RF Device Detection, Identification, and Secure Data Transmission using a Free-Space Optical Communications System

Zackary Zuniga, Bill Noel, Sarah Thibaut, Zachary Flores

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

***Abstract* — RF-enabled device detection is used in numerous applications today, from locating unauthorized transmitters to safeguarding wireless networks. Spectrum sensing with a software-defined radio observes emissions across crowded bands and environments. Our approach applies rule-based detection and feature extraction to reduce noise and isolate candidate emitters. Continuous I/Q streams are converted into nominal events via power/PSD tests and short-time analyses, then cross-checked with camera captures for verification. This paper's primary focus is a portable pipeline that combines an ADALM-Pluto SDR, Raspberry Pi control, triggered imaging, and Free Space Optical Communication with On-Off Keying telemetry for low-bandwidth reporting. While no single configuration dominated all settings, results show that fusing RF cues with visual confirmation reduces false positives and improves timeliness. The design and early tests presented here substantiate this approach.**

***Index Terms* — Free Space Optical Communication, passive radar, software-defined radio (SDR), geolocation, Image recognition, Light fidelity.**

I. Introduction

Conceived and developed to make practical use of several complementary modern technologies, the RF-Enabled Device Detection and Identification System is a complete, high-level electrical and computer engineering design project. The system, engineered for portable field operation, integrates an ADALM-Pluto software-defined radio and dual-band antennas for passive spectrum sensing, a Raspberry Pi 5 for orchestration and data processing, a GPS module for geotagging, and a machine-vision camera for event-triggered imagery. Custom Python software continuously acquires I/Q samples, computes power spectral density and short-time frequency features, and applies thresholding and rule-based detection to isolate candidate emitters across VHF/UHF and 2.4 GHz bands. When an RF event satisfies confidence criteria, the Raspberry Pi triggers the camera to capture corroborating images and optionally runs a lightweight embedded vision model to tag contextual cues (e.g., handheld radio, access point). A compact Free Space Optical Communication link using on–off keying (OOK) with a high-brightness LED provides a covert, low-bandwidth optical channel for discretely removing summaries of paramount data (time, frequency, geolocation, and visual tags) to a photodiode/ESP32 receiver for relay to a remote dashboard. Power management and thermal considerations enable battery-backed operation, while data handling includes timestamping, and structured JSON logs for downstream analysis. The originality of the system lies in its fusion of passive RF cues with visual confirmation and optical telemetry, yielding a low-cost, low-probability-of-intercept workflow that reduces false positives and accelerates operator awareness. From raw I/Q collection to event scoring, image capture, Free Space Optical Communication transmission, and remote visualization, the design demonstrates an end-to-end pipeline that transforms opportunistic emissions into actionable, georeferenced intelligence suitable for campus security studies, spectrum surveys, and contested-environment sensing.

II. System components

A. Processing & Control

The heart of the project is a Raspberry Pi 5 orchestrating RF capture, event logic, image capture, and Free Space Optical Communication transmission. An optional PYNQ-Z2 FPGA board is considered for DSP/ML acceleration, featuring on-board DRAM, block RAM, DSP slices, and a dual-core ARM for heterogeneous workloads. The chosen microcontroller unit for the receiver is an ESP32. This low-powered device is sufficient to run a demodulation algorithm to recover the signal sent through the Free Space Optical Communication link at varying lengths. Power requirements and supply options for both boards are noted for field use.

B. Software-Defined Radio (SDR)

RF sensing is performed with an ADALM-Pluto SDR. The design emphasizes narrow-band collection for higher sensitivity and lower processing overhead, with the option to switch bands or augment to two synchronized receivers for simultaneous multi-band capture and geolocation experiments. The ADALM-Pluto SDR is compatible with MATLAB and multiple spectrum analyzing software applications such as SDR Sharp. Our chosen SDR has open-source documentation for connection with the Raspberry Pi 5.

C. *Antennas*

Two monopole antennas are employed: a dual-band 2.4/5.8 GHz monopole for Wi-Fi/UAS links and a VHF/UHF (~470 MHz) whip for handheld radios. Using focused, band-appropriate antennas increase detection reliability while keeping the form factor compact; a matched-pair layout supports phased stretch goals.

