parts. To do this, we conducted tests using Zemax software. For each proposed lens system, we simulated how the light would behave. These tests not just confirmed our handwritten calculations, but also provided us insight on how spaced out our lenses need to be from each other and how far away the sensor needs to be from the last lens.

For whole system testing, we captured and processed images of static targets such as a ThorLabs LabSnacks box, a coca cola can, and many more items. These tests were all conducted in a lab so the results may vary when tested outside and at farther distances. Though the first set of testing was done without the camera array holder, we managed to get some decent results, considering we were imaging relatively small objects at about 2 meters away, still not at the level we expect however.

After the testing, we made a key discovery in our image processing pipeline. Since our camera array system utilizes multi-field images meaning field of views change rather than position in relation to the object we realized we would need to use AI to be able to properly recognize similar features within the fields and align them using non-linear means. Further testing is necessary to make better use of the AI software and to produce better results overall.

C. Quantitative Results

RF system

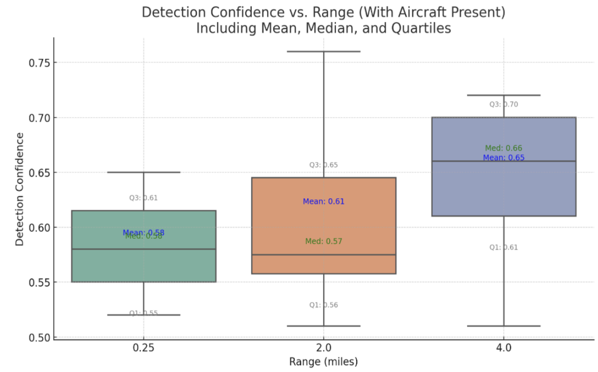

Fig. 6. Quantitative Results with Aircraft

From the quantitative data at 0.25 miles, we observe a mean confidence of 0.58 with an accuracy of 70%. At 2.0 miles, the mean confidence increases to 0.61 with an accuracy of 80%. Finally, at 4.0 miles, the mean confidence reaches 0.65, again with an accuracy of 80 percent.

As shown in Fig. 6, the mean confidence generally increased with distance, along with improved accuracy. This results in an overall accuracy of 76.66% and an average confidence of 0.64 for the case in which an aircraft is present.

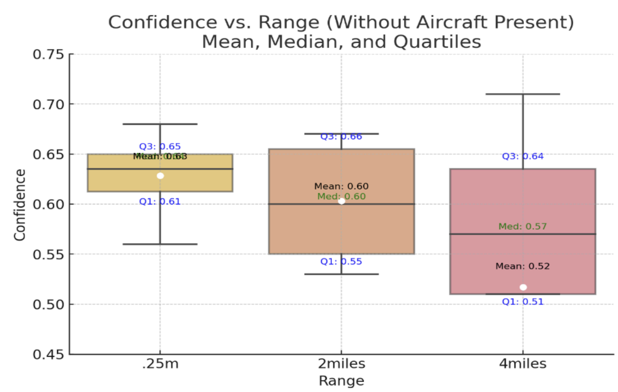

Fig. 7. Quantitative Results without Aircraft

At 0.25 miles, the mean confidence was 0.63 with an accuracy of 60%. At 2.0 miles, the mean confidence dropped slightly to 0.60, while accuracy increased to 90%. At 4.0 miles, the mean confidence further decreased to 0.52, with an accuracy of 80%.

As shown in Fig. 7, confidence values tended to decrease as distance increased. However, this trend is not necessarily causal. It's likely that the system's confidence scores were less precise when no aircraft was present, even though its predictions remained highly accurate. Additionally, due to limitations in our testing setup particularly that antenna orientation was not controlled across different ranges the confidence scores may have been affected, contributing to the observed decline.

Conclusion

From the quantitative results, we observed an overall accuracy of 75% and an average confidence of .63. These are strong results for our machine learning and hardware setup. All testing was conducted at MCO, demonstrating not only the proof of concept for detecting anomalies but also the model's ability to identify aircraft based on disruptions in the RF airspace. The results are promising and indicate that performance can be further improved with additional data and expanded testing.

Camera Array

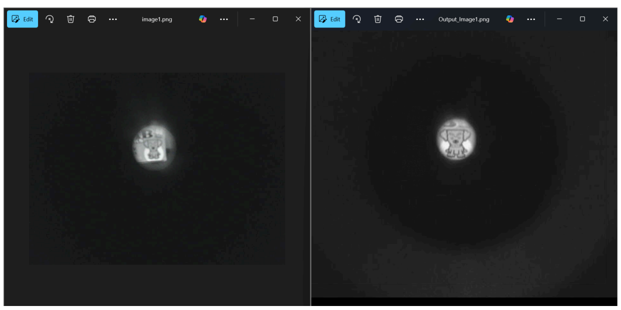

Fig. 8 Left image shows 1 low resolution input. Right image shows an upscaled image.

Fig. 8 shows one example where we were able to upscale the image on the right slightly, however, the reason AI needed to be introduced was so that the super resolution image would have the same field of view as the widest lens in the system. So, with the LabSnacks box cartoon, we did see an upscaling of the image and the resolution of the image did increase by 3x which is what we want but not how we want it. More work and research is going into figuring out which AI tool will work best for us. For now, this is the output we can achieve.

Since our camera array can be further enhanced with more cameras, we decided to store the numerical data pertaining to each image in an excel sheet, as shown in Fig. 9. For our purposes, the upscaling magnitude stays the same at 3x hence the consistency below. Should more cameras be added to the system, this value could be variable and could produce very highly upscaled outputs. Below is an example of how the spreadsheet is set up.

