2HAD-RAD: Portable Life Detection Radar

[**H**itt, **H**ankinson, **A**ndersen, **D**'Souza – **Rad**ar]

Group 32 Authors:

Ryan Andersen	Bailey Hitt	Shannon Hankinson	Joshua D'Souza
Electrical Engineering	Electrical Engineering	Electrical Engineering	Computer Engineering

Review Committee:

Dr. Masoumeh Kalantari-Khandani	ECE	Lecturer
Dr. Enxia Zhang	ECE	Assistant Professor
Dr. Mark Maddox	ECE	Instructor
Dr. Mike Borowicz	ECE	Associate Professor

Advisor:

Professor Sreeram Sundaresh, Ph.D.

Department of Electrical and Computer Engineering, University of Central Florida

EEL4914: Senior Design I

August 28, 2025

UCF | College of Engineering and Computer Science

UNIVERSITY OF CENTRAL FLORIDA

Table of Contents

List of Tables

List of Figures

CHAPTER 1

1.1. Executive Summary

CHAPTER 2

2.1. PROJECT BACKGROUND AND MOTIVATION

In the aftermath of disasters such as earthquakes, structural collapses, or fires, the most critical priority is locating and rescuing survivors as quickly as possible. Traditional search and rescue techniques rely heavily on human senses, trained canines, and thermal imaging systems. While these tools have saved countless lives, they each present significant limitations. Canine units require extensive training and can experience fatigue or environmental disorientation, while thermal cameras alone are less effective in smoke-filled spaces, under debris, or when heat signatures are blocked. Similarly, acoustic detection techniques may fail in noisy post-distortion environments. These constraints underline the need for advanced, multimodal technologies that can complement conventional tools and improve the speed, reliability, and safety of rescue operations.

Recent advances in sensing technologies have enabled the development of compact, low power devices capable of detecting human presence and activity in challenging environments. Radar based sensing is particularly promising for search and rescue applications because it is unaffected by lighting conditions, operates effectively in smoke or dust, and offers the capability to detect micro-motions associated with life functions, such as chest movement from breathing. However, radar is not always sufficient, as complex environments and clutter can obscure signals or cause false positives. To address these challenges, integrating radar with complementary sensing methods can significantly enhance the detection of confidence and situational awareness.

This project proposes a handheld device that can detect humans using a combination of radar, ultra-wideband (UWB), infrared (IR), inertial, and barometric sensors. Each sensing modality provides a distinct type of information that enhances the reliability and robustness of human detection. Radar serves as the primary detection mechanism, while UWB aids in localization and ranging. IR provides thermal contrast in accessible scenarios, and barometric data supports altitude estimation indoors or underground. If time permits, the addition of seismic sensing allows the system to detect low frequency vibrations from footsteps, tapping, or shifting debris, even when line-of-sight is not available. Depending on mission requirements, the seismic capability can be implemented using a low noise MEMS accelerometer integrated into the handheld unit or a dedicated geophone module connected externally for higher sensitivity.

By integrating these sensing technologies into a single, portable system and applying advanced digital signal processing (DSP) and sensor fusion techniques, this project aims to enhance human detection in complex rescue scenarios. The system is designed to function primarily in line-of-sight conditions but includes a stretch goal of achieving limited through material detection, expanding its utility in real-world emergency response operations. Ultimately, this approach seeks to support first responders with a tool that is faster, more reliable, and more versatile than current solutions.

2.2. PROJECT FUNCTION

The proposed system is a handheld, battery powered device designed to support first responders in locating survivors during disasters such as earthquakes, building collapses, or fires. The device integrates multiple sensing modalities such as a millimeter-wave radar, ultra-wideband (UWB), infrared (IR), and an inertial navigation system (INS) with barometric sensing, to provide complementary data that improves detection accuracy and situational awareness.

The radar subsystem is the primary detection mechanism, capable of capturing micro-motions associated with human activity. High frequency radar signals penetrate smoke and airborne debris and reflect moving targets, allowing the system to detect subtle motions such as breathing. The signal returned undergoes analog signal conditioning and digitization before being processed with digital signal processing (DSP) algorithms to extract motion patterns and reduce environmental noise and clutter. The UWB subsystem enhances the radar by providing precise ranging and localization capabilities. UWB signals enable accurate distance measurements and tracking of movement, improving spatial awareness in both indoor and outdoor environments. This is particularly valuable in collapsed structures or underground areas where GPS signals are unavailable.

The IR subsystem provides thermal imaging data to highlight human presence where line-of-sight is available. Unlike radar, which operates independently of light and temperature, IR sensing utilizes temperature contrast to confirm the presence and approximate location of a human target, particularly in smoke-filled or visually obstructed environments. The INS and barometric subsystems add spatial context by tracking the device's movement and altitude changes, enabling responders to maintain positional awareness as they navigate multi-story buildings or subterranean environments. When paired with a global navigation satellite system (GNSS) receiver, the system can further improve positional accuracy by referencing absolute geographic coordinates when satellite signals are available.

Sensor data is fused in real time within the microcontroller unit (MCU), which runs DSP pipelines, sensor calibration routines, and data fusion algorithms such as Mahony or Madgwick filters for attitude estimation, and extended Kalman filters (EKF) for pose estimation and sensor integration. The device assigns confidence scores to detections by cross validating radar, thermal, and ranging data, reducing false positives, and improving detection reliability.

A user interface (UI) on the handheld device displays detection confidence, relative location, and activity levels, while a companion Android application communicates over Bluetooth Low Energy (BLE) to visualize detections and mark victim locations on a map. This connectivity enables rapid sharing of situational data among rescue teams. The device is powered by a rechargeable battery system, ensuring portability and continuous operation in field conditions.

While the primary objective is accurate line-of-sight detection, the system is designed with scalability in mind. Future iterations will explore through material detection, extending the device's effectiveness to scenarios involving drywall, wood, or light rubble. By uniting multiple sensing technologies into a single platform, this project aims to provide first responders with a powerful, field deployable tool that enhances search and rescue operations in time-critical environments.

2.3. Current Commercial Technologies and Existing Projects

2.3.1. Portable Sign of Life Identifier (Department of Homeland Security)

The Portable Signs of Life Identifier was a research effort funded by the Department of Homeland Security (DHS) focusing on evaluating compact and portable tools that could detect subtle respiration and heartbeat signals through walls or rubble [7]. The project explored radar and acoustic sensing methods capable of identifying minor chest movements caused by breathing or cardiac activity. The study emphasized portability, rapid deployment, and resilience to environmental interference, aligning closely with the design objectives of the 2HAD-RAD system.

While the DHS report demonstrated that the radar is capable of successfully detecting vital signs using Doppler radar, it also discussed limitations of the system such as the high-power demand, environmental noise, and false positives triggered by non-human motion. The 2HAD-RAD project builds on these findings by leveraging more power efficient embedded radar hardware, modern signal processing techniques, and a multi sensor approach to reduce false detections. By integrating radar data with infrared and inertial sensors, the system aims to confirm that detected motion correlates with human body temperature, thus improving detection confidence in cluttered environments.

2.3.2. Detection of Presence of Life Through Walls (DHS Research Report)

The DHS publication Detection of Presence of Life Through Walls expanded on earlier work by analyzing prototypes based on radar in urban search and rescue scenarios [8]. These prototypes used continuous wave (CW) and pulse-based radar to detect motion and vital signs through common building materials. The study highlighted several engineering requirements critical for field deployment, minimizing false alarms, maximizing detection range, and ensuring regulatory compliance with FCC standards.

Drawing from these insights, 2HAD-RAD explores ultra-wideband (UWB) radar as the preferred sensing technology. UWB transmits extremely short electromagnetic pulses across a wide frequency band, offering high spatial resolution and the ability to penetrate most nonmetallic materials [9]. This makes it ideal for locating survivors in collapsed structures or obscured environments. Whereas earlier DHS prototypes were bulky and power intensive, the 2HAD-RAD system integrates a compact radar front end with an ESP32-WROOM microcontroller, enabling local processing, reduced power draw, and improved portability. This shift from research grade instrumentation to a fully embedded system supports field ready use while maintaining technical rigor.

2.3.3. Integration of Multi Sensor Detection Technologies

Multiple academic and industrial studies have explored multi modal life detection by combining radar with other sensing domains such as infrared (IR), and acoustic. Multi sensor fusion improves detection reliability by verifying that observed movement correlates with a valid heat or vibration signature. The 2HAD-RAD system adopts a similar sensor fusion strategy, merging radar data with IR temperature readings and barometric or inertial feedback. This approach enhances both accuracy and interpretability, ensuring that life detection outputs are based on multiple physical indicators rather than radar reflections alone.

Unlike earlier systems relying only on radar, which required external computers for visualization, 2HAD-RAD integrates a dedicated LCD interface to display immediate feedback such as “Human Detected” or “Area Cleared”. The system’s design philosophy prioritizes speed, portability and user clarity, which are essential in high stress and time sensitive emergency operations. The inclusion real time signal filtering and sensor fusion to differentiate human activity from background motion further improves performance by distinguishing human respiration from background motion, offering a modernized and efficient evolution of earlier prototypes.

2.3.4. Commercial Through Wall Radar Systems: Camero Xaver and Vayyar Imaging

Several commercial radar solutions demonstrate the capabilities of current industry technology. The Camero Xaver 100 and 400 series, developed in Israel, are professional grade through wall imaging radars that use UWB impulse radar to detect movement and breathing through concrete and other solid barriers [11]. These systems offer high accuracy and real time visualization but are limited by large form factors, high cost, and significant power consumption, making them impractical for compact or battery-operated systems.

Similarly, Vayyar Imaging produces radar on chip modules based on frequency modulated continuous wave (FMCW) and UWB techniques that provide 3D imaging and real time presence detection [12]. While Vayyar modules are smaller than the Xaver systems, they rely on proprietary signal processing and high-speed computing, increasing integration complexity for lightweight embedded designs.

By integrating a compact radar module with thermal and inertial sensors on a single custom PCB, the 2HAD-RAD system delivers the same core capabilities as large commercial life-detection platforms, such as identifying human presence through barriers, but in a far smaller and more power efficient design. Its lightweight, affordable cost architecture makes it well suited for search and rescue kits, field testing, and educational applications, bridging the gap between complex commercial systems and practical embedded solutions.

2.3.5. Infrared Assisted Radar Concepts

Recent academic research has explored combining radar with infrared thermography to enhance detection accuracy, particularly in environments with high clutter or minimal motion [13], [14]. These hybrid systems use radar to detect movement and IR sensors to verify heat signatures corresponding to living organisms. Studies have shown that thermal data can significantly reduce false positives caused by moving objects such as fans, machinery, or debris.

The 2HAD-RAD project adopts this hybrid principle by incorporating an IR temperature sensor (MLX90640) alongside the radar subsystem. The IR module provides contactless temperature verification, confirming that detected motion corresponds to human body heat (typically 30 - 40 °C). This design choice improves performance in both open and obstructed fields of vision, while maintaining a compact and portable size. By integrating IR, radar, and inertial data within a unified embedded framework, 2HAD-RAD effectively combines vital sign detection, thermal validation, and environmental context for higher reliability than single sensor designs.

2.3.6. Summary

Both government and commercial studies confirm that detecting human life through walls is achievable but technically complex. Early DHS projects proved radar could measure vital signs, while also revealing the challenges of interference and high-power use. Meanwhile, modern commercial systems like Camero Xaver and Vayyar Imaging demonstrate impressive capability but remain expensive and bulky.

The 2HAD-RAD project takes these lessons a step further by merging radar, infrared, and inertial sensing into a compact embedded design. It delivers a handheld, low power, and affordable life detection system that’s both practical and field ready. This multi sensor approach not only enhances accuracy but also makes the technology more accessible to academic researchers and first responders, representing a meaningful step toward the next generation of rapid, portable life detection devices.

2.4. PROJECT OBJECTIVES AND GOALS

The objective of 2HAD-RAD is to design and build a portable, integrated multi-sensor device that assists rescuers in detecting and localizing human presence during emergency and disaster response. Radar serves as the primary sensing modality, supported by complementary technologies such as UWB, infrared, and inertial/barometric sensing. Research and manufacturer demonstrations have shown that radar can detect subtle life signs, such as breathing or even heartbeat, though this project emphasizes broader line-of-sight detection, with limited through-material capability considered as a stretch goal. The device combines sensor fusion, modern signal processing into a unified system, delivering rapid and reliable feedback in real time under field conditions.

2.4.1. GOALS

- Basic:
  - Integrate a commercial radar module to detect motion and respirations detection in open space.
  - Confirm human presence by detecting body heat signatures using an infrared sensor.
  - Provide relative location and orientation tracking using an INS with barometric feedback.
  - Display detection results on an LCD interface providing immediate visual feedback.
  - Power the entire device with a rechargeable battery subsystem that ensures stable operation.
- Advanced:
  - Maintain performance in smoke-filled or cluttered environments.
  - Improve relative location tracking precision by using GNSS.
  - Extend radar detection by integrating a UWB sensor for sensor fusion.
  - Optimize the power delivery for extended operational runtime.
  - Enhance LCD readability with clear formatting and detection indicators.
  - Transmit data from a detection device to an android mobile application.
- Stretch:
  - Extend radar detection for stationary human targets in open-air testing.
  - Incorporate a machine learning Layer to improve classification accuracy,
  - Demonstrate detection capability through common building materials.
  - Add user-adjustable controls to the UI to modify radar sensitivity, detection modes, and display a simplified thermal heat map.
  - Enhance data transmission to the Android application in real-time and include positional mapping.
  - Integrate a battery monitoring subsystem that displays voltage, charge percentage, and estimated runtime on the LCD.
  - Incorporate a hazard awareness layer that flags potential risks using fused sensor data.

2.4.2. OBJECTIVES

- Basic:

  - Demonstrate radar-based detection of a human subject within a line-of-sight range of at least 2 meters.
  - Confirm human presence using an infrared temperature sensor under line-of-sight conditions at distances up to 1 meter.
  - Provide relative location using INS and barometric sensing with an accuracy of ≤ 2 meters.
  - Display detection feedback on an LCD interface within 5 seconds of confirmed detection.
  - Add a red LED to indicate when a human is detected in less than 1 second.
  - Incorporate buttons to mark the location of a victim, clear/reset, and power the device.
  - Operate the complete device on a rechargeable subsystem for at least 2 hours of continuous use.
- Advanced:
  - Demonstrate consistent radar detection performance in cluttered or noisy environments and maintain ≥ 90% detection accuracy.
  - Improve fused localization accuracy within ≤ 1 meter when GNSS data is fused with INS outdoors.
  - Extend radar detection by fusing UWB data to detect stationary targets beyond 3 meters.
  - Enhance LCD readability in field conditions to ≥ 90% legibility under ambient light with clear indicator labels.
  - Transmit data from the handheld device to an android mobile application through BLE within 5 seconds.
  - Improve operational runtime by at least 25% compared to baseline while maintaining ≤ 10 W power draw.
- Stretch:
  - Demonstrate detection of human presence through 1 inch of drywall material in controlled indoor tests with ≥ 80% accuracy.
  - Incorporate a hazard awareness layer that flags 3 potential risks using fused sensor data with ≥ 80% detection accuracy.
  - Add user adjustable controls to UI for radar sensitivity, detection modes, and heat map visualization and update display within 5 seconds of adjustment.
  - Add real-time battery monitoring to charge percentage and estimated runtime on the UI.
  - Enhance android application with real-time positional mapping with ≤ 2 second latency.
  - Detect ground vibrations (knocks or footsteps) within a 2-meter radius using a seismic sensor.

2.4.3. REQUIREMENT SPECIFICATIONS

Parameter	Description	Target Value
Radar detection range	Detect a human subject in LOS	≥ 2 m (LOS)
Detection accuracy	Overall detection performance for radar and IR.	≥ 90%
Average system power	Average electrical power during continuous scanning	≤ 10 W
Battery life	Continuous operation time on rechargeable battery during scans.	≥ 2 hours
UI response	Time for LCD and buzzer/LEDs to respond after confirmed detection.	≤ 5 second
Infrared detection	Confirm human presence by detecting body heat signatures.	within ≤ 1m LOS
Location	Track victims' location when detected using INS, GNSS (when applicable) and barometer.	≤ 2 m

Table 1: Overall System Specification

Component	Parameter	Description	Target Value
Radar module	Integration latency	Time from RF transmission to usable baseband output	≤ 200 ms

Component	Parameter	Description	Target Value
Radar module	Detection range	Detect human motion under LOS conditions	≥ 2 m
IR sensor	Human heat confirmation	Detect knocks/footsteps through materials	≤ 1 m
IR sensor	Interface and conditioning	Digital bus, add small series resistors for edge damping.	I2C/SPI with 22-47Ω series, proper decoupling
INS with barometer	Relative localization	Dead reckoned displacement over short paths	≤ 2m error over 10 m
GNSS (advanced)	Fusion anchor	Absolute position when available	≤ 1 m
UWB (advanced)	Ranging precision	Sensor fusion with radar to extend range	≥ 3 meters
LCD	Feedback latency	Time from confirmed detection to display	≤ 5 s
LED	Turn on latency	Visual alert	≤ 1s
Wireless	Telemetry latency	From detection device to mobile application	≤ 5 s
MCU	Logic and supply	Core voltage and primary input/output for links and wireless	3.3 V, processor class at least 240 MHz
MCU	Interfaces	Serial links for INS, IR sensor, barometer, ranging for radio and telemetry.	at least 2 universal asynchronous receiver-transmitter links, 1 I2C and 1 SPI interface bus
MCU	Processing/signal processing	Filtering, detection logic, sensor fusion, and UI updates	Display update within 5 s and LED within 1 s

Table 2: Component Specifications

2.5. RELATED STANDARDS

Our team anticipates conforming to the following national and/or international standards:

- FCC Part 15: including but not limited to subparts A, B, C and F

- IEEE 802.15.4: UWB communications standard
- IEC 62133 / UL 2054: Li-ion battery safety
- ETSI EN 305 550: European equivalent to FCC Part 15 to cover 57-64 GHz short-range devices
- ESTI EN 302 065: Covers 3.1 - 10.6 GHz band.
- USB-IF USC Type-C and Power Delivery standards
- IEC 62368-1/UL 60950 for DC input safety
- IEEE 52-2019: Standard letter designations for radar-frequency bands
- IEC 62366: Application of using engineering for medical devices.
- API 35 for new android applications
- Build OWASP MASVS (Mobile application security verification standard)
- Follow WCAG 2.1 level AA

2.6. PROJECT BLOCK DIAGRAM

Figure 3 illustrates the signal flow and major subsystems of the proposed handheld human detection device. The design integrates radar, UWB, IR, inertial and barometric sensors into a unified system, all managed by a central MCU. The figure represents a single PCB implementation where sensing, processing, user interface, and power management are consolidated into a compact and portable platform.

At the core of the system is the TI IWR6843AOP radar module, which transmits and receives a 60 GHz frequency modulated continuous wave (FMCW) chirp. Reflected signals return as baseband I/Q data, which pass through an analog to digital converter (ADC), a hardware accelerator, and a C674x DSP core for low level signal processing. The radar has the capability of detecting human micro-motions such as breathing or small body movements, even in environments with low visibility caused by dust, smoke, or debris. Radar outputs are then sent to the ARM R4F control core and passed to the MCU for higher level processing, fusion, and decision making.

The sensors subsystem complements the radar by integrating additional modalities that improve detection accuracy and situational awareness. To expand the radar’s sensing capabilities and strengthen overall system performance, the device incorporates an UWB subsystem, which is specifically chosen for its ability to complement radar data by improving ranging precision and providing additional spatial information. Previous research and industry solutions have demonstrated that combining UWB with radar improves localization and reduces uncertainty in complex environments. In this project, UWB measurements will help refine detection zones and can aid in limited through material sensing, extending the radar’s effectiveness in cluttered or partially obstructed scenarios. The IR sensor provides thermal contrast, confirming the presence of humans when line-of-sight is available and highlighting heat signatures even in visually obstructed environments. The INS, combined with a barometric pressure sensor, enables inertial navigation by tracking device orientation, motion, and elevation changes, critical for maintaining positional awareness in multi-story buildings or underground areas.

All sensor data converge at the MCU, which serves as the central processing hub. The MCU performs multiple roles: it coordinates data acquisition from each sensor, applies DSP algorithms and filtering, and executes sensor-fusion logic to combine inputs into a unified detection decision. Algorithms such as Mahony or Madgwick filters are used for attitude estimation, while extended Kalman filters (EKF) may be implemented for fusing IMU, UWB, and barometric data into a coherent position estimate. Detection of confidence scores is generated by cross validating radar, thermal, and ranging data to reduce false positives and improve decision reliability.

The user interface subsystem provides intuitive feedback to the operator. A display presents key metrics such as detection of confidence, relative location, and sensor activity. LEDs and a buzzer provide quick visual and audible alerts when potential human presence is detected. Additionally, BLE connectivity links the device to an Android application, which visualizes detection data, marks victim locations, and enables rapid communication among response teams.

The power subsystem is built around a rechargeable lithium-ion battery, supplying regulated 3.3 V power to all modules through a buck converter and low noise linear regulators. The power circuit includes charging management and protection features to ensure safe and continuous operation in the field. Finally, the division of labor table shows how subsystem responsibilities are distributed among team members according to their expertise. This structure allows specialized focus on radar processing, sensor integration, power delivery, user interface design, and communication modules, while enabling collaborative development and integration of the full system.

Figure 2 (software block diagram) further illustrates the flow of information within the device. The MCU centralizes data from all sensing subsystems and processes it through layers of filtering, calibration, and fusion. LED indicators respond to radar detection results, and combined sensor confidence outputs. Meanwhile, localization and movement data from the INS, UWB, and barometric sensors feed into the mapping interface, which updates the user's display and Android application in real time. Together, these hardware and software layers enable a coordinated, multi-modal detection system designed to improve the speed, accuracy, and reliability of search and rescue operations.