RF sensing is performed with an ADALM-Pluto SDR. The design emphasizes narrow-band collection for higher sensitivity and lower processing overhead, with the option to switch bands or augment to two synchronized receivers for simultaneous multi-band capture and geolocation experiments.

D. Camera

Alongside the antenna, there is a CMOS camera added for image capture and object detection. The chosen antenna is an omni-directional type; therefore, direction is chosen with the CMOS. This CMOS has a 66° field of view. The field of view is the sensor's limiting range of sensing.

E. *Free Space Optical Communication Transmitter (Optical Emitter & Driver)*

For free-space optical exfiltration, the team is utilizing white LEDs over visible emitters due to cost, longevity, and temperature robustness. Representative IR parts span 850–1550 nm with 20–46 MHz bandwidths; the transmitter is driven by a simple OOK modulator/driver stage.

F. *Free Space Optical Communication Receiver (Photodiode & Front-End)*

Reception uses a photodiode matched to the LED wavelength, prioritizing adequate bandwidth and active-area for alignment tolerance. The report catalogs visible and NIR photodiodes suited to 850–940 nm or 1550 nm operation for reliable OOK detection.

G. *Optics*

Both the emitter and receiver employ optical elements (collimating/focusing lenses) to project and collect light efficiently. The report notes lens-design considerations and equations for collimating LED output; the enclosure allocates dedicated optical paths for the LED and cameras.

H. *Power System*

Field operation is supplied by a $LiFePO_4$ battery pack feeding regulated rails (5 V/3.3 V). $LiFePO_4$ is favored for cycle life, safety, flat discharge, and solar-charging compatibility, supporting long unattended deployments with stable voltage for compute, RF, and optical subsystems.

III. SYSTEM CONCEPT

The RF-Enabled Device Detection and Identification System integrates passive spectrum sensing, event-triggered imaging, and low-bandwidth optical telemetry to convert opportunistic emissions into actionable, georeferenced intelligence. The concept is organized around three cooperating subsystems: (i) an RF sensing path that detects energy and extracts features, (ii) an imaging path that verifies context on demand, and (iii) a communications path that exfiltrates compact event summaries over a covert Free Space Optical Communication link. A Raspberry Pi 5 orchestrates timing, buffering, scoring, and interprocess messaging across these subsystems.

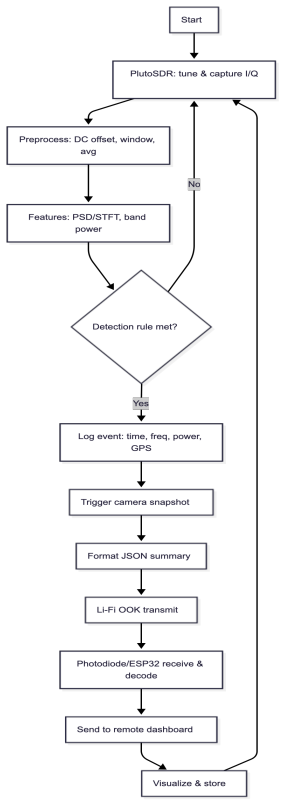

Fig 1. The system logic process from start to finish

Operational workflow. As shown in Fig. 1, the ADALM-Pluto SDR continuously captures complex I/Q samples in targeted bands (FRS/GMRS and 2.4 GHz Wi-Fi in the current prototype). Each buffer is preprocessed to remove DC bias and stabilize variance, then converted into frequency-domain features (e.g., Welch PSD, band-power, short-time statistics). A lightweight detection rule evaluates these features against adaptive thresholds to declare a candidate RF event. Upon a positive decision, the controller logs the event (time, center frequency, peak/band power, GPS fix) and triggers the camera to capture a corroborating image. The system packages a concise JSON summary and transmits it via on–off keying (OOK) over an infrared LED. A photodiode/ESP32 receiver decodes the payload and relays it to a remote dashboard for visualization and storage; the local node then resumes acquisition without operator intervention.