	A	B	C	D	E	F
1	Date & Time	Original Resolution	Nearest Neighbor	Bicubic	Back Projection	Upscale Factor
2	MIDTERM DEMO 1	800x600	2400x1800	2400x1800	2400x1800	3
3	MIDTERM DEMO 2	800x600	2400x1800	2400x1800	2400x1800	3
4	MIDTERM DEMO 3	800x600	2400x1800	2400x1800	2400x1800	3
5	2025-06-29 15:27:59	800x600	2400x1800	2400x1800	2400x1800	3
6	2025-06-29 16:23:27	800x600	2400x1800	2400x1800	2400x1800	3
7	2025-06-29 16:34:17	800x600	2400x1800	2400x1800	2400x1800	3
8	2025-06-29 16:43:48	800x600	2400x1800	2400x1800	2400x1800	3
9	2025-06-29 16:49:26	800x600	2400x1800	2400x1800	2400x1800	3
10	2025-06-29 16:54:11	800x600	2400x1800	2400x1800	2400x1800	3
11	2025-07-06 21:37:00	650x460	1950x1380	1950x1380	1950x1380	3

Fig. 9 Spreadsheet depicting numerical values and when the image processing took place

D. Limitations

For the RF system, the primary limitation was the equipment provided by the sponsor, as it was supplied at no cost to our group. This also influenced our design approach, which focused on the use of passive radar.

Passive radar depends heavily on the sensitivity of the receiving equipment. The signals we aim to detect are extremely low in power, sometimes as weak as -120 dBm. As a result, both the directionality of the antenna and the sensitivity of the hardware significantly impact system performance.

Another key limitation is the amount of data we are able to collect. Since aircraft flyovers are relatively rare events, accumulating a large and diverse dataset is challenging. This limits the training and performance of our machine learning model, which relies on substantial amounts of labeled data to improve accuracy.

For the camera array, there is still not a real-time super-resolution pipeline as it is technically still an area undergoing lots of research, though it has made leaps in the last few years. [3] Although we say affordability is one of the nuances of our project, that itself is a limiting factor and usually directly impacts sensor quality. Additionally, since the optical system requires having line of sight of the anomaly, weather and any other line-of-sight disturbances will be a major factor that will determine the success of the optical system.

VII. Discussion & Future Works

A. RF System

Our system has demonstrated promising results and capabilities. This is especially notable given its simplicity. With the integration of machine learning, the system successfully detects aircraft as well as other signals, including radio communications.

There are several opportunities for improvement. One enhancement would be enabling the system to determine a line of bearing to the target by comparing the amplitude of reflected signals from different angles. This would likely require some form of mechanical actuation to scan across a range of orientations.

Additionally, power optimization is essential to ensure reliable operation across diverse conditions and climates.

Another potential upgrade involves implementing a reference antenna. Traditional systems often use a secondary antenna to provide a known signal for comparison. By incorporating this approach, we could eliminate the need for mechanical movement. Instead, through signal processing and clever mathematical techniques, the system could estimate the target's range and bearing in a solid state configuration. This would represent a significant advancement over the current design.

Improved directional accuracy and range estimation would also allow the detection of smaller, more elusive targets, such as drones, increasing the system's utility across a broader range of real world applications.

B. Camera Array

While our system is able to show the capabilities of super-resolution techniques, there is much to be desired. Future improvements to the system such as adding more cameras for higher reconstruction capability, making more magnification lenses and holders, building multiple arrays to introduce different perspectives, and various other improvements on the software end. Since super-resolution is still a growing technology [4] we are also limited in the real-time processing we are able to do. Ideally, the system would be fully self-sufficient and could use adaptive optics, wave-correction plates, and other more professional grade equipment. Using higher grade equipment and mechanical assistance could also unlock the potential to have the ability to localize the aircraft in relation to the system.

Additionally, like the RF side of the project, there is room to expand the camera array side into the identification of aircrafts. This would also be done with artificial intelligence. With enough training, the AI could take the super resolution output image and determine key characteristics like plane type with a certain amount of confidence. [5] This could also be posted on the user interface so that the user is given more information on the aircraft flying into the detection zone.

Further work could even include anti-drone capabilities like the real-world products have. This would add another layer of interest to a potential user and allows them to control at the very least illegal drone activity.

VII. Conclusion

We successfully designed, assembled, and tested a low-cost, dual-modality surveillance system that combines RF anomaly detection with super-resolution optical verification. The system meets the key goals of affordability, scalability, stealth, and ITAR restrictions. Our results and future work can further demonstrate the strong potential for practical use in defense, environmental, and research applications. Future enhancements will focus on real-time centralized processing, localization, and deployment readiness.

Acknowledgement

The authors wish to acknowledge the sponsor of this project, Roy H. Wubker Jr. through the Department of Defense's Southern Command (SOUTHCOM). As well as our reviewers: Drs. Gong, Heinrich, and Kuebler.

References

[1] Hensoldt TwInvis Passive Radar System, Hensoldt, 2022.
[2] Thales Elvira Multisensor System, Thales Group, 2023.
[3] S. Park et al., "Image Super-Resolution: A Survey," IEEE Trans. Signal Process., vol. 20, no. 3, pp. 210–229, 2021.
[4] Z. Wang et al., "ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks," ECCV, 2018.
[5] R. Gonzalez and R. Woods, *Digital Image Processing*, 4th ed., Pearson, 2018.