2HAD-RAD Single PCB Block Diagram

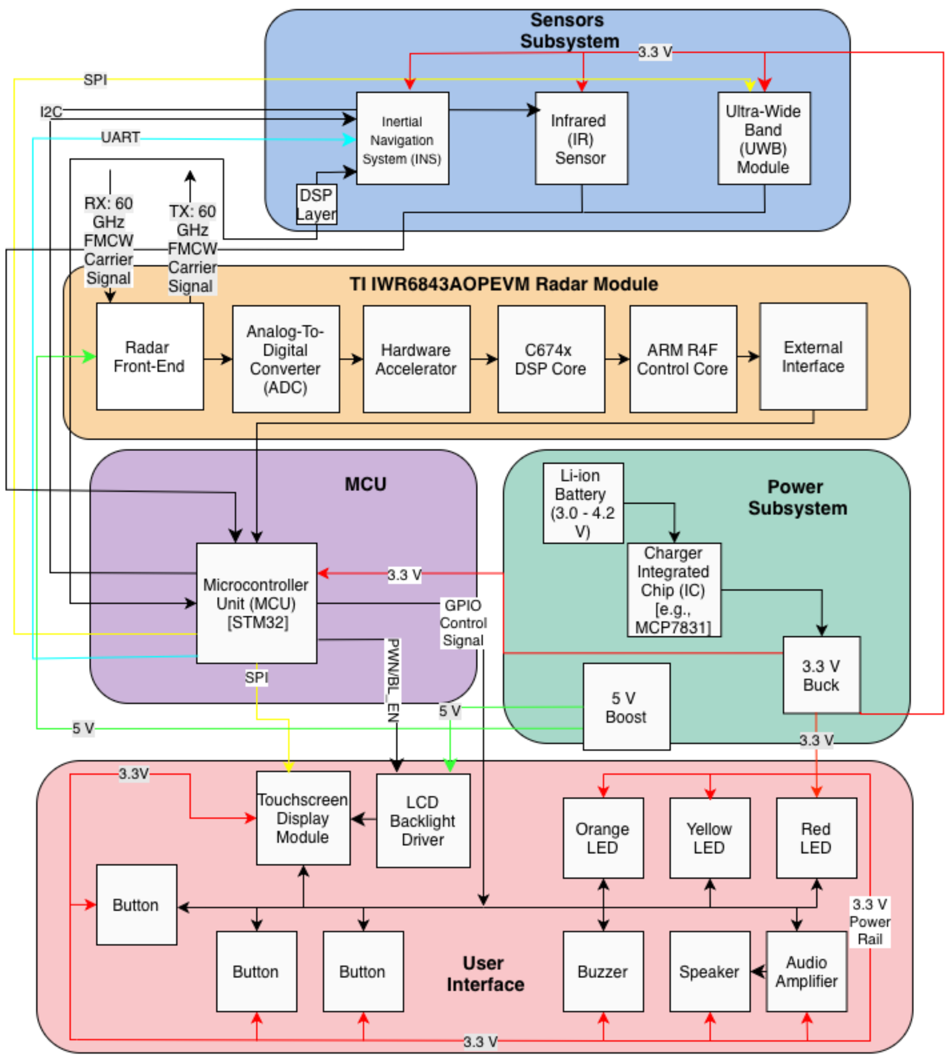

Division of Labor

Subsystem	Primary Lead	Secondary Support
Power Subsystem	Bailey Hitt (EE)	Shannon Hankinson (EE)
Sensors Subsystem	Ryan Andersen (EE)	Bailey Hitt (EE)
Radar Module	Shannon Hankinson (EE)	Ryan Andersen (EE)
MCU	Joshua D'Souza (CpE)	Bailey Hitt (EE)
User Interface	Joshua D'Souza (CpE)	Ryan Andersen (EE)
Radio	Shannon Hankinson (EE)	Ryan Andersen (EE)

Figure 1: Hardware Block Diagram for 2HAD-RAD

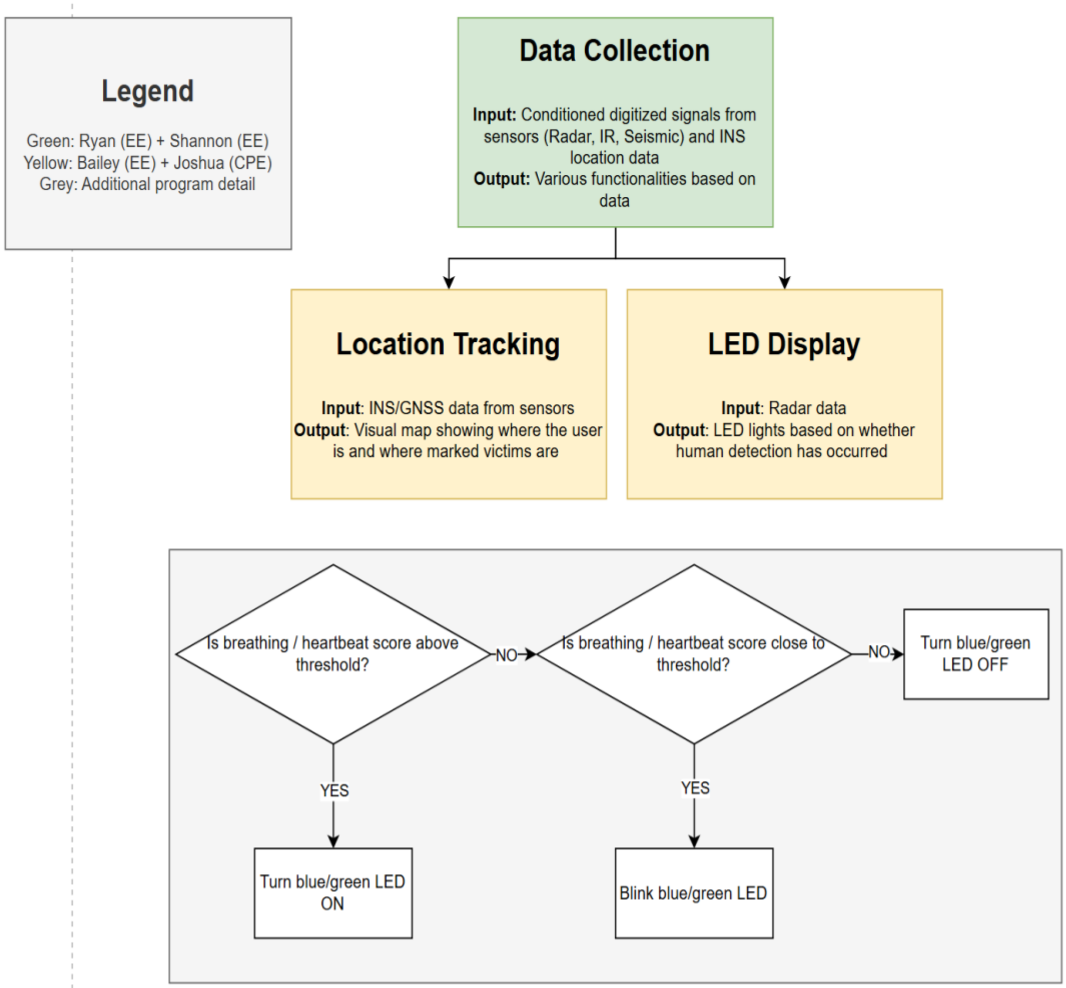

Figure 2: Software Block Diagram for 2HAD-RAD

2.7. HOUSE OF QUALITY

Symbol	Description
↑	Positive Correlation
⇈	Strong Positive Correlation
↓	Negative Correlation
⇊	Strong Negative Correlation
+	Positive Desirability
-	Negative Desirability

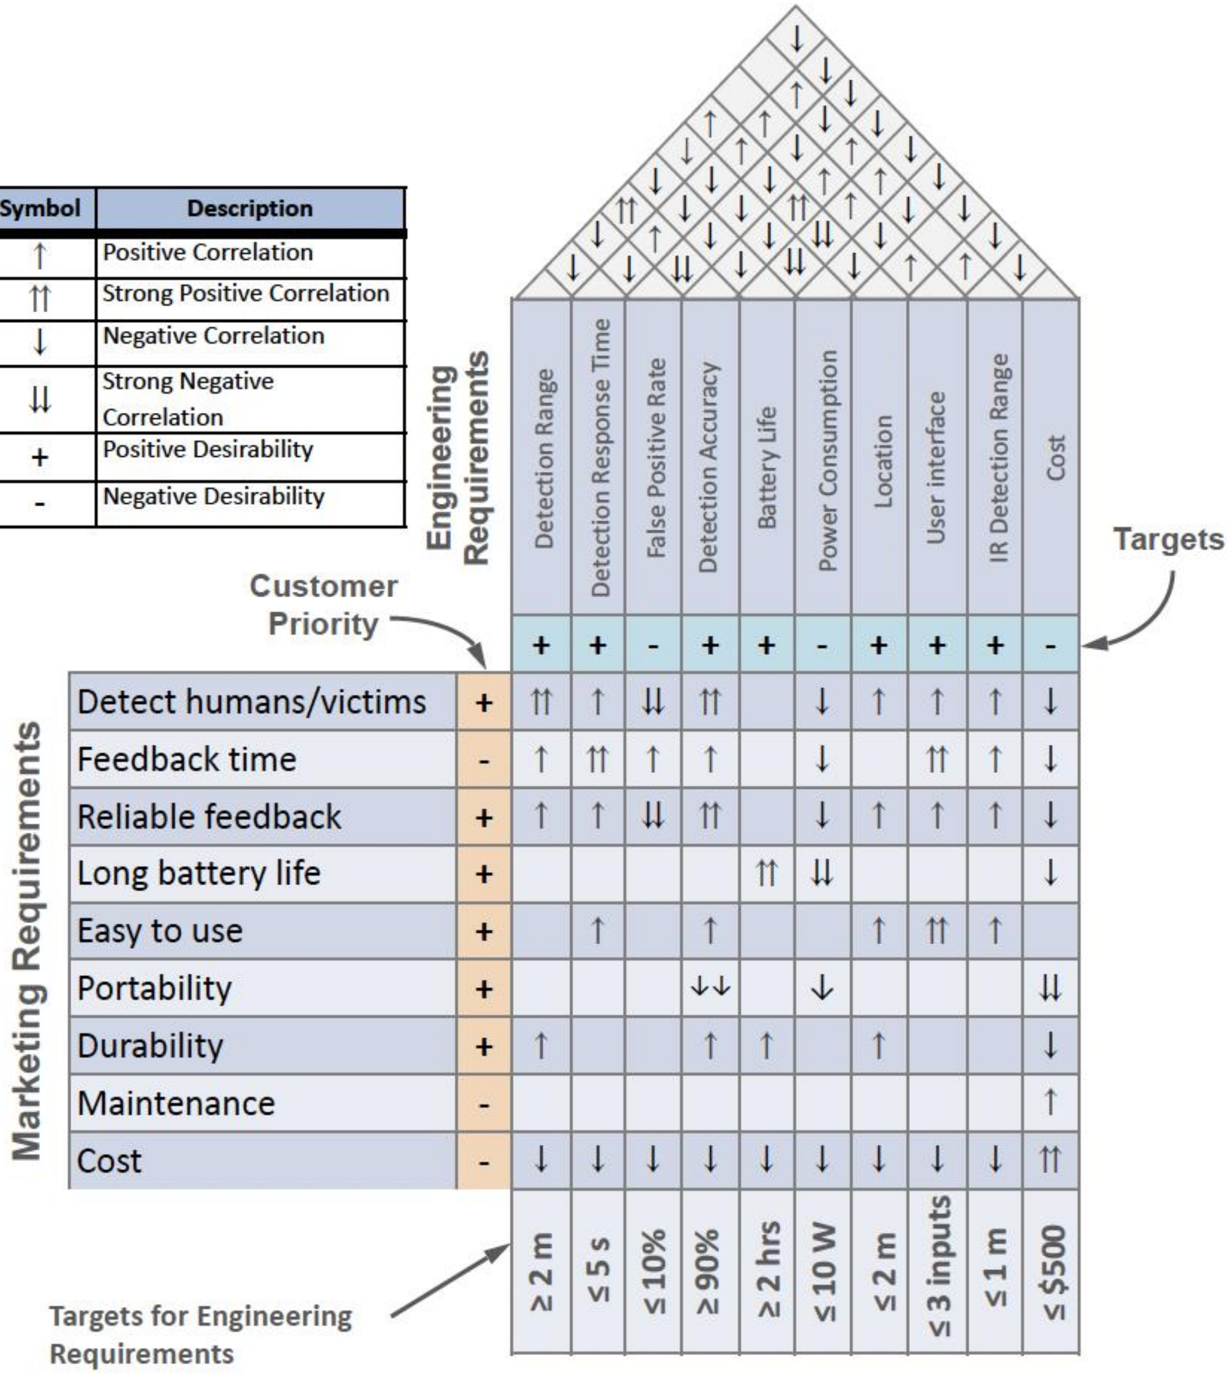

Figure 3: House of Quality

2.8. PROJECT BUDGET AND FINANCING

Our team will self-fund this project, so it is in our best interest to keep the total budget as conservative as possible. Here is the estimated total bill of materials (BOM) as can be foreseen currently:

Item	Quantity	Unit Cost (USD)	Total Cost (USD)
Radar Front-End Module	1	$200.00	$200.00
Custom PCB Components	1	$75.00	$75.00
Microcontroller / Processing Unit	1	$75.00	$75.00
Battery Pack	1	$40.00	$40.00
Battery Charger & Regulator Circuit	1	$20.00	$20.00
User Interface	1	$30.00	$30.00
Enclosure	1	$40.00	$40.00
Miscellaneous	1	$20.00	$20.00
Total Estimated Cost			$500.00

Table 3: Bill of Materials

2.9. INITIAL PROJECT MILESTONES

Our team formed in the Spring semester of 2025 when we began brainstorming. However, our serious efforts did not begin until the start of the Fall semester of 2025. Therefore, the timeline begins at the start of the Fall 2025 semester.

Item	Start Date	Anticipated End Date	Duration
Project Brainstorming	8/19/2025	8/26/2025	0 Weeks
Project Scope Finalized	8/26/2025	9/6/2025	1 Week
Individual Research	8/26/2025	9/6/2025	2 Weeks
Initial Project Proposal	8/26/2025	9/6/2025	1.5 Weeks
30-Page Milestone	9/6/2025	9/27/2025	3 Weeks
60-Page Milestone	9/27/2025	10/17/2025	3 Weeks
90-Page Milestone	10/17/2025	11/7/2025	3 Weeks
120-Page Milestone	11/7/2025	11/25/2025	3 Weeks
Final Draft	11/25/2025	12/3/2025	1 Week

Table 4: Senior Design I Project Documentation Timeline

Item	Start Date	Anticipated End Date	Duration
Individual System Design	9/5/2025	10/2/2025	4 Weeks
Individual System Testing	10/2/2025	10/30/2025	4 Weeks
Breadboard Prototyping	10/30/2025	11/21/2025	3 Weeks
PCB Design/Ordering	11/21/2025	12/12/2025	3 Weeks

Table 5: Senior Design I Project Design Timeline

Item	Start Date	Anticipated End Date	Duration
PCB Testing	1/9/2026	1/23/2026	2 Weeks
System Integration/Testing	1/23/2026	2/16/2026	4 Weeks
Practice Project Demo	2/16/2026	2/27/2026	2 Weeks
Finalize Documentation	2/27/2026	3/13/2026	2 Weeks
Practice Final Presentation	3/13/2026	3/20/2026	1 Week
Final Presentation	TBD	TBD	TBD

Table 6: Senior Design II Project Design Timeline

2.10. PROTOTYPE

To envision what our final project will look like, we rendered a blueprint of the physical model and a prototype for the user interface. To ensure portability and ease of use in the field, the physical model was designed to resemble a tablet, the size and shape of a kindle. The device features a power/reset button, status LED indicators and a USB-C power port. This design allows for reliable use by first responders in low visibility environments.

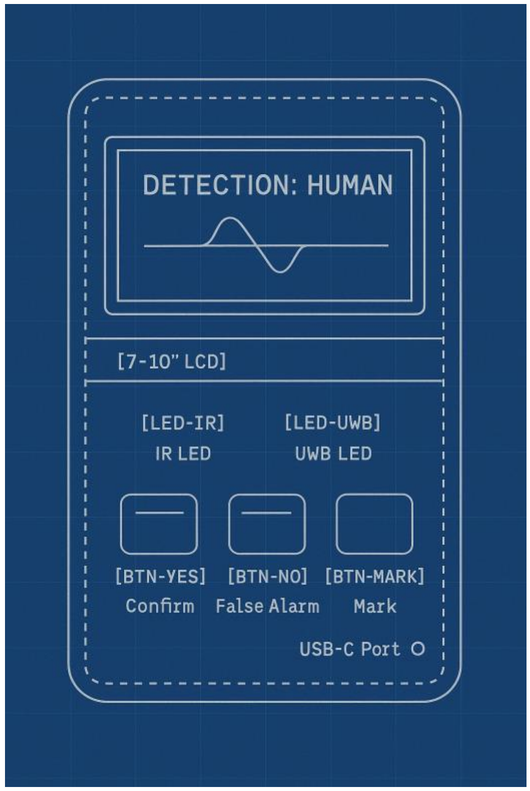

Figure 4: Portable Life Detector Prototype Blueprint

Furthermore, the prototype illustration for the user interface can be seen in the figure below. The interface highlights real time alters, thermal detection overlays and mapping functions to provide rescuers with immediate and actionable information.

Figure 5: Real Time User Interface Illustration

CHAPTER 3

1.1. Location/Tracking

Several sensing technologies can contribute to pinpointing locations, each with unique capabilities and limitations. In this project, we researched various methods of locating victims without the use of wireless communications. In this project, we evaluate Inertial Measurement Units (IMUs), Inertial Navigation Systems (INS), Ultra-wideband (UWB) radio, and Infrared (IR) and Radar modules as potential solutions.

3.1.1. Inertial Measurement Unit

An IMU has two separate configurations: a six-axis and a nine-axis. The six-axis IMU combines a three-axis accelerometer and a three-axis gyroscope. The nine-axis uses the same two components but adds a three-axis magnetometer. The additional component provides a reference to magnetic north, providing more accurate heading and attitude data to help correct sensor drift over time.

The six-axis configuration measures linear acceleration and gravity using the accelerometer while the gyroscope measures angular rate. Together these components are ideal for measuring tilt, detecting motion, and orientation changes. In this project, we want to implement a way to mark the location of a victim. The six-axis will provide a general location of the victim due to the orientation drifting over time, the yaw will not be precise, and the stored heading may be relative. However, it still delivers strong performance for many basic functionalities needed in search and rescue situations. It is sufficient for tasks like marking a location, logging path segments and relative orientation awareness. The six-axis IMU draws less power, offers a simpler calibration, and is immune to magnetic field corruption.

The nine-axis configuration adds the measurement of Earth's magnetic field vector. This allows the magnetometer to correct the drift periodically the gyroscope introduces. An academic study "Statistical Sensor Fusion of a 9-DoF MEMS IMU for Indoor Navigation" shows that using accelerometer, gyroscope, and magnetometer together in a robust unaltered calibration framework reduces orientation error significantly in environments with magnetic distortion, improving orientation accuracy by 80-90%, compared to gyroscope and accelerometer alone. However, magnetometers require careful calibration and are vulnerable to some local disturbances such as steel framing, machinery, or power lines.

3.1.2. Inertial Navigation System

An Inertial Navigation System (INS) combines motion sensing with computational processing to provide continuous estimates of position, velocity, and orientation. It consists of an Inertial Measurement Unit (IMU) typically a six-axis or nine-axis sensor which is integrated with a processor that executes the strapdown mechanization equations. The IMU's three-axis accelerometer measures linear acceleration, allowing the system to detect motion, tilt, and displacement, while the three-axis gyroscope measures angular rate to track rotations and directional changes. A three-axis magnetometer, when included, references the Earth's magnetic field to correct yaw drift, thereby ensuring consistent heading even during prolonged operation. Together, these sensors form a self-contained, nine-axis configuration capable of maintaining relative motion tracking without any external signals.

To enhance long term accuracy, many INS modules also include a Global Navigation Satellite System (GNSS) receiver such as GPS, Galileo, or BeiDou. The GNSS provides absolute position fixes, while the INS interpolates user movement between these updates using its high-frequency inertial data. When the Global Navigation Satellite System (GNSS) signal becomes unavailable, such as inside buildings, tunnels, or debris, the inertial navigation system (INS) continues to estimate motion autonomously. It relies on internal sensors and algorithms, such as the Kalman filter, to model and correct sensor drift, bias, and scaling errors.

In the context of this project, the INS plays a critical role in providing relative location tracking for search and rescue (SAR) operations. As rescuers move through a disaster zone, the INS continuously logs their motion path and orientation. When a victim is found, the rescuer can precisely mark their location relative to the rescuer, even if satellite coverage is weak or completely lost, by marking it in local coordinates. Once the rescuer returns to an area with GNSS visibility, these relative positions can be correlated with absolute coordinates to generate a map of victim locations. This hybrid approach ensures that every detected victim can be precisely revisited or relayed to other team members, even in GNSS denied environments.

By integrating the high update rate and drift-corrected attitude of the INS with periodic GNSS fixes, the system offers a dependable dead-reckoning navigation solution that ensures accuracy even during transitions between outdoor and indoor environments. In short, the INS ensures that responders always know where they have been, where they found victims, and how to guide others back, without relying solely on external infrastructure.

3.1.3. Ultra-Wideband (UWB)

Ultra-Wideband (UWB) is a short-range radio frequency (RF) technology that transmits extremely brief pulses over a wide frequency spectrum, typically between 3 and 10 GHz. UWB's wide bandwidth allows for precise time-of-flight (ToF) distance measurements, often within a range of 10–30 cm. This makes UWB highly suitable for determining relative distances and movements between devices. Unlike Wi-Fi or Bluetooth, UWB's broad spectral spread allows it to resist multipath interference from walls, debris, and reflective materials, ensuring stable signal integrity in cluttered environments.

While conventional UWB positioning systems rely on fixed anchors to calculate position, such infrastructure is often unavailable in disaster zones. Instead, this project considers UWB as a complementary sensing layer to the INS, where it could enhance relative localization among rescuers or between multiple handheld devices. For example, when a victim is located, the location is marked using the INS, UWB could help refine that relative position by providing ranging data between rescuers' devices in a local network. This cooperative positioning approach could improve mapping consistency when multiple responders are searching in different areas, especially when GNSS signals are unavailable.

Though UWB alone cannot serve as a global navigation method without fixed anchors, integrating its short-range precision with the INS's continuous trajectory tracking offers a promising pathway for improved relative mapping, spatial awareness, and team coordination in the field.

3.1.4. IR and Radar Module

IR systems detect infrared radiation emitted by warm bodies against cooler backgrounds. There are two types, passive and active IR. Passive IR or thermal cameras detect temperature differences and generate images or heat maps. Active IR or structured IR beams project IR light and measure reflections to detect presence or movement. If the IR is calibrated with distance, it can provide both direction and approximate distance. However, this is limited to line of sight only and heat sources can create false detections.

Radar transmits radio waves, measures reflections from objects, and detects motion from doppler shifts or locates targets through ToF and range resolution. Specialized bio-radar systems can detect chest motion through obstacles. Radar modules can estimate the range and angle of a detected person, which translates into a location estimate relative to the radar. Radar may be able to pinpoint a location of a victim through light obstacles but not through heavy or metallic barriers. Although both modules may be useful to aid an IMU or INS, it is not ideal as is.

3.1.5. Barometric Pressure Sensor

Barometric pressure sensors measure atmospheric pressure and convert it into height above sea level using the barometric formula. This conversion enables a system to estimate changes in elevation. In search and rescue operations, this capability is crucial for providing vertical context. For instance, it helps distinguish between floors in a collapsed building or identify when a rescuer ascends or descends stairwells. When combined with the INS and GNSS data, the barometer completes a three-dimensional tracking framework, enabling rescuers not only to log horizontal movement but also to determine the altitude of each detected victim location.

In the proposed system, when a rescuer marks the location of a victim using a handheld device, the barometer measures the ambient pressure at that precise point. This reading, combined with INS and GNSS data, generates a 3D coordinate. This coordinate provides both the relative (INS based) and absolute (GNSS based) positions, along with the vertical dimension obtained from the barometer. Such integration ensures that responders can later interpret whether a victim is located on a specific floor or depth level, crucial for prioritizing rescue routes and coordinating multi-level operations.

For this project, the barometric sensor must balance accuracy, stability, and low power consumption to suit portable, battery powered hardware. High-resolution devices like the Bosch BMP390 and Infineon DPS310 can achieve sub-meter altitude precision while drawing only a few microamps, making them ideal for continuous, low-drift monitoring. Their compact size and I2C/SPI interfaces simplify direct integration with the INS subsystems. Conversely, lower cost options like the Bosch BMP280 still provide meter level accuracy, which may be adequate for rough altitude tracking but could limit effectiveness in multi-story or subterranean search scenarios where precision vertical mapping is required. In essence, the barometric pressure sensor transforms the INS from a planar tracker into a comprehensive 3D localization system. This transformation provides rescuers with not only a map of the victims' locations but also their relative height or depth in relation to the rescuer's position.