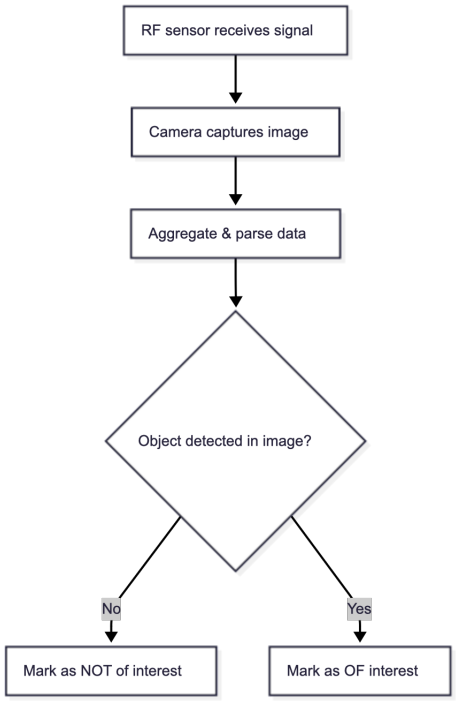

Fig 2. Object detection logic, including determination as to whether object is of interest

Decision logic. Fig. 2 summarizes the fusion rule used to focus operator attention. If the RF subsystem detects an emission but the image classifier (or simple object-presence test) does not confirm a relevant object in the scene, the event is marked **not of interest** and retained only for background statistics. If both RF and image cues are present, the event is marked **of interest** and promoted for alerting, map pins, or follow-up collection. This AND-style gating reduces false positives from benign emitters and environmental noise while preserving recall on targets that are both transmitting and visually present.

Interfaces and data artifacts. Internal processes communicate over in-memory queues and lightweight files:

- **I/Q buffers:** complex int16/float32 arrays from the SDR.
- **Features:** PSD vectors, band-power scalars, and detection scores.
- **Event records (JSON):** {timestamp, freq_hz, power_db, gps, img_uri, decision}.
- **Free Space Optical Communication frames:** short, checksummed payloads suitable for OOK at tens–hundreds of kb/s.

Performance goals and constraints. The architecture targets sub-second end-to-end latency from RF event to decision, with camera trigger and JSON serialization amortized below the SDR buffer period. Battery operation and thermal headroom constrain duty cycles for imaging and Free Space Optical Communication transmit bursts; threshold adaptation maintains sensitivity while limiting spurious triggers in crowded bands.

Security and reliability considerations. Optical exfiltration provides a low-probability-of-intercept channel, while signed JSON and sequence numbers mitigate spoofing and replay. GPS timestamps and image hashes support chain-of-custody for offline analysis. Watchdog timers on the controller and receiver recover from dropped frames or peripheral faults without operator action.

IV. SDR AND FREQUENCY DETAIL

Before the decision logic of the system can be explained in full, a clear understanding of the RF sensing portion is required. The objective is twofold: (1) detect narrowband handheld radios in the FRS/GMRS band, and (2) detect broadband Wi-Fi activity on 2.4 GHz channel 6. The ADALM-Pluto SDR is used for acquisition and visualization (SDR++), while a Raspberry Pi processes the samples for power-spectral features and thresholding.

A. RF Capture Configuration

The SDR is configured per target signal:

- **Handheld radio (FRS/GMRS).** Center frequency near **462.6875 MHz (Ch 6)**; narrow receive bandwidth (1 MHz) to focus on the FM carrier and reduce noise power. A moderate manual gain is used to avoid clipping at close range.
- **Wi-Fi hotspot (2.4 GHz).** Center frequency **2.437 GHz (Ch 6)**; receive bandwidth of 20 MHz and sample rate in the **MS/s** range to cover the full 20 MHz 802.11 channel. In SDR++, tests shown in Fig. 3 used manual gain of 10 dB.

Fig 3. RF Detection Spike due to turning on phone hotspot on SDR++

All levels shown are **dBFS** (decibels relative to the ADC full-scale), so absolute dBm is not implied; the noise floor in the shown setups sits roughly between −110 and −95 dBFS depending on bandwidth and gain.

B. Handheld Radio Characterization

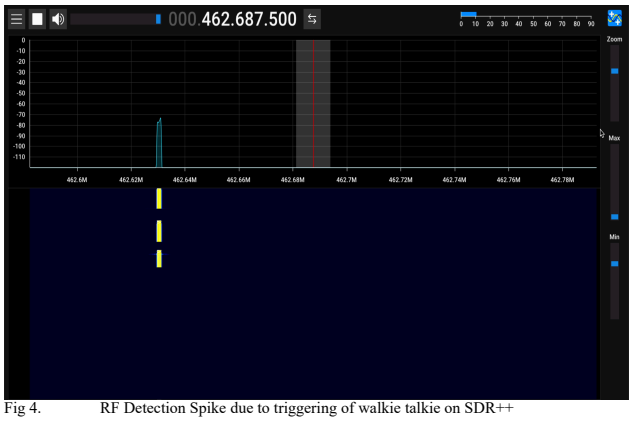

Fig 4. RF Detection Spike due to triggering of walkie talkie on SDR++

Spectral behavior. Fig. 4 above shows the instantaneous spectrum and waterfall when the push-to-talk (PTT) is engaged. A **narrow spike** appears at the tuned frequency with short **vertical bursts** in the waterfall corresponding to PTT activation. The spike's width is dominated by the receiver filter; the true emission is narrowband FM with energy concentrated at the carrier and nearby deviation.

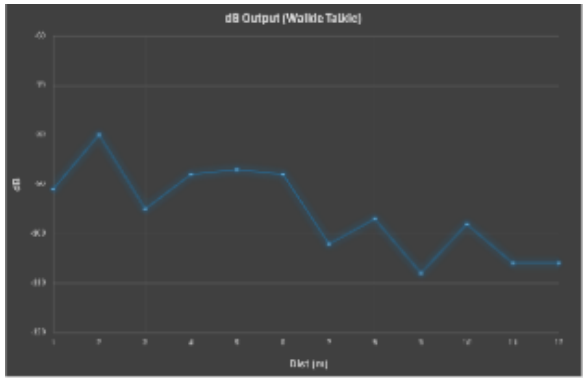

Fig 5. dB output for walkie talkie detection vs distance from SDR

Distance sweep. Fig. 5 above plots the received level vs. distance from the SDR. For each mark (1–12 m), the radio transmitted briefly while the receiver recorded the peak/average level over a short window. The overall trend is a decrease in received power with distance, consistent with log-distance path loss indoors. The non-monotonic dips and bumps(e.g., around 2–4 m and 7–10 m) arise from multipath and body/hand shadowing; small changes in position can cause constructive or destructive interference at the antenna. At very close range, AGC or front-end compression can also flatten peaks, which is why moderate manual gain is preferred for this sweep.

Interpretation. These observations justify a narrowband detector around 462.6875 MHz using band-power or peak tracking. In practice, we treat short, tall excursions above the local noise estimate as candidate "talk" events and hand them to the fusion logic.

C. Wi-Fi Hotspot Characterization

Spectral behavior. Fig. 3 shows the spectrum when a hotspot is enabled on channel 6. Unlike the handheld, Wi-Fi occupies a wide 20 MHz plateau with notches/peaks from OFDM subcarriers and pilots. Activity is quasi-continuous while the AP beacons and when traffic is present.

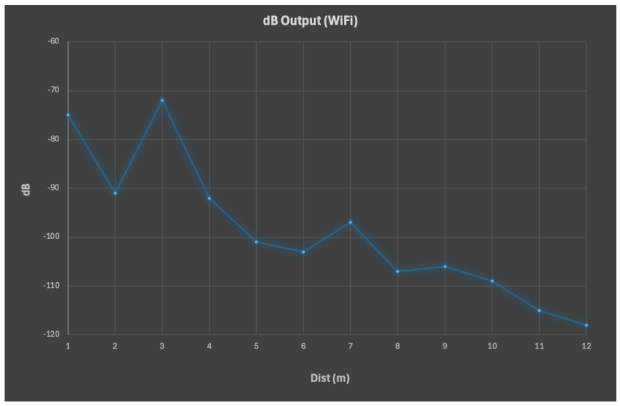

Fig 6. dB output for phone hotspot detection vs distance from SDR

Distance sweep. Fig. 6 plots received level vs. distance for the hotspot. The curve shows a clearer downward trend than the handheld but still exhibits ripples (notably near 7–10 m) due to indoor multipath and human movement. Because Wi-Fi is broadband, the measured level reflects the integrated power across the 20 MHz channel; using a wider measurement bandwidth raises the noise floor but better captures the total Wi-Fi energy.