3.1.6. Magnetometers

Magnetometers add an essential function to location tracking systems by measuring the Earth's magnetic field and providing absolute heading information. While accelerometers and gyroscopes can describe motion and orientation, they suffer from drift over time, especially in yaw, which directly affects navigation accuracy. A three-axis magnetometer corrects this issue by acting as a compass reference, which is particularly important indoors or in urban areas where GNSS reception is unreliable. For a search and rescue tool, this means a rescuer can maintain correct orientation while mapping a path or marking the location of victims. The choice of magnetometer must consider sensitivity, noise, and how well the device tolerates interference from nearby metal structures or electronics. Compact and efficient devices such as the Bosch BMM150 are widely used for basic applications, while higher-end parts like the MEMSIC MMC5983MA include features such as set/reset coils that help cancel residual magnetic offsets. In the context of this project, selecting a magnetometer with strong accuracy and resistance to interference, such as the LIS3MDL or MMC5983MA, would provide the most reliable heading data to support both victim location and team navigation in challenging environments.

3.2. Location Part Selection

For this project, we are selecting a 6 or 9-axis INS module, for which we plan to integrate a barometric pressure sensor. The INS provides continuous motion tracking through its internal IMU, while the barometer adds precise altitude and floor-level detection. The integrated GNSS enables absolute positioning outdoors, complementing the INS and preventing long term drift. This combination allows the handheld device to operate reliably in both indoor and outdoor environments, providing rescuers with accurate location information even when a victim cannot be immediately rescued.

3.2.1. Inertial Navigation System

3.2.1.1. M5Stack M135 Module

The M5Stack M135 is a ready to use positioning board that integrates a u-blox NEO-M9N-00B GNSS module, a BMI270 six-axis IMU, a BMM150 three-axis magnetometer, and a BMP280 barometric pressure sensor. This INS module is all in one and supports multi-constellation GNSS: GPS, GLONASS, Galileo, BeiDou, QZSS, and can update navigation up to 25 Hz. Its GNSS accuracy is about 1.5 meters under good conditions. The IMU provides selectable accelerometer ranges up to +/- 156.96 m/s^2 and gyroscope range up to +/- 2000º/s. The magnetometer has a field resolution of approximately 0.2μT. The barometer covers pressure ranging from 300 to 1100 hPa, usable for detecting altitude or floor changes.

The module also includes a backup battery to retain data during power loss. The M135 comes with a small dual in-line package (DIP) switch block on its PCB. These switches provide direct hardware level control without needing to reflash firmware. The pins that the GNSS and sensors connect to may be reassigned if needed. The GNSS module defaults to a 9600 baud rate, but this may be increased using the DIP for higher throughput GNSS data logging. The IMU, magnetometer, and barometer sit on the same I2C bus. This module is specified to operate at altitudes up to 80,000 m, sustain accelerations of about 39.2 m/s^2, and function at velocities as high as 500 m/s. The module also requires a fixed supply voltage of 3.3 V and has physical dimensions of 54 x 54 x 13.1 mm, with a total weight of 14.4 g. However, this module does not include any algorithms onboard, but we can apply them ourselves on an external MCU.

3.2.1.2. STMicroelectronics Teseo-VIC3DA Module

The ST Teseo-VIC3DA is a compact dead-reckoning GNSS module that integrates an AM330LHH six-axis IMU with ST’s Teseo-DRAW firmware for on module sensor fusion. It supports multi-constellation GNSS tracking including GPS, GLONASS, Galileo, BeiDou, and QZSS. Under open sky conditions, its GNSS accuracy is approximately 1.5 meters circular error probable. This module can output dead-reckoning fixes up to 30 Hz, while the raw GNSS update rate is 1 Hz. The integrated IMU offers selectable accelerometer ranges up to +/- 156.96 m/s^2 and the gyroscope ranges up to +/- 2000º/s. the VIC3DA does not include a magnetometer or barometric sensor on-chip, so external sensors would be required to achieve a nine-axis solution or altitude aiding capabilities.

The device also incorporates features such as map-matching feedback and an odometer input for wheel ticks, which can be used to enhance dead-reckoning accuracy when the GNSS signals are obstructed. Firmware updates are supported, and a temperature compensated crystal oscillator maintains frequency stability over -40º C to +85º C. This module requires a supply voltage of 3 to 3.6V and has a very low standby current consumption. The physical dimensions are 16 x 12.2 x 2.42 mm in a 24-pin LCC package. While it delivers integrated sensor fusion through ST’s algorithms, users who need a magnetometer or barometer must supplement the module with external devices to achieve a complete INS configuration.

3.2.1.3. U-blox NEO-M9V

The u-blox NEO-M9V is a dead-reckoning GNSS module designed for high-efficiency positioning in both automotive and embedded applications. It integrates a six-axis IMU directly into the module and employs an onboard Kalman filter algorithm to blend IMU and GNSS data for smoother navigation in weak-signal conditions. It supports multi-constellation GNSS reception (GPS, Galileo, GLONASS, BeiDou) and provides navigation solutions at update rates up to 25 Hz. Under nominal conditions, horizontal positioning accuracy is typically within 1.5 meters.

The onboard IMU supports accelerations up to +/- 156.96 m/s² and angular rates up to +/-2000°/s, making it suitable for tracking dynamic motion. Unlike the M135, the NEO-M9V does not integrate a magnetometer or barometer, though the module is designed to accept external I2C connected sensors for these functions. Operating limits allow for high dynamic use cases, including maximum altitudes of 50,000 meters and velocities up to 500 m/s, consistent with GNSS industry constraints.

The module requires a supply voltage of 2.7 to 3.6 V and has a very low power consumption of approximately 60 to 90 mW in continuous tracking mode, making it ideal for battery-powered field devices. Its compact footprint measures 12.2 x 16.0 x 2.4 mm in a standard u-blox NEO package. While the NEO-M9V delivers robust dead-reckoning and multi-band GNSS capabilities at a low power draw, the lack of onboard magnetometer and barometric pressure sensing means external components are required for a complete nine-axis INS solution. Although, the GNSS can determine altitude in outdoor areas.

3.2.1.4. U-blox ZED-F9R-03B

The u-blox ZED-F9R-03B is a high-performance, dual-band GNSS dead-reckoning module that combines a six-axis IMU, a multi-constellation dual-band GNSS receiver, and an onboard Kalman filter to deliver continuous navigation solutions. It supports GPS L1/L2, Galileo E1/E5b, BeiDou B1I/B2I, and GLONASS L1 signals, enabling robust positioning even in dense urban or obstructed environments. The module achieves a navigation output rate of up to 30 Hz, with real-time kinematic support for centimeter-level accuracy when correction data are provided.

The integrated IMU supports accelerometer ranges up to +/-156.96 m/s² and gyroscope ranges up to +/-2000°/s, allowing the device to maintain high-precision dead-reckoning during GNSS outages. Like the NEO-M9V, the ZED-F9R does not include a magnetometer or barometer but supports the addition of external I2C sensors for nine-axis capability and altitude aiding. Operational limits are consistent with other GNSS modules, supporting altitudes up to 50,000 m, accelerations up to 39.2 m/s^2, and velocities up to 500 m/s. However, this module is only capable of determining the altitude in outdoor conditions due to no barometer including in the package.

The ZED-F9R-03B requires a 2.7 to 3.6 V supply and has a typical power consumption of about 85 mA at 3.0 V at about during continuous tracking. Its footprint is 17 × 22 × 2.4 mm in the ZED LGA form factor, offering more PCB real estate than the NEO series but with higher performance. Although it has a higher cost relative to the VIC3DA and NEO-M9V, the ZED-F9R provides dual-band robustness, RTK capability, and superior accuracy, making it the strongest candidate for advanced field deployments where precision is critical.

Module	Sensors on board	GNSS	Power	Supply	Size(mm)	Cost
u-blox ZED-F9R-03B	3-D IMU (includes sensor fusion)	Dual band multi-constellation, RTK/PPP-RTK support, up to 30 Hz	About 0.255 W	2.7-3.6 V	17x22x2.4	$149-$164
u-blox NEO-M9V-20B	Six-axis IMU	L1 multi-constellation, Automotive/untethered dead-reckoning, up to 50 Hz (default 1Hz)	About 0.147 W	2.7-3.6V	16x12.2x2.4	$35
ST Teseo-VIC3DA	6-axis IMU	L1 multi-constellation with automotive dead-reckoning FW	0.161 W	3.3 V	16x12.2x2.42	$19-$21
M5Stack M135	Barometer, 6-axis IMU, magnetometer	L1 multi-constellation, up to 25 Hz	0.111 W	3 V	54x54x13.1	$45

Table 7: INS Comparisons

3.2.1.5. INS Selection Justification

The M5Stack M135 module was chosen for the inertial navigation subsystem of this project due to its all-in-one integration, compact form factor, and readiness for rapid prototyping. The M135 consolidates every core element of a full nine-axis INS into a single board, integrating a u-blox NEO-M9N GNSS receiver, BMI270 six-axis IMU (accelerometer + gyroscope), BMM150 magnetometer, and BMP280 barometric pressure sensor. This configuration provides immediate access to complete position, orientation, and altitude data without the need for multiple breakout boards or separate communication buses.

The M135 minimizes wiring complexity and eliminates the need for external data fusion hardware because the sensors are already internally connected through a shared I2C interface and GNSS UART. The compact DIP switch design of the module allows developers to directly remap sensor pins, enabling plug-and-play prototyping on breadboards or custom PCBs without the need for an additional breakout adapter. This makes the M135 exceptionally practical for the early stages of system integration, where flexibility and quick iteration are critical.

Functionally, the M135 enables victim tracking by integrating GNSS for global positioning, IMU and Magnetometer for continuous dead-reckoning and orientation, and Barometer for estimating vertical height. When a rescuer identifies a victim, the system can precisely record a three-dimensional coordinate, x–y, derived from the fusion of GNSS and INS, and z from barometric altitude. This forms an accurate local reference that can later be correlated with absolute coordinates once satellite coverage is restored.

In summary, the M5Stack M135 offers a self-contained, reliable, and power efficient navigation subsystem that drastically simplifies development and sensor fusion efforts. Its integrated design, prototyping convenience, and robust component suite make it the most cost effective and technically sound choice for implementing a portable 3D localization platform in this search and rescue system.

3.2.2. Barometer

3.2.2.1. Bosch BMP390

The Bosch BMP390 is one of the most precise barometric pressure sensors currently available for compact designs. It comes in a 2.0 x 2.0 x 0.8 mm package and offers a relative accuracy of about +/- 3 Pa, which corresponds to roughly +/- 0.25 m in altitude. Absolute accuracy is rated around +/- 0.5 hPa over a wide temperature range, while its temperature coefficient of offset (TCO) is only +/- 0.62 Pa/K, which helps maintain stable readings when the environment changes. The device operates over a pressure range of 300 to 1250 hPa and can sample at up to 200 Hz. At a 1 Hz update rate, it consumes only about 3.2 μA, with a sleep current near 1.4 μA, making it well suited for low power portable devices. While it is slightly more expensive than older parts, its noise performance and accuracy make it a strong candidate for applications where detecting floor changes or subtle altitude shifts is important.

3.2.2.2. Infineon DPS310

The Infineon DPS310 is a digital barometric pressure sensor designed for low power consumption and strong thermal stability. It uses a capacitive sensing element, which makes it less sensitive to sudden temperature changes, such as when a user moves between indoor and outdoor environments. The device is packaged at 2.0 x 2.5 x 1.0 mm and provides a relative accuracy of +/- 0.06 hPa (about +/- 0.5 m) and an absolute accuracy of +/- 1 hPa. It supports both I2C and SPI interfaces and can achieve measurement rates up to around 157 Hz, depending on the oversampling settings. Typical current is only a few microamps in low frequency operation, making it an attractive choice for battery powered applications. One advantage of the DPS310 is its internal 32 sample FIFO buffer, which reduces the load on the microcontroller. With its combination of precision, stability, and low cost, the DPS310 represents a practical option for adding altitude awareness to an INS system.

3.2.2.3. Bosch BMP280

The Bosch BMP280 is packaged in a 2.0 x 2.5 x 0.95 mm LGA, it provides a relative accuracy of about +/- 0.12 hPa (+/- 1 m altitude) and a typical absolute accuracy of +/- 1 hPa. Its TCO is higher than that of the BMP390, around 1.5 Pa/K, which means its readings can drift more in changing temperatures. Power requirements are very low, drawing about 2.7 μA at a 1 Hz update rate, with standby currents below 0.3 μA. Interfaces include I2C and SPI, both operating at relatively high speeds. Although its accuracy does not match newer sensors, it is inexpensive, widely available, and sufficient for many applications where floor level precision is not required.

3.2.2.4. ST LPS22HB

The ST LPS22HB is a compact barometer, measuring at 2.0 x 2.0 x 0.76 mm, that emphasizes robustness and flexibility. It covers a pressure range of 260 to 1260 hPa and achieves a typical relative accuracy of +/- 0.1 hPa. Power consumption depends on the mode, ranging from about 3 to 12 μA at 1 Hz. In addition to its low power features, the LPS22HB includes built-in FIFO storage and programmable interrupts, allowing it to signal the microcontroller when thresholds are crossed, reducing the need for constant polling. It is also designed to withstand high over pressure and extreme shock levels up to 22,000 g, making it suitable for rugged environments.

Sensor	Relative Accuracy	Absolute Accuracy	Pressure Range	Power	Output Data Rate	Supply Voltage	Size (mm)
Bosch BMP390	+/-0.25 m	+/-4.2 m	300-1250 hPa	0.0106 mW	200 Hz	1.65–3.6 V	2.0 x 2.0 x 0.8
Infineon DPS310	+/-0.5 m	+/-8 m	300-1200 hPa	0.0165 mW	157 Hz	1.7-3.6 V	2.0 x 2.5 x 1.0
Bosch BMP280	+/-1 m	+/-8.3 m	300-1100 hPa	0.0089 mW	157 Hz	1.71-3.6 V	2.0 x 2.5 x 0.95
ST LPS22HB	+/-0.8 m	+/-0.8-8.3 m	260-1260 hPa	0.0099-0.0396 mW	75 Hz	1.7-3.6 V	2.0 x 2.0 x 0.76

Table 8: Barometric sensor comparisons.

3.2.2.5. Barometer Selection Justification

While the M5Stack M135 already integrates a Bosch BMP280 barometric pressure sensor, additional research was conducted to evaluate higher precision and lower power alternatives, such as the Bosch BMP390, Infineon DPS310, and ST LPS22HB. These standalone sensors were reviewed to evaluate all our options and identify potential improvements in altitude accuracy, thermal stability, and integration complexity. The newer BMP390 and DPS310, for instance, offer sub-meter precision and advanced temperature compensation, making them attractive options for fine grained vertical mapping in multiple story search scenarios. Since the integrated BMP280 in the M135 offers sufficient performance to meet the project’s resolution requirements, effectively distinguishing floor-level changes and elevation shifts, it was decided to retain the onboard component. This choice minimizes the need for additional wiring, conserves power and simplifies firmware development. It still enables reliable three-dimensional localization when combined with INS and GNSS data. The evaluation revealed that while higher-end sensors could slightly enhance vertical accuracy, the integrated barometer provides the most practical balance of performance, efficiency, and integration simplicity for this prototype stage.

3.2.3. Magnetometer

3.2.3.1. MEMSIC MMC5983MA

The MMC5983MA is a high performance three-axis magnetometer that uses anisotropic magneto resistive (AMR) technology. It offers an 18-bit resolution and extremely low RMS noise of about 0.4 mG, making it one of the most precise magnetometers available in this class. The measurement range is +/- 8 Gauss, and the device supports output data rates up to 1000 Hz. A unique feature of the MMC5983MA is its built-in set/reset function, which helps cancel residual magnetic offsets that can build up in metal-rich environments. It operates from 1.8 to 3.6 V and supports both I2C and SPI interfaces. Its small footprint (3 × 3 mm) and low current draw make it feasible for embedded systems, though its cost is somewhat higher than simpler devices.

3.2.3.2. ST LIS3MD

The LIS3MDL from STMicroelectronics is a widely adopted three-axis magnetometer with selectable full-scale ranges of +/- 4, +/- 8, +/- 12, or +/- 16 Gauss. It provides 16-bit output data and supports both I2C and SPI interfaces, with output data rates up to 1000 Hz. Power consumption depends on the mode, ranging from around 40 µA in low power configurations to 270 µA in ultra-high-resolution mode at 20 Hz. The sensor is rated for operation between −40 ºC and +85 °C and comes in a compact 3 x 3 mm package. It also includes built-in self-test and interrupt features, making it easier to integrate into systems that require reliable event driven heading measurements. The LIS3MDL represents a balanced option in terms of accuracy, flexibility, and cost, and it is well supported in software libraries and hardware platforms.

3.2.3.3. QST QMC5883L

The QMC5883L is a low-cost magnetometer that uses AMR technology and provides a three-axis magnetic field measurement with a typical heading accuracy of 1 to 2° after calibration. Its operating voltage range is 2.16 to 3.6 V, and it supports I2C communication. The maximum output data rate is about 200 Hz, which is sufficient for many basic navigation tasks but lower than higher-end devices. The small size and inexpensive price make it attractive for student projects and prototypes, but performance can vary depending on the manufacturer, so calibration is essential for reliable results.

3.2.3.4. Bosch BMM150

The BMM150 is a compact three-axis magnetometer that is often paired with Bosch IMUs in sensor fusion applications. It measures magnetic fields up to +/- 1300 µT in the X and Y axes and up to +/- 2500 µT in the Z axis. Typical heading accuracy is about +/- 2.5° under good calibration conditions. The sensor supports I2C and SPI interfaces and offers both forced and normal operating modes, with a maximum output data rate of 30 Hz. Current consumption is around 170 µA at 10 Hz in low-power mode, and the package measures only 2 x 2 mm. Although its maximum data rate is lower than other options, its integration within the Bosch sensor ecosystem makes it easy to pair with IMUs such as the BMI270. This makes it a practical choice for projects where power efficiency and compact integration are more important than high-speed sampling.

Sensor	**Range**	**Noise/Res.**	**ODR**	**Power**	**Supply Voltage**	**Heading Accuracy**
MEMSIC MMC5983MA	+/- 8 G	0.4 mG RMS	1000 Hz	2.2 mW	1.8-3.6 V	+/- 0.5º (well calibrated)
ST LIS3MDL	+/- 4/8/12/16 G	3.2 mG typical	1000 Hz	0.12-0.89 mW	1.9 – 3.6 V	+/- 1-2º
QST QMC5883L	+/- 8-20 G	0.5-2 mG	200 Hz	0.33 mW	2.16-3.6 V	+/- 1-2º (needs calibrated)
Bosch BMM150	+/- 1300 µT (XY), +/- 2500 µT (Z)	0.3 µT res.	30 Hz	0.56 mW	1.62-3.6V	+/- 2.5º typical

Table 9: Magnetometer sensor comparisons

3.2.3.5. Magnetometer Selection Justification

Similarly, while the M5Stack M135 features a Bosch BMM150 three-axis magnetometer, a comparative review was conducted to explore alternative magnetometers that could improve heading accuracy and interference tolerance. This review also provided insights into our available options. Alternatives such as the ST LIS3MDL, MEMSIC MMC5983MA, and QMC5883L were analyzed for factors including magnetic sensitivity, noise, and offset compensation. High performance devices like the MMC5983MA demonstrated superior resolution and active offset cancellation features, while the LIS3MDL provided broader range settings and strong library support for embedded platforms. However, these benefits come at the cost of additional wiring, calibration steps, and power consumption. The BMM150 is already internally aligned with the M135's IMU on the same bus, ensuring consistent axis orientation, compact integration, and minimal setup complexity. The onboard magnetometer provides sufficient accuracy for heading correction and drift mitigation in fusion with the accelerometer and gyroscope. Therefore, the team concluded that retaining the integrated BMM150 provided the most efficient and practical solution, fulfilling the system's requirements for orientation tracking, drift correction, and absolute heading without introducing additional hardware dependencies.

3.3. Life Detection Modules

In search-and-rescue applications, the ability to detect human presence through common building materials such as drywall, wood, and concrete is critical for locating victims who may not be visible to thermal or optical sensors. Traditional sensing solutions, such as radar modules and infrared cameras, provide valuable information but can be limited in their penetration depth or sensitivity to stationary individuals. To enhance the overall system, this research investigates alternative modules and sensors, particularly those that use ultra-wideband (UWB) or related technologies, that are capable of transmitting signals through obstacles and detecting subtle human motion, such as respiration. The primary objective of this section is to evaluate candidate components that complement existing radar and IR systems, providing redundancy and increased reliability when searching for victims in complex structural environments.

3.3.1. Ultra-Wideband (UWB) Modules

Beyond localization, UWB also offers valuable capabilities for life detection, particularly when fused with millimeter-wave radar. UWB modules transmit pulses across a wide frequency range, usually between 3.1 and 10.6 GHz. These pulses can penetrate non-metallic materials like drywall, wood, and light concrete. Lower frequency components within this band achieve greater penetration depth, while higher frequencies offer finer spatial resolution. This dual behavior allows UWB to estimate the distance and presence of a target behind obstacles, even when direct line of sight is blocked.

When combined with radar in a sensor fusion framework, each technology compensates for the other's limitations. Radar offers high-resolution Doppler and micro-motion signatures, including the subtle chest motion caused by breathing. On the other hand, UWB provides depth and material penetration data. By merging both outputs through algorithms such as an Extended Kalman Filter or weighted fusion model, the system can both detect life signs and estimate victim location through barriers.

This hybrid UWB–radar subsystem is designed to feed its data directly into the INS-based mapping system, where the detected signal's relative position and range can be logged as a 3D pin. The result is a multi-modal detection and localization framework capable of confirming human presence and recording precise spatial context, even in GNSS-denied or visually obstructed environments.

UWB modules are also well-suited for field deployment due to their compact size, energy efficiency (operating at 100–150 mW active), and already supported robust communication protocols (IEEE 802.15.4z). With proper signal processing techniques like clutter suppression and motion extraction, UWB significantly enhances the radar's ability to differentiate between living targets and noise. This improvement leads to enhanced sensitivity and reliability in the handheld device's life detection capabilities.

3.3.2. Low Frequency RF Modules

Unlike higher-frequency systems (24 GHz or 60 GHz radar), low-frequency RF modules take advantage of improved penetration properties in dense materials. At these lower frequencies, attenuation in concrete and wood is significantly reduced, enabling detection through walls greater than 30 cm thick in some cases. By operating in sub-GHz bands, the tradeoff is spatial resolution. As bandwidths are typically narrow, limiting localization precision to meters rather than centimeters. They provide an alternative when UWB signal strength is insufficient to penetrate through thicker barriers. The modules are often larger due to the size of the antenna. The power usage is usually between 50 to 100 mW. These systems may not retrieve small motions like respiration as effectively as UWB, but their penetration depth is a significant advantage.

3.3.3. Seismic Sensors

While RF-based systems such as radar and UWB provide strong potential for through-wall detection, they can be limited when faced with dense, rebar-reinforced concrete or when operating in environments with high electromagnetic noise. In such scenarios, a seismic sensing modality serves as a highly effective complement by directly coupling to structural surfaces and capturing mechanical vibrations that propagate through walls, debris, or flooring. These vibrations, generated by human activity such as knocking, shifting limbs, or even subtle body motion, travel through materials that often block radio frequency signals.

A promising approach for integrating seismic sensing into the handheld device is to embed a MEMS accelerometer within a spring-loaded contact pad on the housing. The spring maintains continuous pressure against the surface, ensuring strong mechanical coupling and consistent transmission of structural vibrations into the sensor. This design greatly improves the sensor's sensitivity compared to a PCB mounted accelerometer alone and reduces interference from operator induced noise. As vibrations travel through the structure, they are mechanically transmitted through the spring assembly to the MEMS sensor, which converts them into acceleration signals. These signals can then be digitally filtered to isolate the low frequency components associated with human movement. Advanced signal processing, such as band-pass filtering and fast Fourier transforms (FFT), can extract spectral signatures of footsteps, tapping, or shifting debris, which are then fused with radar and UWB data to improve detection confidence and reduce false positives.