Interpretation. These measurements support a broadband detector centered at 2.437 GHz that monitors integrated band power (or the median of the PSD across the channel) rather than a single-bin peak. Transient rises over the rolling noise estimate indicate the hotspot is active.

D. *Practical Notes and Limits*

1. **Bandwidth vs. SNR.** Narrow bandwidths (handheld case) lower the integrated noise and yield taller spikes; wide bandwidths (Wi-Fi case) increase noise power but capture the full signal.
2. **Gain strategy.** Moderate manual gain avoids clipping at close range and reduces AGC-induced variability between distance points.
3. **Indoor variability.** The non-monotonic behavior in both charts is expected indoors; repeating the sweep with small lateral offsets often shifts the ripples.
4. **Trigger thresholds.** For deployment, we compute a local noise floor (median/percentile over time) and trigger on excess band power with a short hold-time to avoid one-off glitches.

V. OPTICAL SETUP

For our optical design, we have a LED light source with a peak wavelength of 432nm that is collimated with lenses and focused onto a photodiode on the receiver.

When it comes to collimation of the LED, there is a tradeoff of power compared to distance. As the NA of the lens increases, the throughput increases as more photons from the LED are collected. Since NA and focal length are inversely proportional, as the focal length decreases, the divergence of the light increases. A decision of power compared to divergence must be made. We have chosen to have a lower divergence and therefore a small throughput of the LED.

Our LED is a Lambertian source, so it emits over a wide cone with a half-angle $\theta_1 \approx 15°$. The beam diameter after collimation with a lens of focal length f_1= 16mm is:

$$D_{beam} = 2f_1 tan(\theta_1) \quad (1)$$

therefore,

$$D_{beam} = 2(16mm)tan(15°) \approx 8.6mm \quad (2)$$

the collimated beam leaving the first aspherical lens is 8.57mm in diameter. We added another collimating lens that is planoconvex with an effective focal length of 150mm. This allows the distance between the planoconvex lens and the focusing lens to have adjustable spacing at different ranges, thus the light reaches farther distances without dispersing too much. As the light reaches the aspherical focusing lens, the spot size of the beam is focused onto our photodiode with an active area of 2mm^2. In order to achieve this, we used formulas to calculate the resulting spot size when the light reaches the end of the lens system. The geometric spot diameter, ignoring aberrations and diffraction:

$$D_{\text{spot}} = D_{\text{source}} \times m \quad (3)$$

including small divergence,

$$D_{\text{spot}}^2 \approx (D_{\text{source}}\, m)^2 + (2 f_2 \tan(\theta_{\text{res}}))^2 \quad (4)$$

where $m = f_2 / f_1$, is the imaging magnification, θ_{res} equals the residual divergence after collimation (≈ a few degrees since the LED is not exactly at f_1). When placing the focusing lens of focal length f_2 at distance d_2 from the collimator, the system roughly images the LED diameter onto the photodiode. The diagram below depicts our optical setup for the transmitter and receiver.

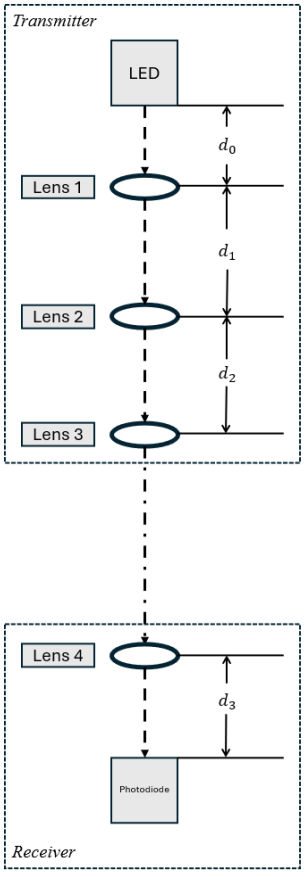

Fig 7. Optical Transmitter and Receiver layout (d0 = 4cm, d1=23.5cm, d2 = 17cm, d4 = 15cm.