While MEMS accelerometers generally offer lower sensitivity to ultra-low frequency vibrations compared to geophones, this spring-mounted approach significantly narrows that gap and maintains the advantages of compact size, low power consumption, and direct digital interfacing. It also supports rapid deployment in field conditions: responders can quickly press the device against a wall or surface and begin scanning for signs of life without needing to deploy separate hardware.

Alternatively, a geophone remains a strong option for scenarios requiring maximum sensitivity to very low frequency vibrations. In such a configuration, the geophone would be deployed as a separate, externally placed module, positioned in contact with the ground or structural surface and connected to the handheld system via a shielded cable. This stationary setup minimizes noise from operator movement and preserves the geophone's exceptional low frequency performance, though it sacrifices portability and integration simplicity.

Ultimately, the choice between these two seismic sensing strategies will depend on project priorities and testing outcomes. A spring mounted MEMS accelerometer offers a fully integrated, compact solution ideal for rapid scanning and earthquake response scenarios, while an external geophone module provides unmatched sensitivity for long duration or stationary monitoring. Both represent viable pathways to enhance the system's capability to detect human presence in RF challenging environments.

3.3.4. Radar Module

The radar module serves as the core sensing element of our system, enabling the detection of subtle physiological motions such as respiration and heartbeat potentially even when line-of-sight is obstructed by obstacles like drywall, debris, or smoke. Millimeter-wave (mmWave) radar operates in the 60–81 GHz range, using short wavelengths that provide high spatial resolution, strong material penetration, and robustness to environmental conditions such as dust, darkness, and temperature changes, characteristics that are essential in search and rescue operations where traditional sensing modalities like cameras or infrared fail. By transmitting a continuous frequency modulated signal and analyzing its reflections, the radar can determine both the range and velocity of objects and possibly detect micro motions on the order of millimeters caused by human breathing and heartbeat. These capabilities enable it to distinguish living individuals from static objects and provide continuous monitoring even in low visibility or cluttered environments. Because radar-based sensing is passive, noncontact, and unaffected by ambient light, it is particularly suited to detecting trapped or unconscious victims where other sensors might miss critical signs of life. For these reasons, the radar module is the cornerstone of our system architecture, and careful evaluation of available radar technologies was a critical step in our design process.

3.3.5. Infrared Thermal Sensor

An infrared (IR) thermal sensor detects heat as opposed to visible light, which allows it to sense the presence of warm objects such as people, animals, or machinery. It works by measuring the infrared radiation naturally emitted from surfaces and converting that energy into temperature readings. Each measurement represents the thermal intensity of a small region in the sensor's field of view, which can then be analyzed to determine relative heat patterns. In essence, an IR sensor functions as a thermometer array at a distance, eliminating the need for contact, allowing it to capture real time temperature data that provides insight into environmental and biological activity.

We selected an infrared thermal sensor because it enhances our system's ability to detect human presence based on body heat rather than motion alone. While radar provides distance and movement information, the IR module contributes temperature-based verification, improving the reliability of detection in scenarios where movement may be minimal. In this project, the IR sensor will be programmed to recognize temperature ranges corresponding to human body heat, typically between 30°C and 40°C. When readings fall within that range, the system will display a "Human Detected" message on the LCD screen. The stretch goal for this subsystem involves using the same thermal data to generate a color-based heat map visualization on the display. The sensor communicates with the ESP32WROOM microcontroller using the I2C protocol, ensuring simple wiring, low power operation, and straightforward data acquisition. By incorporating this module, the system gains an additional sensing layer that supports both accuracy and environmental adaptability in field applications.

3.3.6. CO-2 Sensor

A CO_2 sensor is particularly effective for life detection because it can identify human respiration by detecting the exhaled carbon dioxide in the air. Every breath a person takes releases carbon dioxide (CO_2) at concentrations significantly higher than the ambient level, resulting in a measurable increase in CO_2 levels when someone is present, even behind barriers like rubble or walls. In rescue or monitoring systems, non-dispersive infrared (NDIR) CO_2 sensors are commonly employed due to their accuracy and selectivity. These sensors detect the infrared absorption characteristics of CO_2 molecules. This makes them valuable for detecting trapped survivors in confined spaces because carbon dioxide tends to accumulate when ventilation is restricted. However, their detection range is often constrained by factors such as airflow, diffusion rate, and sensor sensitivity. Despite this, a CO_2 sensor offers direct evidence of metabolic activity, making it a reliable indicator of life in enclosed or low-visibility environments.

3.4. Life Detection Module Part Selection

The proposed life detection system prioritizes the radar module as the primary sensing technology because it has a proven ability to detect minute physiological motions, such as breathing or heartbeat, even through non-metallic materials like drywall or debris. Radar's non-contact nature and resilience to lighting and temperature variations make it a dependable core sensor for continuous monitoring. To enhance this, an infrared (IR) sensor is integrated to capture thermal signatures indicative of human presence. This sensor provides valuable surface level temperature data that complements radar readings, particularly when line-of-sight is available. The ultra-wideband (UWB) module is chosen for sensor fusion with the radar system, as research suggests that combining UWB's precise time-of-flight measurement with radar's Doppler sensitivity significantly improves detection accuracy and material penetration performance. UWB's low power and wide spectrum signals also help distinguish living subjects from static objects by providing more precise resolution of distance and motion. Lastly, a seismic sensor is set as a stretch goal, which could be incorporated later to detect low-frequency vibrations or knocks transmitted through solid surfaces. While not crucial for the initial prototype, incorporating this feature in future iterations could enhance detection capabilities to ground coupled signals. This improvement would lead to more accurate survivability assessments in complex rubble environments. This multi-sensor architecture strikes a harmonious balance between penetration depth, precision, and environmental adaptability, providing a solid foundation for the development of next generation life detection systems.

3.4.1. Ultra-Wideband Modules

3.4.1.1. DWM3000EVB

The DWM3000EVB is an Arduino shield form factor evaluation board built around the DWM3000 UWB module, which uses the DW3110 transceiver IC. Operating in the 6.5 to 8 GHz range, this module can achieve centimeter-level ranging accuracy while claiming to provide modest penetration through barriers such as drywall and wooden structures. With a power consumption of approximately 100 mW, it is suitable for battery powered handheld systems. However, it is primarily designed for communication and localization tasks; adapting it to detect human respiration or micro-motions would require additional signal processing algorithms.

3.4.1.2. ULM3-STM32 + DWM3000

The ULM3-STM32 + DWM3000 stack combines a STM32 microcontroller with the DWM3000 UWB transceiver in a single embedded platform. Supporting IEEE 802.15.4z protocols, the stack delivers about +/- 5 cm ranging accuracy in the 6.25 to 8.25 GHz band with an average active power draw of 120 mW. This integration reduces system complexity by eliminating the need for an external processing unit, making the design more compact and suitable for handheld applications. It also includes an onboard OLED display and offers a detection range up to 40 m in an open area when communicating at 6.8 Mbps.

3.4.1.3. DWM1000

The DWM1000 is a fully integrated ultra-wideband transceiver module developed by Qorvo. It is based on the DW1000 IC, which complies with the IEEE 802.15.4-2011 UWB standard (HRP PHY) and supports time-of-flight ranging and location services. The module integrates the RF front end, crystal oscillator, and antenna in a compact package measuring approximately 23 mm x 13 mm x 2.9 mm. According to the datasheet, the device supports data rates of 110 kbps, 850 kbps, and 6.8 Mbps, with line-of-sight (LOS) accuracy ranging from +/- 10 cm under favorable conditions. Its operating frequency spans from 3.5 to 6.5 GHz and it is powered from a 2.8 to 3.6 V supply. Active transmit current is listed as 46 mA at 3.3 V, while receive current is 50 mA, translating to approximately 150 to 165 mW of active power consumption.

3.4.1.4. DWM3000

The Qorvo DWM3000 is a surface mounted ultra-wideband transceiver module that integrates the DW3110 chipset and complies with the IEEE 802.15.4z standard. It is designed to support secure ranging and localization through time-of-flight measurements, with LOS accuracy on the order of +/- 10 cm as specified in the product documentation. The UWB module transmits extremely short pulses, which spread across a wide frequency band and allow the signals to penetrate materials like drywall with relatively low loss. The module operates within 6.5 to 8 GHz with supported data rates of 850 kbps and 6.8 Mbps, enabling both high-speed and low power communication modes. It is supplied from a 2.5 to 3.6 V source and achieves typical active currents in the tens of milliamperes range, corresponding to power consumption near 100 to 150 mW during transmission or reception. With a compact footprint of 13 x 13 x 2.9 mm, the DWM3000 provides an integration ready form factor that can be directly interfaced to an external microcontroller through standard SPI communication. This combination of size, accuracy, and compliance with modern UWB standards makes the module suitable for embedded applications requiring precise ranging or human detection in multipurpose sensing systems.

Module	**Frequency**	**Power Consumption**	**LOS**	**Data Rate**	**Supply Voltage**	**Cost**
DWM3000EVB	6.5 to 8 GHz	100mW	+/- 10 cm	850 kbps / 6.8 Mbps	2.5 to 3.6 V	$29.50
ULM3-STM32 + DWM3000	6 to 8.5 GHz	120 mW active	+/- 5 cm	6.8 Mbps	External 3.7 to 5 V	$88.06
DWM1000	3.5 to 6.5 GHz	150 to 165 mW	+/- 10 cm	110 kbps/850 kbps/6.8 Mbps	2.8 to 3.6 V	$25-$30
DWM3000	6.5 to 8 GHz	100 to 150 mW	+/- 10 cm	850 kbps / 6.8 Mbps	2.5 to 3.6 V	$25

Table 10: UWB product comparison

3.4.1.5. UWB Selection Justification

The DWM3000EVB was selected as the preferred UWB platform for this project due to its ease of prototyping and hardware compatibility. As an evaluation board with an Arduino shield form factor, it can be directly interfaced with common development platforms such as the ESP32-WROOM without requiring custom breakout adapters or fine-pitch soldering. This reduces hardware integration time and allows the team to focus on signal processing and algorithm development rather than PCB rework. Functionally, the board maintains the performance of the core DWM3000 module, providing centimeter-level ranging accuracy in the 6.5–8 GHz band while consuming power suitable for battery-operated systems. It’s on-board antenna, standard pin headers, and documented

SPI interface simplify firmware integration and make it possible to validate ranging and potential life detection applications quickly. Compared to standalone modules like the DWM3000 or DWM1000, the EVB provides a ready-to-use prototyping environment that facilitates iterative testing and future migration to custom PCB integration once the fusion framework with radar and INS sensors is validated.

3.4.2. Geophones

3.4.2.1. Geophone SM-24

The SM-24 is a vertical geophone element, widely used in seismic and structural monitoring. It is designed to detect vertical vibrations starting at about 10 Hz and can pick up motion up to a few hundred hertz. It is sensitive enough to detect small ground motions and is very consistent across units. It has been labeled as a “workhorse” sensor in both academic research and industry. For the application of this project, the SM-24 would excel at capturing knocks, footsteps, or vibrations traveling through floors or rubble as these often fall within the 10 to 200 Hz range. It does not require any external power source and is encased in aluminum which makes it ideal for portable integration. However, it is not ideal for extremely slow motions below 10 Hz, such as subtle body movements or shifts. It performs best when it is firmly mounted on a stable surface with tilt angles under 10 degrees to avoid damping errors.

3.4.2.2. SM4 Horizontal Geophone

The SM4 variant is oriented for horizontal motion sensing rather than vertical. It’s often used for shear wave or lateral vibration detection in structural and seismic systems. According to manufacturer catalogs from Sercel and similar suppliers, the SM-4 share core coil and damping parameters to SM-24, delivering comparable sensitivity around 28-30 V/m/s but tuned for horizontal axes. It begins responding at 10 Hz, maintaining strong performance up to 160 through 190 Hz before other unwanted responses start to show up. This configuration is valuable in environments with multi-directional disturbances such as building collapses or uneven terrain. Housed in a compact, weather resistant case, it requires secure orthogonal mounting to the vertical unit, and like its sibling, it falters below 10 Hz for very low velocity events.

3.4.2.3. EG-4.5-II Geophone

The EG-4.5-II is a low frequency vertical geophone engineered for enhanced sensitivity to subtle seismic events, featuring a natural frequency of 4.5 Hz that expands its usable bandwidth between 4.5 and 200 Hz. In comparison to the previous choices, this geophone offers a sensitivity of about 40 V/m/s and has a moving coil design that minimizes internal damping, low distortion, and high signal to noise ratios. This model could be ideal for detecting micro-motions from breathing, crawling, or distant footsteps. In our project, this component would provide a critical low pass extension to the SM-24’s capabilities. It remains compact and portable but is slightly larger than the SM series.

Specification	SM-24	SM-4	EG-4.5-II
Orientation	Vertical	Horizontal	Vertical and Horizontal available
Natural Frequency	10 Hz	10 Hz	4.5 Hz
Sensitivity	28.8 V/m/s	28.8 V/m/s	28.8 V/m/s
Damping (oc)	0.25	0.25	0.6
Max Tilt Angle	10º	25º	15º
Bandwidth	10-240 Hz	10-180 Hz	4.5-140 Hz

Table 11: Geophone comparison

3.4.2.4. Geophone Selection Justification

For seismic sensing, we opted for the SM-24 vertical geophone due to its ruggedness, and suitability for our primary structural vibration band. In a search and rescue scenario, the SM-24 is highly effective at capturing knocks, footsteps, shifts in debris, and other structural signals that travel through floors, walls, and rubble. This practical and low-power modality serves as a reliable complement to UWB/radar detections. While it is not optimized for sub-10 Hz micro-motions, its strengths align with our immediate objectives: detecting impulsive and low-audio events through contact coupling, integrating those detections into our INS pin-drop workflow, and guiding responders back to likely victim locations. Due to its mechanical integration requirements and secondary priority, this component is currently considered a stretch goal. It will be incorporated if time and resources permit during later development stages.

3.4.3. Radar Modules

3.4.3.1. Infineon BGT60TR13C

The Infineon BGT60TR13C is a compact 60 GHz FMCW radar sensor designed for operation in the unlicensed ISM band under FCC Part 15, offering an attractive balance of size, power, and baseline sensing capability for embedded applications. It features one transmit and three receive channels and includes an antenna-in-package (AiP), which reduces the need for external antenna design and tuning compared to bare die radar solutions. Its small 6.5 x 5 mm footprint makes it suited for highly space constrained systems, and its typical active power consumption of 0.5–0.6 W fits comfortably within our system’s power budget. The device supports the high sampling rates necessary to achieve sub 200 ms end-to-end latency and operates across a wide temperature range (−40 °C to +85 °C), exceeding our environmental requirements.

However, the BGT60TR13C outputs only raw intermediate frequency (IF) data over an SPI interface and does not include on-chip digital signal processing. As a result, all range fast Fourier transforms, doppler processing, and micro doppler extraction must be implemented on an external processor such as the ESP32. This shifts significant computational burden to the host MCU, potentially increasing firmware complexity and impacting real-time performance, especially when additional tasks such as wireless communication and user-interface updates are considered. Although the built-in AiP simplifies integration, antenna orientation, enclosure design, and PCB layout still strongly influence the resulting beam pattern and detection range. Achieving consistent vital sign detection through drywall at 2 m, one of our stretched design goals, would require careful antenna positioning, extensive signal processing optimization, and potentially external amplification or filtering stages.

From a development perspective, Infineon's software support is limited primarily to low level drivers and evaluation tools. Application-level algorithms for respiration and heartbeat detection would need to be built from the ground up, increasing development time and risk compared to more integrated solutions. While the BGT60TR13C is appealing for its small size, low power, and low unit cost, these benefits are offset by the additional design complexity it introduces. It is best suited for designs where custom signal processing pipelines and antenna systems are acceptable or where minimizing cost is prioritized over rapid prototyping and system level simplicity.

3.4.3.2. Infineon BGT60ATR24C

The Infineon BGT60ATR24C is a 60 GHz FMCW radar sensor designed for applications that require higher spatial resolution and multi target detection. Compared to the BGT60TR13C, it provides two transmit and four receive channels, enabling beamforming and angular estimation that improve localization accuracy. This capability is particularly useful for distinguishing multiple targets and resolving subtle motion signatures. The device's compact package is well suited for embedded designs, and its wide automotive-grade temperature range (−40 °C to +125 °C) exceeds the system's operational requirements. With typical active power consumption around 1 W, it remains within the 10 W power budget and supports the high sampling rates necessary to achieve sub-200 ms latency.

Despite these advantages, the BGT60ATR24C lacks an AiP, requiring external patch antennas fabricated on high frequency, low loss substrates. These antennas must be precisely dimensioned and tuned to achieve the desired beam pattern and gain, which significantly increases RF design complexity and cost compared with modules that integrate antennas. Errors in substrate selection, trace impedance, or antenna layout can degrade system sensitivity, detection range, and angular resolution. Additionally, because the device outputs only raw IF signals over an SPI interface and does not include on chip digital signal processing, all radar processing (including range FFTs), doppler extraction, micro motion isolation, and classification, must be implemented on an external processor or dedicated DSP. This requirement places considerable computational demand on the host microcontroller, such as the ESP32, and may necessitate additional processing resources to maintain real time performance.

Infineon provides basic driver support and evaluation tools for the BGT60ATR24C but does not supply application-level software or prebuilt algorithms for our stretch goal for vital sign detection. Consequently, respiration and heartbeat monitoring would require developing the entire signal processing and classification pipeline from the ground up, which would increase development time and technical risk. While the BGT60ATR24C's additional channels and potential for enhanced angular resolution offer clear advantages over the BGT60TR13C, these benefits are offset by more complex RF design requirements, the need for specialized materials, and substantial firmware development. Although the device could theoretically meet the range and latency requirements, including detection through drywall, achieving this level of performance would require significant engineering effort and would likely extend the project timeline beyond the scope of a prototype stage student project.

3.4.3.3. TI IWR6843AOP EVM

The Texas Instruments IWR6843AOP evaluation module is a fully integrated 60–64 GHz FMCW radar solution designed to accelerate the development of advanced sensing applications such as vital-sign detection and presence monitoring. It incorporates three transmit and four receive channels with a factory-calibrated antenna-on-package (AoP), eliminating the need for external antenna design and significantly reducing RF layout complexity. In addition to the radar front end, the module integrates a C674x digital signal processor (DSP) and an Arm Cortex-R4F microcontroller core, enabling on board signal processing that includes range FFTs, doppler estimation, and micro doppler extraction in real time. This on chip processing capability allows the radar to deliver processed data directly over UART or USB interfaces, reducing the computational load on the host microcontroller and simplifying system-level integration.

The IWR6843AOP EVM supports end-to-end latencies of under 200 ms, which meets the project's real-time performance requirements. It operates within the unlicensed 60 GHz ISM band under FCC Part 15 and is rated for industrial temperature ranges from −40 °C to +85 °C. With an active power consumption of approximately 2.0–2.5 W, it fits comfortably within the system's 10 W power budget, and its 9 x 6 cm footprint is acceptable to incorporate into the overall design. Moreover, Texas Instruments provides an extensive software ecosystem for this platform, including the mmWave SDK and mmWave Studio GUI, which streamline configuration, data acquisition, and algorithm development. Pre-validated reference labs for vital sign and occupancy detection allow rapid prototyping without requiring the development of signal processing algorithms from scratch. TI has also demonstrated respiration detection through drywall at distances of 2 m or more, validating the module's suitability for stretch objectives of this project. Although its cost is higher than that of discrete radar ICs, the IWR6843AOP's high level of integration, mature software support, and proven performance substantially reduce development time and technical risk, making it a strong candidate for rapid deployment in a prototype system.

3.4.3.4. TI AWR1843BOOST EVM

The Texas Instruments AWR1843BOOST is a closely related development platform that offers many of the same hardware and software advantages as the IWR6843AOP EVM. It integrates three transmit and four receive channels along with an on-chip C674x DSP and Arm Cortex-R4F core, enabling real time signal processing, range doppler estimation, and micro motion analysis directly on the device. The AWR1843BOOST supports similar end-to-end latency and leverages the same mmWave SDK and mmWave Studio GUI, providing equivalent software support and development workflows. With an active power consumption of approximately 2.5 W and a footprint of about 10 x 6 cm, it also aligns with the goals and objectives of this project.

The primary difference between the AWR1843BOOST and the IWR6843AOP lies in its operating frequency band. The AWR1843BOOST operates in the 76–81 GHz automotive radar band, which offers slightly higher range resolution due to the shorter wavelength but is subject to stricter regulatory constraints and licensing requirements. This frequency range is optimized for automotive sensing applications, such as adaptive cruise control and collision avoidance, rather than detecting human presence. As a result, while the AWR1843BOOST delivers similar performance and processing capabilities, its operation outside the unlicensed 60 GHz ISM band makes it less aligned with the objectives of this project. Additionally, its higher cost compared with the IWR6843AOP further reduces its suitability for a cost-sensitive, portable prototype. Overall, although technically capable, the AWR1843BOOST is less optimal for the intended use case due to regulatory limitations and application focus.

Radar Module	Infineon BGT60TR13C	Infineon BGT60ATR24C	TI IWR6843AOP EVM	TI AWR1843BOOST
Frequency (GHz)	60	60	60–64	76–81
Channels (Tx/Rx)	3-Jan	4-Feb	4-Mar	4-Mar
On-Chip Processing	No	No	C674x DSP + ARM R4F	C674x DSP + ARM R4F
Power (W)	0.5–0.6	1	2.0–2.5	2.5
Range Capability	≈ 2 m (requires DSP)	≥ 3 m (external antennas)	≥ 2 m through drywall	≥ 2 m (automotive band)
Size (mm)	6.5 × 5.0	8.0 × 8.0	90 × 60	100 × 60
SDK Support	Drivers only	Drivers only	Full mmWave SDK + GUI	Full mmWave SDK + GUI
Approx. Price (USD)	$25 (chip only)	$45 (chip only)	$299 (EVM board)	$349 (EVM board)

Table 12: Radar Module Comparison Table

3.4.3.5. FMCW Principle and Range-Doppler Processing

FMCW radar is a widely used technique for range and velocity sensing in automotive, industrial, and biomedical applications due to its high accuracy, continuous operation, and its relatively low power consumption. Unlike pulsed radar, which transmits data using discrete bursts of energy, FMCW radar continuously transmits a frequency modulated signal whose instantaneous frequency varies linearly with time, forming what is known as a chirp. The transmitted signal is defined as:

$$s(t) = A_t cos\left[2\pi \left(f_0 t + \frac{B}{2T_c} t^2 \right) \right]$$

where A_t is the transmit amplitude, f_0 is the start frequency, B is the total sweep bandwidth, and T_c is the chirp duration.

As the signal circulates through space and reflects off a target at range R, it experiences a round-trip time delay defined as:

$$\tau = \frac{2R}{c}$$

where c is the speed of light. When the reflected signals are mixed with a copy of the transmitted signal, the result is an intermediate frequency (IF) signal whose frequency is proportional to the time delay, and consequently to the target range. This beat frequency is expressed as:

$$f_b = \frac{B\tau}{T_c} = \frac{2BR}{cT_c}$$

From this relationship, we can derive the range R to the target as:

$$R = \frac{cT_c}{2B} f_b$$

In addition, the range resolution of the radar system is determined by the total sweep bandwidth and is given by:

$$\Delta R = \frac{c}{2B}$$

A larger sweep bandwidth enables a finer discrimination between closely spaced targets. To process the received signal, the radar samples the IF signal and applies the Fast Fourier Transform (FFT) across each chirp. This transforms the time domain signal into the frequency domain, allowing the system to resolve multiple targets as discrete range bins, each corresponding to a unique beat frequency. A secondary FFT, known as the Doppler FFT, is administered to extract velocity across multiple chirps to observe how the signal phase evolves. Together, this signal processing approach enables FMCW radar to detect, localize, and track multiple static and dynamic targets simultaneously with high precision. All four radar modules considered implement this FMCW principle, but those without on chip DSP (such as the Infineon devices) require the host MCU to perform these FFT and doppler operations externally, whereas the TI modules execute them internally and deliver processed data directly.