The following table and graph display the data collected and recorded during initial tests of the power output emitted from our light source through our lens system. This test range starts at zero meters and goes up to forty meters.

	0	5	10	15	20	25	30	35	40
Dark Current	7.07	13.09	14.21	13.33	9.68	8.31	9.92	9.64	11.85
Power (μW)	11.09	16.98	17.39	16.48	12.49	10.25	11.25	10.69	12.26
Measured	4.02	3.89	3.18	3.15	2.81	1.94	1.33	1.05	0.41

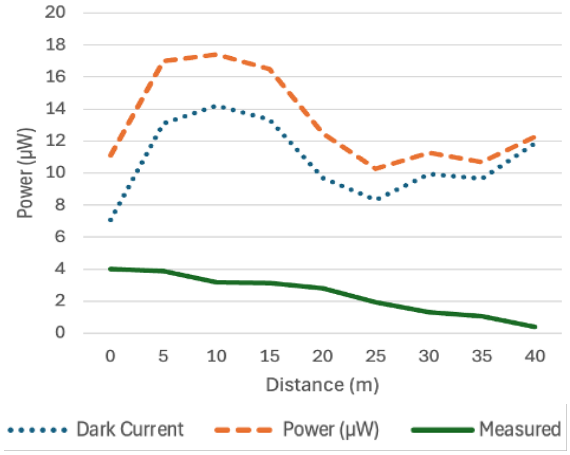

For testing the output power of the light at different distances, we used a power meter to detect our LED that is collimated through our optical setup. The dashed line represents the power recorded when the LED is turned on.

The dotted line represents the ambient light picked up by the power meter sensor while the LED is turned off. And finally, the solid line represents the measured difference between values when powered on versus off.

VI. VISUAL OBJECT DETECTION

The Object Detection system, relying on our ArduCam optical camera, is capable of capturing an image in a frontal 66° horizontal field of view (HFOV). The Object Detection works in a supplementary role to the RF detection, providing further information that can be used to mark collected samples as "Of Interest" or "Not of Interest" and ultimately for the image prediction data to be sent over the Optical Transmitter or retained on the device for further monitoring. The details of the device's machine learning implementation are detailed below.

A. Structure

For this system, the model chosen for the device's computer vision requirements balances the following factors:

1. **Ease of Training**: The model must be capable of being trained quickly, with minimal overhead. Various options used for training such as multi-scaling increase the time spent training depending on which model is used.
2. **Modern Operations**: Newer models with modern operations may improve achievable precision, whereas an older framework may be better suited for custom changes due to its simpler framework.
3. **Latency**: Measured in FPS, the model chosen should not be too slow. The FPS should be above two to give enough room to meet the system specifications. Smaller model size decrease latency and precision.
4. **Model Size**: The number of parameters must be low enough to fit on the Pi. If the model trained is too large, quantization must be used.

The model architecture used to meet the computer vision needs of the system is the Ultralytics YOLOv8, running on the Raspberry Pi 5. YOLOv8 was evaluated and chosen due to the speed and ease of training, modern operations, and support for both smaller models and optional quantization. The final model trained is the YOLOv8s variant, balancing FPS and accuracy with 11M parameters. YOLOv8s is then deployed on the Raspberry Pi 5 for inference with FP16 quantization.

B. Training

YOLOv8 for RF Detection was trained using public datasets [4] and has three classes: Drones, Cell Phones, and Routers. The pre-trained YOLOv8 weights used as a base are from training with the COCO dataset, and the cell phone class has images and labels mainly from the COCO dataset. To improve the diversity of data - mainly by providing more examples of modern phones - more images and labels were added. The drones and router classes were compiled from separate sets. To ensure correct training, any objects present in the disparate datasets but not labeled were given proper labels. The size of the entire dataset used is 8834 training images and 1157 validation images.

The final version is trained for 30 epochs using multi-scaling to learn the features of small objects such as phones and some aerial drone images. Then it is trained another 20 epochs at a reduced learning rate until mAP@50 improvements plateau at 89.7%, averaged across all classes.