3.4.3.6. Micro-Doppler Extraction for Vital signs

Vital sign detection depends on the ability to sense millimeter-scale chest displacements caused by respiration and heartbeat. These minute motions create periodic phase modulations in the returned signal that appear as low frequency oscillations superimposed on the static reflection. The DSP isolates this micro doppler component by applying high pass filtering to remove clutter and low pass filtering to capture the breathing and heartbeat bands (typically 0.1–0.5 Hz for respiration and 1–2 Hz for cardiac motion). The resulting waveform is processed through time domain analysis or FFT based spectral estimation to output real-time breathing and pulse rates. In devices like the BGT60TR13C and BGT60ATR24C, these processing tasks must be implemented manually on the external MCU, whereas the IWR6843AOP and AWR1843BOOST handle micro doppler extraction internally, simplifying integration and reducing software complexity.

3.4.3.7. On Chip Processing and Firmware Framework

The TI IWR6843AOP radar integrates a C674x DSP and an ARM Cortex-R4F microcontroller core to handle signal processing and system control in parallel. The DSP executes computationally intensive tasks such as FFTs, filtering, and peak detection, while the ARM core manages radar configuration, data framing, and communication with the external ESP32 microcontroller via UART. This dual-core architecture allows one processor to continuously compute the next radar frame while the other handles data output, maintaining end-to-end latency below 200 ms. The TI mmWave SDK provides a layered software framework that abstracts low-level hardware control from higher-level application logic. Its components include the mmWave link driver for device-level commands, a control layer for interpreting configuration messages, and an application layer for processing and communication. Pre-validated reference labs within the SDK accelerate development of occupancy and vital-sign monitoring applications, while the mmWave Studio GUI enables rapid configuration, data capture, and visualization. Together, these tools reduce the need for custom firmware development and shorten the path from prototype to functional system.

3.4.3.8. Chirp Design, Data Handling, and Integration

Radar performance depends heavily on chirp design parameters such as slope, bandwidth, and repetition period, which together determine range resolution and frame rate. Increasing bandwidth improves range resolution $\Delta R = \frac{c}{2B}$ but at the cost of computational load. Through experimental calibration, chirp parameters will be tuned to maintain sub 200 ms update intervals while potentially detecting vital signs at distances ≥ 2 m through drywall.

Processed radar data are streamed as Type–Length–Value (TLV) packets over a UART interface. Each TLV block contains metadata such as frame number, object count, range–Doppler matrices, and micro-Doppler waveforms, which the system uses to extract and display motion parameters. Integration between the radar module and the ESP32-WROOM microcontroller forms the backbone of the embedded data pipeline. The ESP32 was selected for its dual-core architecture, low power consumption, and robust support for UART, SPI, I2C, Wi-Fi, and Bluetooth, enabling seamless communication between the sensing front end and the control or cloud layers.

A primary UART channel operating at 921,600 baud is used to transmit radar data from the IWR6843AOP to the ESP32. One core of the ESP32 is dedicated to UART reception and preprocessing, while the second handles higher level tasks such as wireless communication or user interface updates. A secondary UART or SPI interface is reserved for radar configuration, firmware updates, and diagnostic output. SPI offers higher bandwidth for future scalability, but UART is prioritized in this prototype for its simplicity, native support within the ESP-IDF development framework, and sufficient throughput for the system’s expected data rates.

Precise synchronization between the radar and MCU is achieved using interrupt-driven GPIO signaling. A “frame ready” pulse from the radar notifies the ESP32 to initiate a data read, ensuring low-latency data handling and consistent timing alignment across measurement cycles. Additional GPIO status lines provide feedback on radar states such as chirp activity or fault detection, allowing the MCU to respond immediately without continuous polling. The system also employs controlled power sequencing to ensure stable startup: the 3.3 V logic rail is enabled first, followed by the analog front end. A dedicated reset line controlled by the ESP32 reinitializes the radar in case of communication timeouts or thermal faults, further improving system robustness. These integration strategies are significantly simplified by TI’s development kits, which provide defined communication protocols and structured data formats, whereas integrating Infineon’s devices would require custom data framing and additional software layers to manage raw IF data streams.

3.4.3.9. Mechanical, EMI, and Thermal Design Constraints

Operating at 60 GHz imposes stringent electromagnetic and mechanical design requirements. High frequency traces can radiate or couple noise into digital lines, degrading performance. To mitigate this, the radar module will be isolated from digital circuits and enclosed in a grounded metal shield, which also provides mechanical stability and environmental protection. Ground planes beneath the radar are continuous and connected through fences to minimize impedance discontinuities. Analog, digital, and power domains are carefully separated to reduce cross coupling. The radar dissipates approximately 2 W of power during operation, which is managed using thermal vias or small plated through holes that conduct heat into internal copper layers and the enclosure. The chassis will be constructed from aluminum or copper-plated polymer for efficient heat dissipation while remaining lightweight. The system is designed to operate reliably between −40 °C and +85 °C, exceeding the 0–50 °C requirement for the application. Validation testing will include temperature cycling and vibration tests to ensure vigorous performance in challenging environments. While all candidate radars share similar thermal and EMI considerations, the integrated design of the IWR6843AOP reduces RF layout complexity and minimizes the risk of impedance mismatches, which would be more challenging to control when implementing external antennas for the BGT60ATR24C.

3.4.3.10. Radar Selection Justification

After evaluating candidate radar devices, we compared the tradeoffs between integrating a bare mmWave radar IC and using a fully integrated development kit. Bare ICs such as Infineon's BGT60TR13C and BGT60ATR24C offer low unit cost and compact size, but they shift critical design burdens onto the development team, including RF front end design, antenna tuning, and signal processing implementation. These tasks require high frequency PCB layout on specialized substrates with controlled impedance and strict trace tolerances, where even small design errors can severely degrade performance. They also require development of custom signal processing algorithms for vital signs and presence detection from the ground up. In comparison, development kits like TI's IWR6843AOP and AWR1843BOOST integrate these elements into a single platform, which not only eliminates most of the RF and antenna challenges but also offloads significant signal-processing work from the host MCU. The TI mmWave SDK, mmWave Studio GUI, and validated reference labs provide ready to use tools for configuring chirps, visualizing data, and building classification algorithms, enabling rapid system level progress. Unlike Infineon's options, which provide only low-level drivers and require building the processing pipeline from scratch, TI's ecosystem shortens development time by providing pre-validated application code and visualization utilities. Although the development kit has a higher cost, it eliminates hidden costs such as custom PCB fabrication, RF test equipment, and extended debugging time. For a student project with tight timelines, the overall cost to benefit ratio strongly favors the development kit. Furthermore, once a proof of concept is validated, firmware and algorithms developed on the kit can be directly ported to a custom board using the same IC, lowering cost in production.

After evaluating multiple radar options, the Texas Instruments IWR6843AOP EVM was selected as the radar front end for our system. It uniquely combines AoP integration, on chip DSP, and a mature software ecosystem at 60 GHz, features that collectively minimize RF design complexity and accelerate development. Its proven capability to detect vital signs through drywall at distances of 2 m or more aligns directly with our stretched application requirements. Compared with Infineon's BGT60TR13C and BGT60ATR24C, which both require external antennas, off chip signal processing, and extensive firmware development, the IWR6843AOP provides a plug and play platform that integrates these elements into a single, validated solution. Although the Infineon devices offer lower unit cost, their reliance on external DSPs and complex PCB design would significantly increase development time, require specialized fabrication materials, and raise the risk of performance degradation due to RF layout errors. Likewise, while the AWR1843BOOST offers similar hardware integration, its operation in the licensed 76–81 GHz band introduces regulatory challenges and is optimized for automotive radar rather than sensing human presence, making it less aligned with the system's intended application.

The IWR6843AOP EVM enables us to achieve first measurements immediately, focus engineering effort on system level integration and maintain a clear migration path to a custom board for future production. Furthermore, it's built in SDK, structured data output, and well documented communication protocols significantly simplify integration with the ESP32 microcontroller, a level of software support not available for the Infineon alternatives. In terms of long-term scalability, firmware and algorithms developed on the IWR6843AOP can be directly ported to a custom board using the same IC, ensuring a smooth transition from prototype to production without re-architecting the signal processing pipeline. This combination of performance, cost efficiency, regulatory suitability, and ecosystem support makes the IWR6843AOP the optimal radar choice for our prototype and the most practical solution to meet the system’s sensing, latency, and integration requirements.

3.4.4. IR Sensor

3.4.4.1. AMG8833 IR Thermal Camera (Grid EYE)

The Adafruit AMG8833 is an 8×8-pixel infrared thermal sensor array that detects emitted heat energy and converts it into temperature data transmitted over an I2C interface. Operating at 3.3 V logic and drawing less than 10 mA, it is fully compatible with the ESP32WROOM microcontroller and ideal for low power embedded applications. This component will measure temperatures ranging from 0°C to 80°C (32°F to 176°F) with an accuracy of +- 2.5°C (4.5°F). It can detect a human from up to 7 meters (23) feet and has a maximum frame rate of 10Hz. In this project, the AMG8833 is programmed to identify temperatures within the typical human body range (30 °C - 40 °C), triggering a “Human Detected” message on the LCD display when such readings are observed. Its straightforward I2C communication, compact design and existing Arduino and MicroPython library support make it a highly practical choice for integration. This unit costs about $50 making it an affordable option. While it can produce low resolution thermal maps as a part of our stretch goals, its primary advantage lies in its fast and reliable detection of heat signatures for life verification in conjunction with the other radar subsystems.

3.4.4.2. MLX90640 32×24 Thermal Camera Array

The MLX90640 offers a higher resolution 32x24 pixel thermal grid, producing more detailed thermal information than the AMG8833. It also operates over I2C and runs at 3.3 V logic, compatible with the ESP32WROOM. The module supports frame rates from 0.5 Hz to 16 Hz and provides a detection range of roughly 1 m - 7 m with a field of view between 55° and 110°, depending on the lens version. Typical power consumption ranges from 18–23 mA, making it moderately efficient but still suitable for battery powered applications. The additional resolution would allow more refined temperature mapping or a smoother heat map display as a part of our stretch goal. However, this improvement comes with higher computational demands and a cost around $80, which is more than the AMG8833 but still falls within the project budget. For a system prioritizing fast and reliable presence confirmation over detailed imaging, the MLX90640 may exceed the project’s immediate needs, though it remains an excellent candidate.

3.4.4.3. MLX90641 16×12 Thermal Array

The MLX90641 is a midrange alternative that balances resolution and cost between the AMG8833 and the MLX90640. It provides a 16×12 thermal grid, uses the same I2C communication protocol, and operates at 3.3 V logic. With a typical field of view of 55° and a detection range of up to 7 m, it offers better spatial accuracy than the AMG8833 while remaining easier to process than the MLX90640. Power consumption is moderate at around 25 mA, and the sensor includes built-in calibration for temperature stability. For this project, it could enhance detection confidence by mapping larger areas, but the smaller pixel count compared to the MLX90640 limits fine detail for visualization tasks. Its moderate price point of $40 and simpler data handling makes it a strong backup option.

3.4.4.4. FLIR Lepton 3.5 Thermal Imaging Module

The FLIR Lepton 3.5 is a professional grade micro thermal camera core offering 160×120 resolution and radiometric output. It interfaces via SPI through a dedicated breakout board, such as the PureThermal 2, and supports 3.3 V logic. Although the FLIR Lepton 3.5 delivers excellent image quality and can operate through light smoke, it draws a relatively high current of about 150 - 200 mA, which makes it less practical for battery powered or portable systems where energy efficiency is a concern. The Lepton provides the most detailed heat data among the compared units and can detect temperature variations from several meters away with precision under +-2 °C. While this would allow full image thermal visualization on the LCD, integration complexity and price are significantly higher, exceeding $150 - $200 once a carrier board is included. Due to its processing requirements and higher cost, the Lepton is better suited for advanced prototypes where high-definition thermal imaging is necessary.

3.4.4.5. Infrared Temperature Sensor (MLX90614)

The Grove MLX90614 is a single point, contactless infrared thermometer capable of measuring surface temperatures from –70 °C to 380 °C. It operates at 3.3 V - 5 V and communicates via I2C. Power consumption is minimal at roughly 1.5 - 2 mA, making it highly efficient for low power systems. Rather than providing a pixel array, it outputs an average temperature reading of the target in its 35° field of view. While limited to one measurement point, its simplicity, low power usage, and $20 cost make it useful for basic proximity or presence detection based on a single temperature threshold. In this project, it could serve as a compact alternative for early testing but lacks the spatial awareness needed for multi-point detection or mapping.

3.4.4.6. Limitations

The limitation of infrared (IR) thermal sensors is that they cannot detect heat sources through solid materials such as walls, doors, or glass. This is because infrared radiation is absorbed or reflected by these surfaces, preventing the sensor from capturing the temperature of objects behind them. As a result, the IR module is limited to the line-of-sight detection, sensing only the surface temperature of whatever is directly visible within its field of view. To overcome this limitation, our team is exploring the integration of ultra-wideband (UWB) radar and similar radio frequency sensing technologies, which can detect motion and presence through most nonmetallic barriers. By combining radar-based presence detection with the IR sensor's precise thermal verification, the system can achieve a more reliable form of life detection. The radar subsystem would provide distance or motion data even when visibility is obstructed, while the IR module would confirm whether the detected object falls within a human temperature range. The complementary sensing strategy enhances detection accuracy, minimizes false positives, and improves performance across a range of environmental and operational conditions.

Component	**AMG8833(GRID-eye)**	**Melexis MLX90641**	**Melexis MLX90640**	**FLIP Lepton 3.5**	**Grove MLX90614 (single point)**
Array / Resolution	8x8 (64 px)	16x12	32x24	160x120	1 px (spot)
FOV (in degrees)	60 x 60	About 55	55 std or 110 wide	About 57x44	About 35
Interface	I2C	I2C	I2C	SPI	I2C
Interface	SPI	SPI	SPI +I2C	SPI	SPI
Voltage / Logic Compatibility	3.3 V	3.3 V	3.3 V	3.3 V	3.3 V
Current (typ.)	10 mA	25 mA	18-23 mA	150-200mA	1.5-2mA
Outdoor Readability	Fair	Good	Very good	Very good	Fair
Pros	Easy wiring, very low power, good libraries	More pixels, light compute	Best balance of detail vs ease	Pro-grade images	Ultra simple, ultra-low power, cheap
Cons	Low resolution, best for indoor use or shade only	Lower resolution than 90640; costs more than AMG	Heavier processing than AMG, costs more	Expensive, more code/power	No spatial info, not ideal for human detection alone

Table 13: IR Thermal Sensor Comparison Table

3.4.4.7. Selected IR Sensor Justification

All evaluated IR modules are compatible with the ESP32WROOM through I2C communication and operate safely at 3.3 V logic. The main tradeoffs among them involve resolution, range, outdoor reliability, processing demand, and cost. The Adafruit MLX90640 provides a strong balance between performance and practicality, offering a significantly higher 32x24 pixel resolution compared to lower end arrays like the AMG8833 while maintaining straightforward integration and low power requirements. Its wider field of view and built in ambient temperature compensation make it more stable and accurate in outdoor or highlight environments, where simpler sensors can become unreliable. Higher end options like the FLIR Lepton enable professional thermal imaging but require more complex drivers and significantly more power. The Grove MLX90614 remains the simplest and least expensive choice but lacks the spatial awareness necessary for reliable human detection.

For this prototype, the Adafruit MLX90640 was selected because it offers a strong combination of resolution, outdoor reliability, and ease of integration. It communicates over I2C, operates at 3.3 V logic, and draws only about 20 mA, making it directly compatible with the ESP32WROOM without requiring additional hardware. The MLX90640's 32×24 thermal array allows for clearer detection patterns, improving accuracy in identifying human heat signatures even under bright sunlight or varying temperatures. Its built-in calibration and available Arduino and MicroPython libraries make implementation straightforward, enabling fast setup and minimal code modification. The sensor can also support our stretch goals such as displaying low resolution heat maps on the LCD or logging temperature distributions over time. Overall, it provides a significant upgrade in capability over the AMG8833 while maintaining simplicity, efficiency, and full compatibility with the current system design.

3.5. User Interface

The user interface (UI) provides the operator with tactile, visual, and auditory interaction with the system. The device is designed for search and rescue environments, so its interface must remain fully functional even in challenging conditions like low visibility, operation through gloves, and high ambient noise. To meet these demands, the UI incorporates five essential components: pushbuttons, LEDs, a buzzer, a speaker, and an LCD display module.
The pushbuttons serve as primary input controls, allowing operators to initiate scans, toggle modes, or reset the system with ease. LEDs and
the buzzer provide immediate visual and auditory feedback for system states and alerts, while the speaker supports more detailed audio cues or status messages. The LCD display module offers clear, structured data visualization, presenting vital information such as sensor readings, system status, and detected targets in real time. Collectively, these components create a redundant and multimodal interface that guarantees reliable user feedback and control, even when one sensory channel, such as sight or hearing, is compromised by environmental factors.

3.5.1. Pushbutton

Pushbuttons were selected as the primary input mechanism due to their mechanical simplicity, high reliability, and superior tactile feedback in demanding field conditions. Unlike capacitive or touch-sensitive controls, which can become unresponsive when operated through gloves or contaminated surfaces, mechanical pushbuttons provide a distinct physical click that confirms successful actuation. This tactile confirmation is critical in search and rescue scenarios where first responders often wear heavy protective gloves, turnout gear, or bulky sleeves that reduce fine motor control. The ability to engage controls without precise finger positioning or visual confirmation is essential for usability and safety in these environments. Their straightforward electrical interface, based on momentary contact closure, also makes pushbuttons easy to debounce and integrate into low power embedded systems. Additionally, sealed pushbuttons offer excellent resistance to dust, moisture, and vibration, making them well suited to the harsh environmental conditions encountered during field deployment. While membrane keypads or capacitive sensors might reduce moving parts, they would perform poorly under the thick gloves and limited dexterity common among firefighters and rescue personnel, and they are more prone to false triggers from debris or moisture.

3.5.2. LED Indicator

LED indicators were chosen as the primary feedback mechanism for this system because they provide instant response, low power consumption, and strong reliability in demanding field environments. They give immediate visual confirmation of key system states such as power (green), scanning (yellow), and detection (red), which is especially important when audible alerts might be drowned out by ambient noise or when tactile feedback alone is not enough. For first responders working in smoke-filled, dark, or debris covered environments who often while wearing helmets, face shields, and respirators, maintaining visibility is a serious design requirement. LEDs help meet that challenge by producing bright, directional light with wide viewing angles, allowing status information to remain visible even when the device is not positioned directly in front of the user or is viewed through layers of protective gear.

Another reason LEDs are ideal for this application is that they simplify the overall system design. Surface mount device (SMD) LEDs reduce panel wiring, streamline PCB layout, and maintain mechanical durability under shock and vibration, all of which are crucial for a portable device that might be dropped, jostled, or used in extreme conditions. While alternatives like LCD or OLED panels could display more detailed information, they consume more power, are harder to see through masks and face shields, and are generally less rugged in harsh search and rescue conditions.

Brightness is also a key consideration. Luminous intensity in the 20–80 mcd range is usually sufficient for handheld devices at close range, but higher output becomes necessary in smoke, bright sunlight, or other high-contrast environments. Wide viewing angles further improve peripheral visibility and reduce the need to precisely orient the device, a valuable feature when responders are focused on victims rather than equipment. The LED families under consideration provide a range of luminous intensities and maintain wide viewing angles, allowing us to tailor individual indicators for continuous status signaling or high priority alerts.

Finally, package size plays an important role in balancing performance with physical design constraints. Most of the LEDs being evaluated are available in standard 0603 (1.6 × 0.8 mm) and 0805 (2.0 × 1.25 mm) SMD packages. 0603

LEDs are excellent for space-constrained layouts and automated assembly, while 0805 LEDs offer greater luminous output and improved mechanical strength, which is especially useful in ruggedized systems. Choosing components from families that offer both sizes gives us flexibility to trade off brightness, board space, and mechanical durability without sacrificing electrical compatibility or manufacturing efficiency.

3.5.3. Buzzer

A buzzer was included as a core part of the user interface to provide reliable auditory feedback in conditions where visual or tactile cues may be missed. Sound is a critical alert channel in low visibility environments such as collapsed structures or smoke-filled rooms, where first responders are wearing face shields, respirators, and heavy gear that limit their field of view. Unlike visual indicators, sound can penetrate obstacles and immediately attract attention, even when the device is not in direct line of sight. By generating distinct tone patterns, such as short chirps for power on, continuous tones for life detection, or rapid bursts for fault conditions, the system can convey status information automatically without requiring the operator to read a display.

In addition to volume, tone control and sound profile flexibility are important considerations. A buzzer that can be modulated by the microcontroller allows different alerts to be assigned to different events, reducing cognitive load and improving reaction time. Other design factors include current consumption, physical footprint, and mounting compatibility with the device’s enclosure.

3.5.4. Speaker

A dedicated speaker is being considered in the user interface to provide continuous or modulated audio feedback that communicates more complex information than a simple buzzer can. While the buzzer is ideal for quick alerts and system status notifications, the speaker enables real time auditory representation of radar derived signals, such as breathing or heartbeat waveforms. This feature provides first responders with an intuitive sense of life presence even when they cannot look at a display. For example, while navigating debris, carrying equipment, or wearing full protective gear that limits visibility and skill. For this project application, the speaker must balance several requirements: clarity, efficiency, durability, and compact.

3.5.5. Display Screens

An LCD screen is a flat panel display technology that uses liquid crystals to control light and present information such as text, numbers or graphics. The liquid crystals twist or untwist when an electrical signal is applied, altering how much light can pass through. A backlight, or ambient light in some designs, provides illumination while polarizers and color filters shape the light into visible images. In essence, an LCD functions as a light gate controlled by electronics, enabling it to display clear and dynamic visual information.

Although there are several different display technologies available, we chose to focus on LCDs because they align well with our course history, prior experience and existing background knowledge. By leveraging a technology, we are already familiar with, we can efficiently develop a working product by ensuring reliable integration into our system.

For this project, the LCD serves as a critical component of the user interface, providing detailed feedback that cannot be conveyed by LEDs or buzzers alone. At the basic level, the display must clearly present simple text outputs such as “Human Detected” “Hazard” or “Safe Zone” ensuring that the user receives unambiguous information in real time. As the system progresses to advanced goals, the LCD should be capable of showing multiple sensor readings simultaneously, supporting simple icons for hazards, and offering improved clarity for field use. For stretch goals, the display would need to handle more complex visualizations, such as logged event history, battery monitoring, or graphical overlays for sensor data trends. To meet these requirements, the LCD must be compatible with the ESP32WROOM microcontroller, supporting standard communication protocols such as SPI or I2C for seamless integration.

3.6. User Interface Parts Selection

3.6.1. Pushbutton

3.6.1.1. Apem IPR1SAD2

The Apem IPR1SAD2 is a compact, IP67 rated sealed tactile pushbutton designed for rugged environments. It features a 200-gf actuation force, providing a firm, positive tactile response that remains perceptible through the heavy gloves and layered protective clothing worn by first responders. The firm actuation pressure helps prevent accidental activation while still allowing easy operation with limited dexterity. Gold plated contacts resist corrosion and ensure long term electrical reliability, while silicone sealing provides robust protection against dust and moisture ingress. With a mechanical life exceeding 100,000 cycles, the switch supports extended field use without significant performance degradation. Its small footprint simplifies integration into compact enclosures, which is beneficial for handheld or wearable devices. However, its smaller surface area requires more deliberate finger placement than larger industrial switches, which could make it slightly less forgiving in situations where precise contact is difficult. Despite this, the IPR1SAD2 strikes a strong balance between ruggedness, tactile feedback, and glove compatibility, making it well suited to our use case.

3.6.1.2. E-Switch PV6F240SS-331

The E-Switch PV6F240SS-331 is a vandal resistant stainless-steel pushbutton engineered for extreme outdoor and industrial environments. It offers IP67 sealing, high impact resistance, and optional integrated LED illumination for visibility in low light conditions. Its large actuation surface and robust construction make it highly compatible with gloved operation, allowing first responders to activate the control even with limited dexterity or partial contact. This makes it particularly helpful in high stress, low visibility environments. However, the switch's 16 mm mounting diameter increases panel space requirements and overall weight, which could present challenges in compact or handheld systems. Additionally, while its large size improves usability with thick gloves, it may be excessive for a portable, lightweight prototype where multiple controls must be integrated into a small enclosure.

3.6.1.3. C&K PTS645VL58-2 LFS

The C&K PTS645VL58-2 LFS is a standard 6 x 6 mm low profile tactile switch that offers a crisp tactile feel and long mechanical life at low cost. Its compact size and simple mounting make it easy to incorporate into dense circuit board layouts. However, this switch lacks environmental sealing and is vulnerable to dust and moisture access, which significantly limits its suitability for field deployed devices. More importantly, its small size and low actuation force make it difficult to operate reliably with gloved hands or reduced dexterity, a critical requirement in first responder environments. While its low cost and compact footprint make it a useful option for early-stage prototyping or indoor applications, it does not meet the environmental or ergonomic requirements of a ruggedized search and rescue device.

Part Number	Manufacturer	Actuation Force	Sealing	Lifespan	Unit Price
Apem IPR1SAD2	Apem	200 gf	IP67	100k cycles	3.85
E-Switch PV6F240SS-331	E-Switch	350 gf	IP67	50k cycles	7.6
C&K PTS645VL58-2 LFS	C&K	260 gf	None	100k cycles	0.22

Table 14: Pushbutton Comparison Table

3.6.1.4. Pushbutton Selection Justification

The Apem IPR1SAD2 was selected as the primary pushbutton for the system because it offers the best overall balance of durability, tactile response, environmental resilience, and ergonomic suitability for first responder use. Its IP67 sealing ensures consistent operation under exposure to rain, dust, and debris, while its 200-gf actuation force provides a firm, positive response that remains perceptible through the thick gloves and layered protective gear commonly worn in search and rescue operations. Compared with the E-Switch PV6F240SS-331, which offers excellent glove compatibility but at the expense of significantly larger size, heavier weight, and increased panel space requirements, the Apem IPR1SAD2 achieves similar tactile reliability in a more compact and lightweight form factor. Likewise, while the C&K PTS645VL58-2 LFS offers a lower cost solution, its lack of sealing, smaller actuation surface, and limited usability with gloved hands make it unsuitable for field deployment in harsh conditions.

By meeting all mechanical, electrical, and ergonomic requirements without sacrificing compactness or ease of integration, the Apem IPR1SAD2 provides a robust and field ready solution that supports reliable operation in mission-critical environments. Its combination of rugged environmental sealing, glove-compatible tactile feedback, and compact size makes it the optimal choice for the pushbutton interface in this prototype.

3.6.2. LED

3.6.2.1. Kingbright KP-1608 Series (Green, Yellow, Red, Blue, Amber)

The Kingbright KP-1608 series is a family of 0603-package SMD LEDs known for their compact size, energy efficiency, and availability in multiple colors, including green (KP-1608CGCK), red (KP-1608SRCP), yellow (KP-1608SYCK), blue (KP-1608SURCK), and amber. The green variant offers approximately 30 mcd luminous intensity with a 120° viewing angle, providing excellent visibility in low light or indirect viewing conditions, while the red and yellow versions achieve similar performance at about 25–50 mcd and 25 mcd, respectively. Operating at a typical current of 2 mA, these LEDs are highly power efficient and appropriate for battery powered field devices. Their color diversity allows designers to standardize the footprint while customizing status indicators for optimal human factor performance. Compared with higher intensity 0805 LEDs, the KP-1608 series sacrifices some brightness but enables tighter PCB layouts and reduced power draw, making them a strong choice for continuous indicators such as “Power On” or “Scanning.”

3.6.2.2. Lite-On LTST-C191 Series (Red, Orange, Yellow, Green, Blue)

The Lite-On LTST-C191 series is another 0603-package LED family that emphasizes stable hue and reliable performance across a wide current range (2–20 mA). The yellow LTST-C191KSKT variant offers a luminous intensity of about 15 mcd, while the red and green versions provide around 20–30 mcd and 30 mcd, respectively. The series maintains consistent color performance under thermal stress, which is important in devices that may experience rapid temperature fluctuations in the field. Although these LEDs are slightly less bright than the Kingbright or OSRAM options, their thermal stability and color consistency make them suitable for indicators that must remain clearly visible during prolonged operation. The wide color range also enables flexible state assignment, which can be refined during field testing based on responder feedback.

3.6.2.3. OSRAM R976 Series (Red, Yellow, Orange, Blue)

The OSRAM R976 series is a family of high brightness 0805-package LEDs optimized for critical alerts and high visibility indicators. The LS R976 red LED delivers up to 80 mcd luminous intensity with a 120° viewing angle, ensuring immediate visibility even through smoke, tinted visors, and peripheral vision. Yellow and orange variants offer luminous intensities of about 60–70 mcd and 60 mcd, respectively, while maintaining similar electrical characteristics. The slightly larger 0805 package increases mechanical robustness and provides enhanced heat dissipation, improving reliability under continuous use. Compared with 0603 options, the R976 series offers significantly higher brightness and enhanced visual impact, making it ideal for safety critical alerts such as “Detection Active” or “System Fault.” The tradeoffs include increased board space and marginally higher power consumption, but these are acceptable for high priority indicators that demand immediate attention.

Part Number	Manufacturer	Color	Package	Luminous Intensity	Viewing Angle	Forward Voltage	Unit Price
KP1608CGCK	Kingbright	Green	0603	30 mcd	120°	2.1 V	$0.09
KP1608SYCK	Kingbright	Yellow	0603	20 mcd	120°	2.0 V	$0.09
KP1608SURC	Kingbright	Red	0603	40 mcd	120°	2.0 V	$0.10
LTST C191KGKT	Lite-On	Green	0603	25 mcd	120°	2.1 V	$0.07
LTST-C191KSKT	Lite-On	Yellow	0603	15 mcd	120°	2.0 V	$0.07
LTST-C191KRKT	Lite-On	Red	0603	30 mcd	120°	2.0 V	$0.08
LS R976	OSRAM	Red	0805	80 mcd	120°	2.0 V	$0.12
LY R976	OSRAM	Yellow	0805	60 mcd	120°	2.0 V	$0.12
LG R976	OSRAM	Orange	0805	50 mcd	120°	2.1 V	$0.12

Table 15: LED Comparison Table

3.6.2.4. LED Selection Justification

The OSRAM LS R976 series was selected as the primary LED indicator family for this system due to its superior luminous output, mechanical resilience, and reliability in harsh environments. Its 50 - 80 mcd brightness and 120° viewing angle ensure that visual signals remain clearly visible even through helmets, face shields, and smoke, addressing one of the most critical ergonomic challenges faced by first responders in the field. The larger 0805 package provides increased luminous intensity and mechanical durability compared to smaller 0603 alternatives, making it more suitable for mission critical alerts where visibility cannot be compromised.

When compared to the Kingbright KP-1608CGCK or Lite-On LTST-C191KSKT options, which offer lower brightness overall and smaller package sizes optimized for compact electronics, the OSRAM LEDs provide significantly better visibility in low visibility or high glare conditions and superior resistance to vibration and shock. Their viewing angle also reduces the need for precise device alignment, which is particularly valuable in dynamic search and rescue scenarios where the operator’s attention is focused elsewhere.

In addition, the LS R976 series is part of a broader OSRAM LED family available in multiple colors, enabling consistent mechanical, optical, and electrical performance across power, scanning, and detection indicators. This consistency simplifies design and manufacturing while maintaining uniform brightness and visibility across all system states. Together, these characteristics make the OSRAM LS R976 LEDs the most reliable and field-appropriate choice for this prototype.

3.6.3. Buzzer

3.6.3.1. CUI CMT-1203-SMT-TR

This passive piezo buzzer operates at 3 V and delivers an output of approximately 80 dB SPL, which is generally sufficient for handheld devices in noisy environments. Because it is passive, it requires a PWM signal from the MCU to generate sound, allowing full control over tone frequency and pattern. This flexibility is a major advantage in field devices that must communicate multiple system states. Its small surface mount footprint and low power draw make it suitable for portable, battery powered applications. Compared to active buzzers, the CMT-1203-SMT-TR demands more from the control firmware but rewards that complexity with significantly greater alert versatility.

3.6.3.2. Mallory Sonalert PS-1206P

This active piezo buzzer outputs around 85 dB SPL with a fixed 2.9 kHz tone. It is simple to implement, requiring only a DC voltage input and no PWM control, which reduces firmware complexity. However, its lack of tone modulation limits its usefulness in conveying multiple types of alerts. In an application where first responders must differentiate between power, detection, and fault states without looking at the device, this limitation is significant. Although its slightly higher sound pressure level improves audibility, its reduced signaling flexibility makes it less suitable for our multi-state system.

3.6.3.3. Soberton WT-1205

This magnetic buzzer achieves the highest output level at roughly 90 dB SPL, providing excellent audibility even in loud environments. However, this comes at the cost of increased current draw (about 60 mA) and larger physical size, which pose challenges for a compact, battery powered device. It is best suited for stationary systems or those where power availability is not a constraint. While its loudness is beneficial, the tradeoff in size and power consumption outweighs the advantages for a handheld search and rescue tool.

3.6.3.4. SameSky CMI-1295-03TH

The CMI-1295-03TH is a 12 mm diameter, through-hole magnetic audio indicator. It operates at a tone frequency of approximately 2.4 Hz and delivers a sound pressure level of approximately 85 dB at 10 cm when rated at 3 V DC. This device has a compact footprint, standard PC-pin termination, and a 3 V drive voltage, making it an ideal choice for embedded systems that require breadboard prototyping and later PCB integration. It features an internally driven oscillator, which means the firmware needs to actively provide power to generate an audible alert, eliminating the need for PWM. This characteristic makes it a strong candidate for our UI design, where reliable loud audible feedback is essential.

Part Number	Type	Voltage	Sound Level (dB)	Control	Current Draw	Unit Price
CUI CMT-1203-SMT-TR	Passive Piezo	3 V	80	PWM (variable tones)	20 mA	1.42
Mallory PS-1206P	Active Piezo	3 V	85	Fixed	25 mA	2.75
Soberton WT-1205	Magnetic	3–5 V	90	Fixed	60 mA	1.15
CMI-1295-03TH	Magnetic	3 V	85 dB	Fixed	30 mA	1.14

Table 16: Buzzer Comparison Table

3.6.3.5. Buzzer Selection Justification

The CUI CMT-1203-SMT-TR was selected for its balance of audibility, flexibility, and efficiency, which makes it ideal for a portable search and rescue device. Its 80 dB output is sufficient for field use while maintaining low power consumption, and its passive design allows the microcontroller to generate distinct tones and patterns for different system states. Compared with the Mallory PS-1206P, which limits signaling to a single tone, the CUI buzzer enables more intuitive communication of status information without visual confirmation. Likewise, while the Soberton WT-1205 is louder, its high current draw and larger footprint make it less practical for a handheld, battery-powered platform. By combining tone programmability, compact size, and efficient operation, the CMT-1203-SMT-TR provides the most versatile and field-ready alert solution for this application.

3.6.4. Speaker

3.6.4.1. CUI CDS-15108

This 15 mm dynamic speaker with 8 Ω impedance and 0.5 W power handling offers balanced audio output across a wide 450 Hz - 20 kHz range. It is lightweight, compact, and energy efficient, making it a good baseline option for embedded systems. However, its lack of environmental sealing makes it vulnerable to dust and moisture access, which is a significant drawback in outdoor or post disaster conditions. Compared with sealed alternatives, it trades ruggedness for simplicity and low cost.

3.6.4.2. Soberton SP-1608

The Soberton SP-1608 is slightly larger at 16 mm and rated for 0.7 W, this speaker delivers higher volume and improved low frequency response, enhancing the realism of heartbeat and respiration audio output. However, the increased output comes at the expense of higher current draw and reduced battery life. Like the CDS-15108, it lacks sealing, limiting its suitability in environments with water spray, fine debris, or high humidity. While it provides improved acoustic performance, the tradeoffs in power efficiency and durability make it less ideal for a rugged, battery powered device.

3.6.4.3. PUI Audio AS01508AO-WR-R

This 15 mm, 8 Ω dynamic speaker provides 0.5 W output and is rated IP65, offering strong resistance to dust and water. Its mylar diaphragm maintains consistent performance across temperature extremes and resists degradation from environmental contaminants. Despite its compact size and low power draw, it provides clear and accurate reproduction of low and mid frequency signals, making it highly effective for conveying vital sign data. Compared to non-sealed options, it offers superior reliability and operational continuity in field conditions without compromising on sound quality or efficiency.

Part Number	Type	Power	Impedance	Sealing	Frequency Range	Unit Price
CUI CDS-15108	Dynamic	0.5 W	8 Ω	None	450 Hz - 20 kHz	1.05
Soberton SP-1608	Dynamic	0.7 W	8 Ω	None	400 Hz - 18 kHz	1.28
PUI AS01508AO-WR-R	Dynamic	0.5 W	8 Ω	IP65	450 Hz - 20 kHz	2.10

Table 17: Speaker Comparison Table

3.6.4.4. Speaker Selection Justification

Although research was conducted on selecting the right speaker for our project, we ultimately decided to omit this component for several practical and system-level reasons. First, incorporating a dedicated speaker would be redundant given that our design already includes both a piezo buzzer and high-visibility LEDs to provide audible and visual alerts to first responders. These existing indicators effectively fulfill the same notification purpose without adding additional hardware complexity. Additionally, integrating a speaker would introduce extra requirements such as an audio amplifier stage, higher peak current consumption, and additional GPIO allocation on the ESP32, all of which would place unnecessary strain on our already constrained power and pin budget. A speaker would also require more detailed firmware development to generate meaningful audio patterns or voice prompts, which diverts development time away from more critical system features like radar processing and user-interface stability. Considering these factors, we concluded that the buzzer-and-LED alert scheme meets all functional requirements while keeping the design lightweight, power-efficient, and reliable. As a result, the speaker was intentionally removed from the final system architecture.

3.6.5. Display Screens

3.6.5.1. ILI9341 TFT LCD

A compact and well supported display option is the 3.2″ TFT LCD using the ILI9341 controller. It offers a 240×320 resolution and supports up to 262,000 colors which is sufficient for showing multiple sensor readings, status icons, and simple graphics. The ILI9341 includes built in frame memory for smooth refreshing, reducing the workload on the ESP32-WROOM. It supports both 8 bit parallel and SPI communication, though SPI is preferred for this design because it uses fewer GPIO pins and still provides adequate update speed. The module operates at 3.3 V logic, making it directly compatible with the ESP32-WROOM without level shifting. Common libraries such as TFT_eSPI and Adafruit GFX offer full driver support. The LED backlight typically draws 80-120 mA and can be dimmed through PWM control. While this 3.2″ model is cost-effective ($20–$30), its small size limits visibility when displaying multiple data windows, and the standard version is not optimized for outdoor readability. Overall, it’s a proven and straightforward option for compact GUIs.

3.6.5.2. 4.3" 800x490 RA8875 TFT LCD Module

A great option for a larger and higher quality display is the EastRising 4.3” TFT LCD module that uses the RA8875 controller. It has an 800×480 resolution, which gives a lot more screen space compared to smaller displays like the ILI9341. The built in RA8875 chip takes care of most of the display work, such as drawing shapes and handling frame refresh, which helps reduce processing load on the ESP32 WROOM. The RA8875 supports both SPI and 8/16-bit parallel communication modes. In this design, SPI is used to minimize wiring and conserve GPIO pins while maintaining sufficient update speed. The module runs at 3.3 V logic, which matches with the ESP32, and the LED backlight typically draws around 150-250 mA. The resolution and color depth make it a good choice for showing multiple readings, plots, or menu interfaces. However, since it’s not sunlight readable by default, it may require a shade or antiglare layer for outdoor use. The cost is around $40, which is reasonable for its size and capability.

3.6.5.3. 4.3” RA8875 Capacitive Touch Module

This version of the EastRising display adds a capacitive touchscreen, communicating via I2C for touch input while retaining SPI for graphics. Capacitive touch provides smoother, more responsive control similar to a tablet, making it ideal for menus or system adjustments on the screen. It operates at 3.3 V logic and draws roughly the same current as the non capacitive touch model (≈200 mA). Both SPI and I2C are natively supported by the ESP32-WROOM, so this module is fully compatible without additional hardware. The cost is around $50, representing a worthwhile upgrade if interactive touch control aligns with project goals.

3.6.5.4. 4.3” 800x490 IPS TFT LCD

The 4.3" 800×480 IPS TFT LCD module (ER-TFTMC043-3) includes a built-in RA8875 graphics controller, which makes it directly compatible with the ESP32-WROOM through the SPI interface. The RA8875 handles low level display tasks such as drawing, refreshing, and buffering image data, allowing the ESP32 to focus on other system functions like sensor processing and communication. The IPS panel provides excellent color accuracy and wide viewing angles, making it suitable for viewing from multiple positions or in handheld configurations. It operates at 3.3 V logic, so no level shifting is required, and typical current draw from the LED backlight is around 150–200 mA, which can be adjusted using PWM dimming. The 800×480 resolution offers plenty of space for user interface elements such as radar readings, sensor data, and system alerts. While it is not fully sunlight readable, visibility can be improved through anti-reflective coatings or a shaded enclosure. This display is priced around $13–$20 and offers a high-quality image, built in controller convenience, and straightforward integration with the existing SPI based design.

3.6.5.5. 4.3” 740-nit Sunlight Readable TFT

Since this project requires outdoor operations and high ambient light environments, the Crystalfontz 4.3” 800x480 TFT LCD is a standout option. This module is designed for readability with a backlight rated at approximately 740 nits (a unit of luminance), significantly brighter than the standard TFT modules. The display integrates a resistive touch panel and communicates via an SPI interface when paired with a compatible driver board or adapter. The high brightness output does increase power consumption a bit, roughly 250–350 mA, but allows excellent visibility in daylight, making it ideal for field deployments or emergency response applications. While the resistive touchscreen is less advanced than capacitive options, it performs reliably in conditions where gloves or moisture are present. The module’s physical construction is robust, with a thick glass cover and extended temperature tolerance. It costs around $30 to be competitive for a display of this brightness, making it a great candidate for outdoor or portable instrumentation interfaces.

3.6.5.6. 4.3” Capacitive Touch LCD

The Waveshare 4.3” LCD offers a crisp 800×480 IPS display with a smooth capacitive touch interface powered by the GT911 controller. While the touch controller itself communicates over I²C and is fully compatible with the ESP32-WROOM, the display portion uses an RGB/MIPI-DSI interface, which requires a dedicated display controller or a single-board computer such as a Raspberry Pi or ESP32-S3 for operation. The module is bright enough for semi-outdoor use, with a typical current draw of around 200 mA, and the capacitive touch system provides accurate multi-touch input. However, because the ESP32-WROOM lacks a native RGB or DSI interface, this screen is not directly compatible without additional hardware such as an RA8876 or SSD1963 controller board. The unit is priced around $30, making it an attractive option for systems that support high-speed display interfaces or for future hardware revisions that include an RGB-capable microcontroller.

Models	**ILI9341 TFT LCD**	**RA8875 Module**	**RA8875 Capacitive**	**Sunlight TFT**	**IPS with RA8875**	**Capacitive LCD**
Display Size	3.2”	4.3”	4.3”	4.3”	4.3”	4.3”

Models	ILI9341 TFT LCD	RA8875 Module	RA8875 Capacitive	Sunlight TFT	IPS with RA8875	Capacitive LCD
Driver IC / Controller	ILI9341 (onboard)	RA8875 (onboard)	RA8875 (onboard + I2C touch)	SSD1963 (on board)	RA8875 (onboard)	Waveshare integrated controller
Resolution	240x320	800x480	800x480	800x480	800x480	800x480
Interface	SPI	SPI	SPI +I2C	SPI	SPI	SPI
Voltage / Logic Compatibility	3.3 V (ESP32 direct)	3.3 V (ESP32 direct)	3.3 V (ESP32 direct)	3.3 V	3.3 V	3.3 V (ESP32 direct)
Power Consumption	80-120 mA (LED backlight)	150 – 250 mA typ.	160-250 mA typ.	250-350 mA (high brightness)	150 – 200 mA typ.	~200 mA typ.
Outdoor Readability	Limited	Moderate	Moderate	High	Moderate	Moderate
Ruggedization Options	Antiglare film / shade	Antireflective cover glass (optional)	Antiglare film / shade + touch overlay protection	Toughened glass / extended temp range	Matte overlay	Touch overlay
Use Case Fit	Compact, budget friendly UI prototype	Larger display for sensor dashboard and maps	Interactive UI with touch control menus	Outdoor or field use	Color rich display for data visualization	Modern touch GUI for portable devices

Table 18: LCD Part Comparison

3.6.5.7. LCD Selection Justification

The ESP32-WROOM module supports SPI and I2C communication natively, which makes displays using those interfaces the most suitable for this design. While RGB and MIPI-DSI displays such as the Waveshare 4.3″ capacitive LCD offer excellent visual quality, they require additional controller hardware or higher end microcontrollers like the ESP32-S3 to operate efficiently. Therefore, this project prioritizes SPI based modules such as the EastRising 4.3″ IPS RA8875 TFT, which provides a high-resolution 800×480 display, low latency SPI integration, and optional touchscreen functionality through I2C. These features make it directly compatible with the ESP32-WROOM, minimizing complexity while maintaining strong visual performance and moderate power consumption.

3.7. Processing and Control Modules

The processing and control subsystem oversees the operation, coordination, and data management of all sensing, communication, and user-interface components within the life-detection system. This module guarantees accurate capture, processing, and interpretation of real-time signals from radar, UWB, infrared, and optional seismic sensors. It also manages system feedback, power control, and external communication with other rescue or monitoring platforms. To identify the most suitable controller for this application, we meticulously examined and compared various architectures, including Microcontroller Units (MCUs), Microprocessors (MPUs), Systems-on-Chip (SoCs), co-processor hybrid designs, Field-Programmable Gate Arrays (FPGAs), and Single-Board Computers (SBCs). Each offers distinct advantages in terms of computational power, energy efficiency, and design intricacy.

3.7.1. Microcontroller Unit

A microcontroller unit (MCU) acts as the central processing unit in the proposed life-detection system, handling all timing, sensing, and communication tasks in real-time. Microcontrollers (MCUs) combine the processor, memory, and peripherals into a compact device, making them highly efficient for embedded control applications. In this project, the MCU orchestrates data collection from various sensors, including radar, UWB, IR, and optional seismic sensors. It performs lightweight signal processing and fusion, and interfaces with the LCD and user controls. Its low power consumption, compact size, and real-time responsiveness make it an ideal choice for a portable, battery-operated search and rescue device. Among the various options considered, including the ESP32-S3, ESP32-WROOM-32, and Raspberry Pi Pico W, the ESP32 WROOM 32 stood out as the best choice. It has proven reliability, dual-core performance, integrated Wi-Fi/Bluetooth connectivity, and versatile peripheral support. It offers the most balanced combination of processing capability, communication range, and integration flexibility, aligning well with the system's performance and power objectives.