C. Testing and Validation

The initial results for validating the precision of the Object Detection are obtained using a validation set of 1157 images containing 1293 object instances.

Class	Images	Instances	Box(P	R	mAP50	mAP50-95):
all	1157	1293	0.922	0.816	0.897	0.677
drone	347	369	0.96	0.851	0.947	0.597
cell_phone	471	576	0.834	0.716	0.794	0.563
router	329	348	0.972	0.882	0.95	0.872

Figure 8. Detection Results Table

The model broadly succeeds at detecting objects of each class, with the cell phone class as the lowest with an mAP50 of 79.4%. As seen below, the model is capable of detecting aerial drones with a small profile in the captured image. As object detail diminishes further away, confidence appropriately declines.

Figure 9. Example of Drone Prediction

VII. RECIEVER

As the receiver of the Free Space Optical Communication link collects the transmitter signal, it is critical to amplify the signal for demodulation. As we are using basic OOK modulation, the demodulation requires the input from the photodiode to have a significant difference between high and low values (high = 1, low = 0). As we are using the ESP32 as the microcontroller unit, we must amplify the small photodiode signal into a measurable input voltage.[3]

A. Transimpedance Amplifier

Therefore, we have designed a transimpedance amplifier capable of amplifying ultra-low photodiode inputs into measurable voltage inputs. In our design, we use an LTC6268 and the MCP6002 as our operational amplifiers to amplify the low photodiode signal. As the ESP32 requires a 3.3V input voltage, we have designated the LTC6268 and the MCP6002 to also use 3.3V. With the transimpedance amplifier design, we were able to achieve above 0.1mV of signal noise, with a 0.3mV increase at long ranges when a high value is detected. To increase the value, it is critical to increase the resistance of the resistor in the LTC6268 operational amplifier. Below is our transimpedance amplifier design, with our output voltage (VOUT) to connect to a GPIO pin on the ESP32.

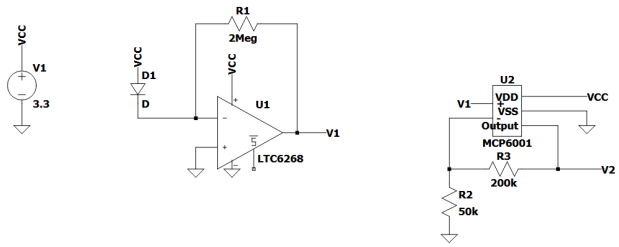

Figure 10. Transimpedance amplifier circuit diagram.

B. Signal Reception

It is critical to collect the waveform that the photodiode is capable of detecting at varying frequencies. This is entirely dependent on the responsivity of the photodiode and the bandwidth modulation of the LED. Therefore, different frequencies of a square waveform were generated through the LED driver and measured with an oscilloscope on the photodiode output with the amplifier.

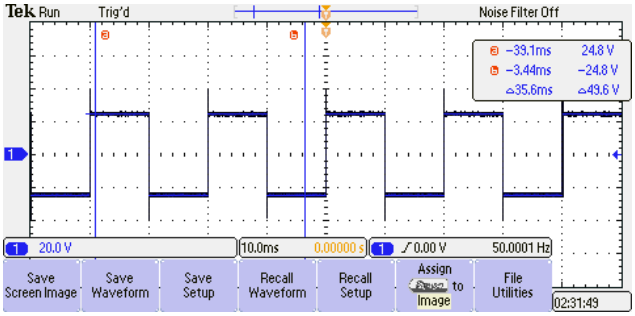

Figure 11. Optical Transmitter square waveform at 50Hz.

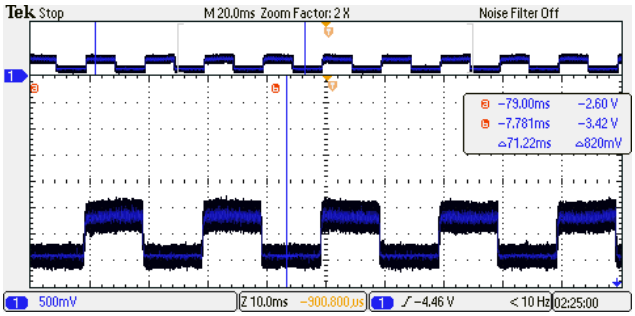

Figure 12. Optical Receiver square waveform at 50Hz.

As we increase the frequency, we can see a significant drop off in the edge of each cycle. To counteract this, we must implement a technique in the demodulation, limiting the time between each peak.

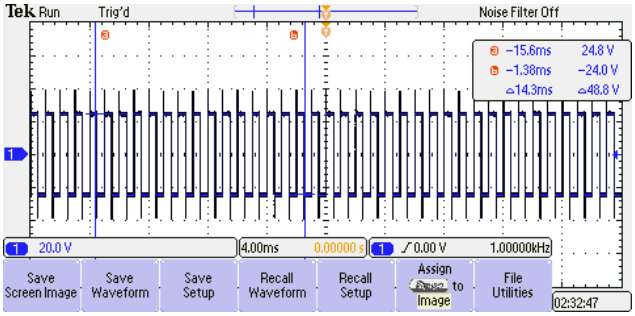

Figure 13. Optical Transmitter square waveform at 1kHz

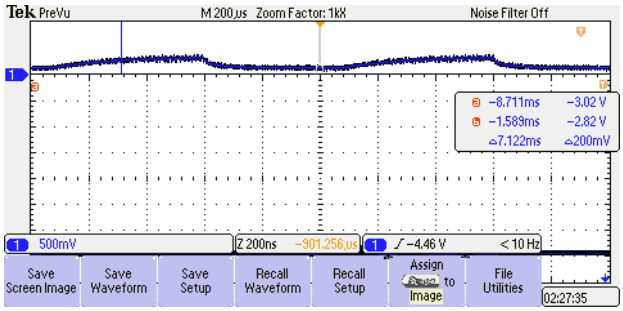

Figure 14. Optical Receiver square wave at 1kHz.

It is critical to amplify the signal as much as possible to increase the distances that a clean signal can be received. If the circuit cannot amplify the signal more, then more photons must be collected.

C. *Increasing Signal to Noise Ratio*

To increase the signal to noise ratio, first the electrical circuit must minimize as much noise as possible. This can be done through printing a circuit board rather than loose wires. Once the circuit board is optimized, then the optics may increase the photons captured. A larger NA lens may increase the photons captured at significant distances where the divergences become increasingly larger. If the diameter of the lens is larger than the size of the collimated light, then increasing the amperes going through the LED will increase the number of photons emitted, therefore more photons can be collected at the receiver.

VIII. Conclusion

This project delivers a small, low-cost system that listens for RF activity, snaps a confirming photo, and sends a short report over an infrared light link. A PlutoSDR and Raspberry Pi handle simple threshold-based detection, and we only flag an event when both the radio and the image agree—cutting false alarms while keeping latency low. The optical (Free Space Optical Communication) link transmits compact summaries to a remote dashboard without using additional radio spectrum.

Acknowledgement

The authors wish to acknowledge the assistance and support of Dr. Chan and Dr. Kar; University of Central Florida.

References

[1] T. Yucek and H. Arslan, “A survey of spectrum sensing algorithms for cognitive radio applications,” *IEEE Communications Surveys & Tutorials*, vol. 11, no. 1, pp. 116–130, 2009.

[2] H. Haas, L. Yin, Y. Wang, and C. Chen, “What is LiFi?,” *Journal of Lightwave Technology*, vol. 34, no. 6, pp. 1533–1544, Mar. 2016.

[3] P. M. E. Vazquez, F. Verias, L. C. Brazzano, P. A. Sorichhetti, “Design method of optical detection systems based on transimpedance amplifiers,” *Rev. Sci. Instrum.,* vol 92 , no. 11. Nov. 2021

[4] Datasets Used:
Phones:https://universe.roboflow.com/exam-detection-a9bsf/mobile-phone-detection-mtsje
Phones: https://cocodataset.org/#home
Drones:https://www.kaggle.com/datasets/dasmehdixtr/drone-dataset-uav/data
Routers:https://universe.roboflow.com/demo-gu3hz/router-juqmt