3.7.2. Microprocessor

A microprocessor (MPU) differs from an MCU in that it relies on external memory and typically runs a full operating system, such as Linux. MPUs like the Raspberry Pi 4, BeagleBone Black, or NXP i.MX series offer significantly higher computational throughput and memory resources, which enable advanced processing tasks such as machine-learning inference, radar image reconstruction, and real-time mapping. However, MPUs demand more power, longer boot times, and larger board space, making them impractical for compact, field-deployable devices. MPUs are an excellent platform for research purposes, such as post-processing sensor data or simulating AI-driven detection algorithms. However, they are less suitable as the primary embedded controller in a portable rescue tool due to power and ruggedness constraints.

3.7.3. System-on-Chip

A System-on-Chip (SoC) seamlessly integrates the flexibility of a Microcontroller Unit (MCU) with the connectivity and compute power of a Multi-Processing Unit (MPU). This integration allows the SoC to combine multiple subsystems, such as wireless radios, memory, and AI acceleration, onto a single silicon die. Devices like the ESP32-S3, Nordic nRF54 series, or Qualcomm QCA4020 fall under this category. System-on-chips (SoCs) are ideal for edge computing applications as they enable onboard AI classification and sensor fusion without the need for external processors. Their high integration level reduces the overall number of components and simplifies PCB design, while also providing enhanced communication options like Wi-Fi, BLE, and Thread. SoCs offer a well-rounded solution for life detection, striking a balance between performance, power efficiency, and wireless versatility. This makes them an attractive research direction for future system iterations that prioritize real-time AI decision-making at the edge.

3.7.4. MCU and Co-Processor Architecture

In a hybrid architecture, a microcontroller unit (MCU) collaborates with a co-processor or secondary core to delegate specific tasks, such as buffering sensor data, encrypting information, or managing power consumption. Examples include microcontrollers with ultra-low-power (ULP) cores, such as the ESP32's internal ULP or the STM32U5 paired with a sensor hub. This distributed workload design enables the main microcontroller to remain in low-power mode while the co-processor manages background monitoring. This is particularly advantageous in long-duration rescue operations where battery life is of utmost importance. Co-processor setups enhance real-time performance by parallelizing tasks, such as managing continuous radar sampling while ensuring responsive user interaction. Although more intricate to program, this architecture improves system efficiency and fault tolerance, making it an advanced yet practical upgrade path for future prototypes.

3.7.5. Field-Programmable Gate Array

A Field-Programmable Gate Array (FPGA) offers unparalleled flexibility through hardware-level reconfiguration, enabling the implementation of custom signal-processing pipelines for radar or seismic data. FPGAs are highly adept at executing multiple tasks simultaneously, such as real-time Doppler filtering, Fast Fourier Transform computation, or multi-sensor synchronization. Devices like the Xilinx Artix-7 or Intel Cyclone 10 can significantly enhance latency and processing throughput compared to software-based microcontroller units. However, they necessitate specialized design tools and hardware description languages, which come with increased power consumption and higher costs. FPGAs are best suited for research and algorithm validation, especially in testing advanced signal-processing or AI fusion models, rather than the initial portable prototype.

3.7.6. Single-Board Computer

A Single-Board Computer (SBC), like the Raspberry Pi 4/5, NVIDIA Jetson Nano, or BeagleBone AI, serves as a comprehensive computing platform capable of executing intricate operating systems and sophisticated software environments. Supercomputers (SBCs) are valuable in research for data visualization, model training, and post-processing of field data. They provide powerful CPUs and

GPUs that can support graphical user interfaces or cloud communication. While they provide unparalleled flexibility for development, they are larger, power-hungry, and less robust than embedded MCUs, which restricts their suitability for handheld or battery-limited devices. In this project, an SBC can function as a ground station or analysis companion, handling AI processing and remote monitoring tasks while the MCU performs low-level sensing and control.

3.7.7. Processing and Control Module Selection

A MCU was selected as the primary processing component for this system because it offers the ideal combination of performance, efficiency, and integration for an embedded life-detection device. MCUs, unlike microprocessors or single-board computers, which necessitate external memory and higher power consumption, integrate the processor, memory, and essential peripherals into a single, compact chip. This design makes them ideal for portable, battery-powered systems. Its real-time processing capability allows it to manage multiple sensor inputs simultaneously, perform lightweight signal processing, and control user interface elements without latency. MCUs provide exceptional reliability, rapid startup times, and predictable performance, all of which are crucial for a search and rescue tool operating in unpredictable environments. This architecture guarantees efficient resource utilization, simplifies PCB design, and ensures long-term stability. Consequently, it makes an MCU the most practical and mission-ready choice for the core control system.

3.7.7.1. ESP 32

The ESP32 platform, specifically the ESP32-WROOM-32 variant, is a dual-core 32-bit Xtensa LX6 processor running up to 240 MHz with integrated Wi-Fi (802.11 b/g/n) and Bluetooth v4.2 connectivity. It provides up to 30 multiplexed GPIO pins and supports UART, SPI, I²C, PWM, ADC, and DAC interfaces, which makes it ideal for managing multiple sensors simultaneously. Its onboard wireless capability eliminates the need for external transceivers, allowing the MCU to handle both data acquisition and wireless communication within a single compact module.

For this project, the ESP32 offers sufficient processing bandwidth to manage radar data streaming from the TI IWR6843AOP module over UART while concurrently reading temperature and CO_2 data via I²C and coordinating inertial measurements from the M5Stack M135 INS. Under the FreeRTOS environment, one core can be dedicated to radar data reception and pre-processing, while the other manages wireless transmission, display updates, and control logic. This dual-threaded structure ensures deterministic timing with minimal latency, which is critical for continuous vital-sign detection and user feedback. With a nominal 3.3 V supply and typical active current of 150–250 mA, it comfortably fits within the handheld system's 10 W power budget and supports battery operation for extended missions. The extensive Espressif SDK ecosystem (ESP-IDF and Arduino frameworks) also minimizes firmware-development time and provides stable Wi-Fi and BLE libraries suitable for field deployment.

3.7.7.2. STM32

The STM32 architecture, produced by STMicroelectronics, encompasses a broad family of ARM Cortex-M-based microcontrollers ranging from low-power M0 devices to high-performance M7 variants operating at up to 400 MHz. These MCUs are renowned for precise timing, extensive analog capabilities, and industrial reliability. They are widely used in control systems and high-determinism applications due to their predictable interrupt handling and hardware-level timers.

In the context of this life-detection system, the STM32 could provide accurate timing for radar triggering or synchronized sensor sampling. However, STM32 devices lack built-in wireless connectivity; Wi-Fi or Bluetooth modules must be added externally, increasing board complexity, wiring density, and power overhead. Integrating an external wireless chipset would also require additional firmware development for network communication. While the STM32's analog precision and hardware timers outperform the ESP32 for deterministic control, these advantages are less critical in this design because the radar module already incorporates its own signal-processing DSP and timing control. Additionally, STM32-based development typically demands more extensive firmware infrastructure, requiring embedded C and hardware abstraction layers, which slows prototyping relative to the streamlined ESP-IDF toolchain.

3.7.7.3. RasberryPi

The Raspberry Pi platform, such as the Pi Zero 2 W or Pi 4 Model B, offers general-purpose processing capability, featuring multi-core ARM Cortex-A53/A72 CPUs running between 1.2 and 1.8 GHz, gigabytes of RAM, and a full Linux operating system. These single-board computers integrate Wi-Fi and Bluetooth connectivity by default and support high level programming environments. While the processing power would easily accommodate radar FFT computation, data fusion, and even visualization, the trade-off is significant power consumption (typically 500 mA – 2 A) and a larger physical footprint.

For a handheld, battery-powered device designed for field search and rescue operations, the Raspberry Pi's size, boot time, and power draw make it impractical. Ruggedized embedded systems require regulated 5 V input, active cooling in certain cases, and a secure file system shutdown to prevent data corruption. These features are not compatible with the instantaneous and reliable operation of ruggedized embedded systems under field stress. While the Pi could serve as a development testbed for advanced radar-signal visualization or machine-learning experiments, it exceeds the computational needs of this prototype and compromises energy efficiency and portability.

Platform	Processor	Wireless	Power (Active)	Key Advantages	Limitations
ESP32	Dual-core Xtensa LX6 @ 240 MHz	Wi-Fi b/g/n + BT 4.2	150–250 mA @ 3.3V	Integrated wireless, low cost, compact, proven SDK support	Moderate DSP power, limited ADC precision

Platform	Processor	Wireless	Power (Active)	Key Advantages	Limitations
STM32	ARM Cortex-M4/M7 up to 400 MHz	None (external module required)	80–150 mA @ 3.3V	Precise timing, strong analog control, industrial reliability	No built-in wireless, more complex development
Raspberry Pi	Quad-core ARM Cortex-A53/A72 @ 1.2–1.8 GHz	Wi-Fi + BT onboard	500 mA–2 A @ 5V	High performance, full OS, easy ML/visualization	High power draw, large size, slow boot

Table 19: MCU comparison

3.7.7.4. MCU Selection Justification

After a comparative evaluation, the ESP32 was selected as the primary microcontroller architecture for the handheld system. It delivers the most balanced combination of performance, wireless integration, cost, and power efficiency. The module's dual core structure, native Wi-Fi/Bluetooth support, and extensive peripheral compatibility make it uniquely suited to act as the bridge between the radar sensing front end and the user interface or network layer. While the STM32 provides excellent precision and reliability, it introduces unnecessary complexity and lacks built-in communication capability. The Raspberry Pi, though far more powerful, would drastically reduce battery life and increase hardware overhead beyond the system's design envelope. The ESP32 therefore stands as the most practical, energy efficient, and integration-ready option for this stage of development, enabling the project to meet real-time performance goals without exceeding cost or power constraints.

3.8. MCU Part Comparison

3.8.1. ESP32-WROOM-32

The ESP32-WROOM-32 is a compact, dual-core wireless microcontroller module developed by Espressif Systems and widely used for IoT and embedded sensing applications. It integrates a 32-bit Xtensa LX6 CPU running up to 240 MHz, a ULP RISC co-processor for low-power background operation, 520 KB SRAM, and 4 MB SPI flash. The module supports 2.4 GHz Wi-Fi (802.11 b/g/n) and Bluetooth v4.2 (Classic + BLE), enabling short- and medium-range communication for field devices. It operates within a nominal voltage range of 3.0 V to 3.6 V, with a typical load current of 150–250 mA. In deep-sleep mode, it drops to microamp levels, making it suitable for portable, battery-powered systems.

In this project, the ESP32-WROOM-32 serves as the central processing and communication hub, connecting the radar, infrared, CO_2, and inertial sensors. Its abundant peripheral set, up to 30 GPIO pins multiplexed for UART, SPI, I2C, PWM, ADC, and DAC, provides sufficient flexibility to connect all sensing modules simultaneously without the need for additional expanders. The 240 MHz dual-core processor provides enough computational headroom for moderate signal-processing tasks such as filtering radar outputs, decoding sensor data, or running control logic under FreeRTOS. Wi-Fi and BLE support allow wireless data transfer between rescuers' devices or to a base station, aligning with the project's goal of producing a networked, field-deployable life-detection tool.

The ESP32-WROOM-32 combines high processing power, built-in wireless communication, and low cost in a single, well-supported package. Its dual-core architecture allows one core to manage sensor acquisition and the other to handle communications or display updates, improving real-time responsiveness. The integrated security features, such as, AES, RSA, ECC encryption, and secure boot, make it suitable for scenarios requiring data integrity and secure firmware updates. The development ecosystem is mature, with support from Espressif IDF, Arduino IDE, and PlatformIO, which reduces firmware-development time. Its power efficiency and small footprint also support handheld or wearable implementation without external computer hardware.

Despite its versatility, the ESP32-WROOM-32 is not optimized for heavy digital signal processing. Operations such as FFT-based range-Doppler analysis or high-rate radar filtering can saturate its cores if implemented entirely in software, potentially requiring off-loading to the radar's onboard DSP or external processors. The ADC performance is limited to 12 bits and can suffer from noise, restricting precision analog measurements. Additionally, the module supports only 2.4 GHz Wi-Fi, which can be congested in some environments, and its wireless range may be limited in debris-dense or metallic areas typical of search-and-rescue zones. Careful power-management design is also necessary to ensure long runtime from portable batteries, as active-mode currents near 250 mA can drain smaller cells quickly.

Overall, the ESP32-WROOM-32 offers an excellent balance of processing capability, integration, and connectivity for the prototype stage of this search-and-rescue system. It is powerful enough to manage sensor coordination and wireless communication while remaining compact, inexpensive, and easy to program. However, future iterations aiming for advanced radar processing or extended field endurance may require supplementing it with a dedicated signal-processing coprocessor or higher efficiency power system.

3.8.2. ESP32-S3-WROOM-1

The ESP32-S3-WROOM-1 is an upgraded wireless MCU module from Espressif designed for applications requiring enhanced AI acceleration, higher memory capacity, and improved peripheral interfaces. It features a dual-core Xtensa LX7 processor that operates at up to 240 MHz, providing higher efficiency compared to the LX6 cores in the original ESP32-WROOM-32. Additionally, it incorporates vector (SIMD) instructions that enhance the processing of signals and neural networks. The module typically integrates 8 MB PSRAM and 4–16 MB flash, providing ample buffer space for radar and sensor data. Wireless support includes 2.4 GHz Wi-Fi (802.11 b/g/n) and Bluetooth 5 (BLE + Bluetooth Mesh), with longer range and higher throughput than previous generations. Operating voltage remains 3.0 V to 3.6 V, and active-mode power consumption is about

200–240 mA, like the WROOM-32, while sleep currents remain in the microamp range.

For this project, the ESP32-S3-WROOM-1 would serve as a strong alternative to the WROOM-32, especially if on-device radar processing or machine-learning-based signal classification (such as identifying respiration patterns) becomes a design priority. Its vector instructions and expanded PSRAM make it capable of handling heavier FFT computations, filtering, or data-fusion tasks locally rather than relying entirely on the radar's internal DSP. The larger flash and RAM are particularly useful for implementing image- or heatmap-based output from the IR sensor or logging multi-sensor data for post-event analysis.

Compared with the ESP32-WROOM-32, the S3 offers noticeably better computational throughput for signal-processing tasks and more memory for buffering or ML inference. Its Bluetooth 5 and BLE Mesh support extend wireless range and improve device-to-device connectivity, valuable for multi-rescuer coordination. The S3 also introduces USB OTG, eight capacitive touch inputs, and an improved ADC (13-bit effective), expanding the device's interface options for custom controls or diagnostics. Additionally, its enhanced security features, including a digital signature, HMAC, and secure boot v2, further fortify the device against firmware tampering.

The S3 is slightly more expensive and less power-efficient than simpler ESP32 modules when running AI or high-throughput tasks continuously. Its 2.4 GHz-only radio still limits range and penetration through rubble compared with sub-GHz systems, and the expanded memory footprint increases wake-up current slightly in low-power modes. While Espressif's SDK fully supports the S3, libraries for some peripherals (e.g., USB or camera interfaces) are larger and can increase firmware complexity. Overall, the ESP32-S3-WROOM-1 represents a balanced upgrade for this system, offering superior processing headroom for future radar-signal enhancement or on-board ML inference while maintaining pin-level compatibility and similar power requirements to the original ESP32. For a prototype that may evolve into a more autonomous, AI-aided platform, the S3 is the more forward-looking choice.

3.8.3. ESP32-C6-WROOM-1

The ESP32-C6-WROOM-1 is Espressif's latest generation single-core RISC-V wireless MCU module, emphasizing efficiency, advanced connectivity, and extended-range wireless features rather than raw processing power. It runs a 160 MHz RISC-V core and includes a ULP coprocessor for sensor polling during deep sleep. Unlike its predecessors, this model supports dual-band connectivity, including 2.4 GHz Wi-Fi 6 (802.11 ax) and Bluetooth 5 LE (with 802.15.4 for Thread and Zigbee networks). This makes it one of the most versatile MCUs currently available. The module combines 4 MB of flash memory and 400 KB of SRAM, operates at a voltage range of 3.0 to 3.6 V, and achieves significantly reduced active power consumption (typically between 80 and 120 mA) with microamp-level deep-sleep current. This makes it an ideal choice for systems that require extended battery life.

In the context of this project, the ESP32-C6-WROOM-1 offers strong advantages where low power and robust network communication take priority over heavy local computation. Its Wi-Fi 6 radio supports Target Wake Time (TWT) and improved spectral efficiency, extending battery life and maintaining stable connectivity in noisy, crowded environments such as urban disaster zones. The inclusion of Thread and Zigbee technology allows for mesh networking between multiple handheld devices, which is an attractive feature for coordinating search teams without the need for external access points.

The C6 delivers next-generation wireless performance and exceptional power efficiency, potentially doubling battery runtime compared to the WROOM-32. Its Wi-Fi 6 and Thread/Zigbee support introduce robust mesh networking options for device-to-device communication, and the RISC-V architecture benefits from modern toolchain support and smaller firmware size. The reduced current draw and advanced sleep scheduling align well with long-duration field missions where recharging is limited.

The C6 sacrifices raw computational capability, its single 160 MHz core is considerably slower for heavy DSP or FFT workloads than the dual-core LX6/LX7 variants. It also lacks the large PSRAM expansion found on the S3, limiting its ability to buffer high-rate radar data or run complex fusion algorithms. As a newer platform, some peripheral drivers (especially for mixed SPI/I²C high-speed sensors) are less mature than the S3's ecosystem.

Overall, the ESP32-C6-WROOM-1 is an excellent candidate for future revisions of the project where extended battery life, mesh networking, and secure wireless operation are prioritized over intensive on-board processing. For the current prototype—where real-time radar data handling and local computation remain key, the C6 would likely serve better as a low-power communication node rather than the primary processing MCU.

Module	**CPU**	**Wireless**	**Power (Active)**	**Key Advantages**	**Limitations**
ESP32-WROOM-32	Dual-core LX6 @ 240 MHz	Wi-Fi b/g/n + BT 4.2	150–250 mA	Proven, low-cost, strong I/O, easy integration	Older wireless standard, limited DSP
ESP32-S3-WROOM-1	Dual-core LX7 @ 240 MHz	Wi-Fi b/g/n + BT 5	200–240 mA	Higher performance, extra RAM, AI/DSP support	Slightly higher cost and power
ESP32-C6-WROOM-1	Single-core RISC-V @ 160 MHz	Wi-Fi 6 + BT 5 + Zigbee	80–120 mA	Very low power, mesh networking	Lower processing power, limited memory

Table 20: MCU model comparison table

3.8.4. MCU Final Selection Justification

After evaluating the ESP32-WROOM-32, ESP32-S3-WROOM-1, and ESP32-C6-WROOM-1, the team chose the ESP32-WROOM-32 as the primary microcontroller for this project. This decision was based on its optimal combination of processing power, power efficiency, hardware maturity, and ease of integration. While newer variants like the S3 and C6 offer incremental improvements in AI acceleration and wireless range, the WROOM-32 provides all the essential functionality required for the current system at a fraction of the complexity and cost. Its dual-core Xtensa LX6 processor running at 240 MHz delivers sufficient computational power to handle real-time data acquisition and preprocessing from the radar, infrared, CO_2, and inertial sensors without the need for external coprocessors. The module's extensive peripheral set, comprising three UARTs, multiple SPI and I2C buses, and up to 30 GPIO pins, facilitates the simultaneous operation of all sensing subsystems. This enables seamless communication with the radar module over UART while maintaining stable I2C and SPI channels for the other peripherals.

In the context of this life-detection and localization system, the ESP32-WROOM-32 offers proven reliability and an extensive development ecosystem, both critical for rapid prototyping and field deployment. Its integrated Wi-Fi and Bluetooth connectivity allow the device to transmit sensor data wirelessly to mobile devices or command centers, fulfilling the project's requirement for remote monitoring in search-and-rescue environments. The availability of FreeRTOS under the Espressif IDF framework ensures efficient multitasking, allowing one CPU core to manage radar data streaming while the other handles wireless communication, sensor fusion, or display updates. Power consumption remains moderate, typically between 150 mA and 250 mA during active sensing, making it feasible for battery powered operation while maintaining the responsiveness required for real-time detection.

Although the ESP32-S3 and C6 offer newer architectures, their advantages are not critical to this prototype's success. The S3's additional AI acceleration and memory are underutilized in a system where radar signal processing is offloaded to the IWR6843AOP's onboard DSP, and the C6's single core design sacrifices too much processing capability for marginal gains in network range and efficiency. The WROOM-32, by contrast, provides a stable, well-documented, and field-tested platform that integrates cleanly with the chosen radar and sensor modules. It satisfies all electrical, communication, and performance requirements for the prototype phase, ensuring dependable operation while keeping hardware costs and development time low. For these reasons, the ESP32-WROOM-32 is the most practical and technically appropriate microcontroller for the current stage of the project.

3.9. Power Supply

Proper power management is critical to ensure stable and efficient operation of the system. Each subsystem, the ESP32-WROOM microcontroller, the 4.3″ ILI9341-based TFT LCD display, and the Infineon Radar Baseboard MCU7, requires precise voltage regulation and power budgeting to maintain consistent performance without overloading the supply. This section discusses the power requirements and distribution across the system, detailing how each component draws current under normal operating conditions and how low-power peripherals are incorporated to improve efficiency.

3.9.1. MCU7 Baseboard

The Infineon MCU7 radar baseboard serves as the main interface and power distribution platform for the IWR6843AOP radar module, providing regulated 3.3 V and 5 V outputs to the radar, ESP32-WROOM microcontroller, and other peripherals. According to Infineon's documentation, the MCU7 operates from a 5 V input and distributes power through onboard low-noise regulators to support digital logic, analog front-end circuits, and radar transceiver functions. The total active power consumption of the IWR6843AOP radar mounted on the MCU7 board is approximately 2.0 – 2.5 W, with an additional 0.3 – 0.5 W allocated to the ESP32 and communication interfaces. The system therefore requires a stable 5 V, 600 mA minimum supply, with additional current margin to account for startup surges and simultaneous radar and LCD activity. The MCU7 includes onboard voltage sequencing to ensure that the 3.3 V logic rail is enabled before the radar analog front-end, preventing latch-up or transient errors during power-up. These characteristics define the primary constraints for regulator selection and overall power budgeting in the system design.

3.9.2. LCD Display

The ER-TFTM043A1-7S TFT LCD serves as the primary user interface for the system, displaying real-time data, status indicators, and user prompts. Because the display is one of the more power-intensive components, it's important to evaluate its electrical and operating requirements carefully. Since the ESP32-WROOM uses a 3.3 V logic system and an Intel 8080 style parallel bus, the display will be ordered with the 8080 interface and 3.3 V power configuration. The backlight can be controlled by the onboard SSD1963 controller (J4 short, J3 open) or through an external signal if reversed. At 3.3 V operation, the module typically draws around 450 mA, though this can vary with brightness and refresh activity. The backlight consumes the most power, so it should be budgeted for worst case brightness and scaled down proportionally if dimming with PWM. Additional peripherals such as the touch controller add roughly 3-10 mA, and logic lines stay within 3.3 V compatibility for direct use with the ESP32-WROOM. When combined with the ESP32's own current demand (about 240-300 mA peak during Wi-Fi activity), a total 3.3 V rail capacity of around 1 A is recommended for safe headroom. This ensures stable operation during current spikes and allows flexibility for added peripherals or higher brightness levels.

3.9.3. ESP32-WROOM

The ESP32-WROOM is the main microcontroller and communication module of the system, responsible for processing sensor data, managing peripheral interfaces, and handling Wi-Fi/Bluetooth communication. According to the Espressif ESP32-WROOM-32 datasheet v3.5, the module operates at 3.3 V with current draw varying Wi-Fi mode and transmission power. According to the ESP32 Series Datasheet v5.1 [1], typical transmit currents are about 180 mA for 802.11n, 190 mA for 802.11g, and up to 240 mA for 802.11b at +19.5 dBm. Receive current averages around 95-100 mA. Although only one RF mode is active at a time, the highest typical draw of 240 mA (~ 0.79 W) is used for design to ensure supply stability during transmit peaks. A 25% margin is added, giving a recommended supply capacity of roughly 300 mA. In low-power modes, current decreases to 20-25 mA in modem sleep, 0.8 mA in light sleep, and 100-150 µA in deep sleep.

3.9.4. Inertial Navigation System

The M5Stack M135 INS module operates from a single 3.3 V rail and draws on the order of 0.11 W in continuous GNSS/IMU operation, corresponding to an average current of roughly 34 mA at 3.3 V (0.111 W/3.3 V). To ensure margin for GNSS acquisition bursts, sensor polling, and interface activity on the shared I²C/UART buses, the 3.3 V power budget allocates 50 mA for this module. This allocation aligns with the module's integrated sensor set (NEO-M9N GNSS, BMI270 IMU, BMM150 magnetometer, BMP280 barometer) and supports update rates up to 25 Hz without impacting the system's overall thermal or efficiency targets. For sizing and rail stability, the M135 should be tied to the primary 3.3 V regulator domain (buck-boost) with local decoupling (≥10 µF bulk + 0.1 µF close to the supply pin) and included in the always-on load group.

3.9.5. Ultra-Wide Band

The DWM3000EVB integrates Qorvo's DWM3000 UWB module (DW3000/DW3110 transceiver). While the module brief notes low power consumption, numeric current values are specified in the DW3000 IC datasheet, which provides TX/RX current profiles. Under active ranging, the DW3000 exhibits burst currents with RX peaks on the order of ~70-80 mA and TX peaks ~20-35 mA, depending on supply rails and mode; for power budgeting we allocate a worst-case transient of ~90 mA at 3.3 V (~0.30 W). In sleep modes, current falls to sub-milliamp/µA levels. The EVB exposes a J1 current measurement jumper to verify module draw in system, once the component has been delivered, the measured values from this will be used for final power budgeting.

3.9.6. Infrared Sensor

The MLX90640 is a 32×24-pixel infrared (IR) thermal imaging sensor manufactured by Melexis. It operates from a 3.3 V supply with a typical current consumption of 20 mA and a maximum of 25 mA, corresponding to an estimated power draw between 0.066 W and 0.082 W. Communication is handled through an I2C interface compatible with both 3.3 V and 5 V logic, making it suitable for use with common microcontrollers such as the ESP32. The sensor's low power requirement allows it to run efficiently in embedded applications while maintaining accurate non-contact temperature measurements across its 55° field of view.

3.9.7. Low Power Supporting Components

In addition to the primary subsystems, several low power supporting components are used throughout the design to provide efficient operation and system stability. These include onboard voltage regulators, indicator LEDs, logic level shifters, and sensor interface circuits that operate within the 3.3 V domain. The regulators ensure a consistent supply to all modules while maintaining low quiescent current, minimizing overall standby consumption. Indicator LEDs draw only a few milliamps each and are typically duty cycled or turned off during idle periods to reduce unnecessary load. Level shifting ICs and I2C pull ups operate in the microamp range during static conditions, contributing negligible power usage compared to major components such as the LCD and radar module. Collectively, these low power peripherals enable reliable monitoring and control without significantly impacting the system's total power budget, supporting energy efficient embedded operation.

Component	Voltage (V)	Current (mA)	Power (W)
ESP32-WROOM	3.3	80-300	0.26-0.99
LCD Display	3.3	~ 450	1.49
LCD Backlight	5	~ 280	1.4
Radar Baseboard MCU7	5	400 - 500	2.0 - 2.5
Infrared Sensor	3.3	20-25	0.066-0.082
Ultra-wide Band	3.3	~90	0.3
INS	3.3	34	0.11
Supporting Components	3.3	< 50	0.17

Table 21: Power Distribution Table

3.10. Battery Types

For this project, several rechargeable battery chemistries were evaluated to determine the most effective option for powering our device. The main selection criteria included energy density, weight, cost, safety, discharge capability, and ease of integration. The following subsections summarize each candidate chemistry and its suitability for a compact, portable radar-sensor platform.

3.10.1. Lithium-Ion (Li-ion)

Lithium-ion batteries are widely used in modern electronics due to their high energy density, low weight, and long service life. Their ability to store substantial energy in a small form factor makes them ideal for portable embedded systems where space is limited. They also offer fast charge cycles and relatively low self-discharge rates. However, Li-ion cells require a battery-management system (BMS) for safe charging and protection against over-voltage or short circuits. If damaged or improperly handled, Li-ion batteries may experience thermal runaway or fire, which introduces additional design and safety considerations.

3.10.2. Lithium-Polymer (LiPo)

Lithium-polymer batteries are a subtype of Li-ion chemistry that use a gelled polymer electrolyte instead of a liquid one. This allows them to be manufactured in thin, lightweight, and flexible enclosures, which is advantageous for embedded applications that demand compactness. LiPo packs can deliver high discharge currents suitable for powering both the ESP32 and peripherals such as the LCD backlight and radar baseboard. While they share the same safety challenges as Li-ion batteries, proper balancing and protection circuitry significantly reduce these risks. LiPo batteries were therefore considered a strong candidate due to their power-to-weight ratio and ease of sourcing in custom capacities.

3.10.3. Nickel-Cadmium (NiCd)

NiCd batteries are durable and capable of handling moderate discharge rates, but they have a low energy-to-weight ratio and suffer from a pronounced memory effect, reducing usable capacity over time. Their high self-discharge rate makes them unsuitable for systems that operate intermittently or remain idle between uses. In addition, cadmium is a toxic heavy metal that presents environmental disposal concerns, making this chemistry less desirable for modern embedded systems.

3.10.4. Nickel-Metal Hydride (NiMH)

NiMH cells are safer and more environmentally friendly than NiCd, but they still have lower energy density compared with lithium-based options. They also discharge more quickly when not in use and exhibit voltage drops under heavy loads. For this project, the combination of continuous wireless communication, sensor operation, and display output would cause NiMH packs to deplete rapidly, reducing system runtime and efficiency. Although they offer stability and low risk of thermal events, their performance limitations made them impractical for this design.

3.10.5. Sealed Lead-Acid (SLA)

SLA batteries are valued for their reliability, affordability, and ability to supply high surge currents, but they are considerably heavier than other types. Their low energy-to-weight ratio and bulky form factor make them unsuitable for lightweight portable systems. SLA chemistry remains common in stationary or large-scale applications but would significantly increase the overall mass of this handheld device. Despite their robust safety record and simple charging requirements, SLA batteries were excluded due to weight and volume constraints.

3.10.6. Battery Selection and Justification

Among the battery types evaluated, each presented distinct advantages and disadvantages when considered for our embedded radar-based sensing system. Lithium-ion and Lithium-Polymer (LiPo) batteries provide high energy density and lightweight performance, but come with elevated safety risks, including the potential for overheating and combustion under improper handling. Due to these concerns and faculty restrictions, LiPo batteries were excluded from consideration. Nickel-Cadmium (NiCd) and Nickel-Metal Hydride (NiMH) batteries, while durable and relatively stable, suffer from high self-discharge rates, limited energy density, and memory effects, making them unsuitable for long-duration portable applications. Sealed Lead-Acid (SLA) batteries remain a cost-effective and safe option, but their large mass and low energy-to-weight ratio make them impractical for this compact, handheld system.

After comparing each battery, the Lithium-Ion was selected as the optimal choice for this project. Li-ion cells offer a strong balance of energy density, efficiency, and safety when paired with an integrated battery management system (BMS). The BMS safeguards against overcharging, short circuits, and thermal events, addressing many of the common safety issues associated with lithium-based designs. In addition, Li-ion batteries are widely available, have long cycle life, and support high discharge currents, sufficient for powering the ESP32-WROOM microcontroller, TFT LCD display, Infineon radar module, and supporting sensors simultaneously.

Although slightly more expensive than NiMH or SLA options, the superior performance-to-weight ratio, portability, and compatibility with standard 3.7 V power regulation circuits make Lithium-Ion batteries the most appropriate power source for this design. A single-cell 3.7 V Li-ion pack rated at approximately 2500–3000 mAh provides sufficient energy to sustain the system’s estimated 3-4 W load for roughly 2–2.5 hours of continuous operation, depending on sensor activity and backlight brightness. This configuration ensures a reliable and efficient power solution while maintaining safety and compliance with laboratory standards.

3.11. Voltage Regulation

Voltage regulation is a critical aspect of the power management system, ensuring that all components, including the ESP32-WROOM microcontroller, LCD display, radar baseboard, and sensors, receive a stable and reliable supply voltage. Regulators maintain consistent voltage levels even as the battery voltage changes during discharge, preventing fluctuations that could cause instability or damage to sensitive electronics.

In this project, the system is powered by a single cell 3.7 V Lithium-Ion battery protected by a Battery Management System (BMS). The BMS safeguards the cell against overcharging, deep discharging, and short circuit conditions, while the voltage regulators provide precise voltage conversion for each subsystem. Two regulated outputs are required: 3.3 V for the ESP32, LCD, and low power sensors, and 5 V for the Infineon Radar Baseboard MCU7 and other peripherals requiring higher voltage.

The 3.3 V rail supplies logic and communication circuitry and must deliver up to approximately 1 A during peak operation. The 5 V rail supports the radar module and other peripherals, requiring around 1 A of available current capacity. Both rails will be derived from the 3.7 V battery using high-efficiency switching buck-boost converters, capable of maintaining stable output even as the battery voltage drops during operation.

To maximize efficiency and runtime, voltage regulators with efficiencies above 90% were selected. This minimizes energy that is lost to heat and extends the usable battery life for portable operation. In addition, regulators with low quiescent current were prioritized to reduce standby power draw when the system is idle. Proper thermal management and copper pour areas on the PCB will further help dissipate any residual heat and ensure consistent long-term performance.

The system is designed to be fully rechargeable through a 5 V USB-C connection, which powers a single-cell (1S) Li-ion charging circuit. This charger manages both constant current and constant voltage charging to safely replenish the battery while allowing the device to operate simultaneously. A Battery Management System (BMS) sits between the battery and load, providing continuous protection against over-charging, over-discharging, and short circuit conditions. The USB-C port is configured as a standard 5 V input using CC pull-downs, and the design includes ESD protection and inrush current control on the VBUS line. Together, this setup ensures safe, efficient charging, and reliable operation whether the system is plugged in or running on battery power.

Parameter	**LM34936**	**LM5176**	**LM33620**
Type	Buck-Boost	Buck - Boost	Buck (Integrated)
Input Voltage (V)	~4.5 - 36	~3.5- 42	~3.8 - 36
Output Voltage (V)	Adjustable	Adjustable	Adjustable
Max Output Current (A)	~3	~ 3 - 5	3
Efficiency	≥ 90%	≥ 90%	≥ 90%
Quiescent Current	25 µA	~ 2.6 µA (10 max)	~ 27µA
Thermal Performance	Low heat	Low heat (uses external FETS)	Low heat

Table 22: 3.3 V Regulator Comparison

Parameter	LM3478	LM5155	LM2731
Type	Boost (Controller)	Boost (Controller)	Boost (Integrated)
Input Voltage (V)	2.9 - 40	3 - 42	2.7 - 5.5
Output Voltage (V)	Fixed (5.0)	Adjustable	Adjustable
Max Output Current (A)	1 – 2	2 – 3	< 1
Efficiency	82 - 85%	≥ 90%	75 - 85%
Quiescent Current	100 μA	2.6 μA	16 μA
Thermal Performance	Moderate heat	Low heat	Moderate – high heat

Table 23: 5 V Regulator Comparison

3.11.1. Regulator Selection and Justification

3.11.1.1. 3.3 V Regulator

The LM5176 was selected as the primary 3.3 V regulator because it provides a stable, high-current buck-boost architecture capable of regulating 3.3 V across the full Li-ion discharge range (3.0 - 4.2 V). The device uses external MOSFETs and an inductor, allowing the design to support higher currents and better thermal behavior than integrated regulators. A major contributor to this selection was direct guidance from the professor. Previous senior design teams who attempted to use TPS series integrated converters (such as TPS63061 or TPS61236P) achieved roughly a 20% success rate, assumed to be largely due to thermal failures, startup instability, and difficulty meeting load demands during bring-up. The integrated switches inside TPS devices do limit design flexibility and leave little margin for layout imperfections or varying load conditions, both of which are common challenges in student-built prototypes. By contrast, the LM5176 is a controller-based solution, not an integrated converter. This means our group can choose MOSFETs, compensation networks, inductors, and current limits appropriate for the system's needs, significantly improving robustness and increasing the likelihood of first-revision success. The LM5176 also achieves extremely low quiescent current (2.6 μA typ) in shutdown mode, which allows the battery-powered system to maintain a low idle drain. In conclusion, the LM5176 provides superior stability, scalability, and success probability compared to TPS-series parts used in past projects. Its architecture directly addresses the issues that caused low success rates in earlier senior design groups.

3.11.1.2. 5 V Regulator

The LM5155 was chosen for the 5 V rail because the system requires a reliable boost converter capable of delivering steady 5.0 V from a 1-cell Li-ion battery. With the battery's maximum voltage at 4.2 V, a boost topology is mandatory, and the LM5155's controller-based design allows selection of MOSFETs and magnetics appropriate for the required 2 - 3 A output range. As mentioned in the previous section, this choice was also strongly influenced by the professor's guidance, as he stated that the TPS series have a low success rate. The LM5155 resolves the issues by giving full control over external MOSFETs, inductor sizing, compensation, and switching behavior. Its shutdown quiescent current is extremely low (2.6 µA typ), and efficiencies of ≥ 90% are achievable with proper synchronous rectification. This makes the LM5155 far more forgiving and robust for a battery-powered system built by our design group. In conclusion, the LM5155 offers significantly greater design margin, thermal reliability, and real-world stability than TPS devices historically used in senior design. Choosing LM-series controllers directly aligns with the professor's recommendation to avoid components that have repeatedly failed in past projects.

Header: Charging IC

Since the team has chosen to use a rechargeable li-ion battery to power the device, a charging IC must be utilized to charge the battery. A dedicated charger IC is required to safely manage the charging and power distribution of the single-cell lithium-ion battery while allowing the system to operate from USB power. Lithium-ion cells demand precise voltage and current control throughout their charge cycle to prevent overcharging, overheating, or cell damage, which cannot be achieved with a simple voltage regulator. The charger IC handles this automatically by providing constant-current/constant-voltage charging, integrated protection, and dynamic power-path management. This enables the device to run and charge simultaneously without brown-outs or overloading the USB source, ensuring safe, efficient, and reliable power delivery for the entire system.

Header: Charging IC Options

	TI BQ24074RGTR	**TI BQ24195**	**Microchip MCP73831**
Manufacturer	Texas Instruments	Texas Instruments	Microchip
Part Number	BQ24074RGTR	BQ24195	MCP73831
Topology	Linear	Switch-Mode (Buck)	Linear
Power-Path Support	Yes (Dynamic Power-Path Management)	Yes	No

Max Charge Current (A)	1.5	4.5	0.5 – 1.0
Input Voltage Range (V)	4.35 – 6.4	3.9 – 17	4.4 – 6
Efficiency	Up to 90%	Up to 95%	~70%
Quiescent Current (μA)	55	150	1
Control Interface	Resistor-programmed	I²C	Fixed (set by resistor)
Thermal Regulation	Yes	Yes	No
Key Features	Integrated DPPM, USB input limiting, simultaneous operation and charging	High-efficiency buck charger, adjustable input limits, system power priority	Simple linear charger, low cost, minimal external components
Typical Applications	Portable embedded systems, IoT devices, handheld sensors	High-power handhelds, smartphones, tablets	Basic single-cell Li-ion devices, USB-powered gadgets

Header: Charger IC Selection and Justification

The BQ24074RGTR was selected as the charger and power-path management IC because it met the system's safety, reliability, and simultaneous operation requirements better than any alternative. Its dynamic power-path management (DPPM) feature ensures that the handheld unit can operate from USB-C power while charging the Li-ion cell, automatically adjusting the charge current whenever system demand increases or the adapter's capability is limited. This behavior prevents voltage droop or brown-outs during peak radar and LCD activity. The device's constant-current/constant-voltage charging profile, thermal regulation, and input current limiting provides full protection against overcharge, overheating, and inrush events without requiring firmware intervention. Because it is a linear charger, the layout and bring-up process is simpler than switch-mode designs, yet its 1.5 A programmable charge current is sufficient for the 3.7 V, 2500–3000 mAh Li-ion pack selected by the team. The BQ24074's USB compliance, robust fault protection, and ability to prioritize system load over charge rate made it the most dependable and practical choice for ensuring stable power delivery in the final design.

3.12. Additional Required Power Hardware

Beyond the battery and regulators, the rechargeable power system we will be using in our project requires a USB-C recpetable configured as a 5V sink with pull down resistors, two 5.1kΩ CC, a single cell Li-ion charger IC and a 1S BMS protection board. To prevent backfeed, we will also need to include ESD/TVS on VBUS, an input filter and a small load switch or ideal diode device. Each DC-DC requires it is specific inductors and input/output capacitors per their respective datasheets as well as a short feedback network. Lastly, we will need a battery holder and connector. During the prototyping phase, the system can be safely powered using a variable DC power supply instead of a battery. This allows precise control of both voltage and current while testing individual subsystems such as the ESP32, radar baseboard, and LCD display. The supply can be set to 4.0–4.2 V to simulate a fully charged single-cell Li-ion battery, or 5.0 V to represent the USB-C input voltage. A current limit of 1–1.5 A is recommended to protect components during early development and troubleshooting. The power supply connects directly to the system rails through the breadboard or test points, allowing easy measurement of voltage and current draw. This setup provides a stable and adjustable source for verifying circuit performance and regulator functionality before integrating the Li-ion battery, charger, and BMS.

3.13. Communication Protocols

The 2HAD-RAD system utilizes a well-structured embedded firmware framework, developed in C/C++ using one of the environments. This framework enables real-time coordination of multiple sensing modules, communication protocols, and user interface components operating concurrently under the FreeRTOS task scheduler. Each module, from the radar and UWB sensors to the thermal array and inertial navigation system, necessitates distinct initialization, data acquisition, and filtering procedures. The firmware manages data flow across UART, SPI, and I2C buses, ensuring deterministic timing and low power consumption. The system architecture employs a modular approach, with independent tasks for sensor handling, data fusion, and display updates, facilitating simplified debugging and future expansion.

3.13.1. Radar Module (TI IWR6843AOP)

The radar subsystem utilizes the Texas Instruments IWR6843AOP millimeter-wave radar, which outputs pre-processed range Doppler and presence detection data through a configurable communication interface. In its default evaluation module (EVM) configuration, the radar transmits processed data via UART, which the ESP32-WROOM-32 can readily receive using one of its multiple hardware serial ports. The firmware initializes the UART at the specified baud rate (usually 921,600 bps for radar streaming) and implements a packet-parsing algorithm to decode the binary frame structure defined by TI's mmWave SDK. This interface simplifies integration during prototyping since UART requires only transmit, receive, and ground connections, making it ideal for early-stage validation and debugging using serial monitors or terminal emulators. Each data frame received is parsed for signal amplitude, range bin, Doppler shift, and phase variance, allowing the MCU to identify both large scale motion and micro-movements such as respiration. Temporal filtering techniques, such as moving average and low pass smoothing filters, are applied to suppress noise and minimize false detections caused by environmental reflections or electromagnetic interference.

While UART provides a simple, stable interface, it is limited in bandwidth for higher radar data rates or real-time streaming of multiple frame types. For this reason, future iterations of the system might transition radar communication to Serial Peripheral Interface (SPI), which is also supported by both the IWR6843AOP and the ESP32-WROOM-32. SPI provides significantly higher throughput, up to several megabits per second, enabling faster data transfers and more efficient synchronization between radar sampling and other sensor acquisitions. SPI's full-duplex capability would enable the MCU to simultaneously transmit configuration commands and receive radar data, thereby reducing latency and freeing up processing cycles for data fusion. Integrating the radar through SPI would further align it with the high-speed communication already used by the LCD display and potential UWB module, streamlining timing control and improving real-time performance under FreeRTOS task scheduling.
Radar configuration parameters, including chirp rate, frequency slope, frame duration, transmit power, and detection thresholds, are managed through initialization commands sent from the MCU to the radar over the chosen interface. The firmware schedules radar polling and data parsing within a dedicated FreeRTOS task, maintaining consistent frame timing independent of other sensor activities. The resulting radar output stream, which has been filtered and time-stamped, is then transmitted to the data fusion layer. Here, it undergoes cross-correlation with thermal and UWB signal responses to confirm the presence of a human being behind walls or rubble.

3.13.2. UWB Signal Processing (DWM3000)

The Decawave DWM3000 Ultra-Wideband (UWB) module is integrated not for localization or ranging applications but as an RF sensing extension that enhances radar detection performance through dense or partially obstructed materials. The module communicates with the ESP32-WROOM-32 microcontroller using either the Serial Peripheral Interface (SPI) or Universal Asynchronous Receiver/Transmitter (UART) protocols, both of which are natively supported by the ESP32's flexible peripheral set. During early development, UART communication may be utilized for rapid prototyping due to its straightforward two-wire interface and simplified debugging via standard serial tools. UART allows the team to verify module functionality, transmit configuration commands, and view diagnostic data without extensive firmware overhead. However, for final integration and synchronized operation with the radar subsystem, SPI communication is the preferred interface. SPI offers substantially higher bandwidth and deterministic timing, allowing the microcontroller to retrieve channel impulse response and received signal strength indicator data at a much faster rate. This higher throughput enables near real-time processing by synchronizing UWB sampling intervals with radar frame timing, thereby enhancing data fusion and temporal correlation.

The firmware routines for the UWB module handle initialization, channel configuration, and periodic pulse transmission and reception. Once activated, the ESP32’s SPI bus retrieves raw CIR data representing the energy of reflected UWB pulses. This data is parsed to extract the temporal decay and intensity profile of signal reflections, which can indicate the presence of human-like motion or breathing patterns even through walls or rubble. Unlike typical UWB implementations that compute distance through time-of-flight (ToF) measurements, this firmware focuses on amplitude and energy stability analysis, allowing it to operate as a passive motion detector rather than an active localization tool. Filtering algorithms, including band-pass filters centered around motion-related frequencies (typically below 10 Hz) and envelope detection techniques, are employed to isolate biological or micro-motion signatures from environmental noise. To reduce interference from nearby RF sources, an adaptive notch filter may also be implemented to suppress narrowband noise peaks within the UWB spectrum.

After filtering, the processed UWB signal data is normalized and time-stamped before being synchronized with radar outputs within the data fusion framework. This synchronization guarantees that both radar and UWB sensing modalities work together seamlessly, enhancing the system’s reliability in verifying human presence through material barriers. By correlating energy fluctuations across UWB and radar data streams, the firmware can distinguish genuine human motion from static or environmental artifacts, significantly enhancing detection confidence. Future firmware updates may implement dynamic SPI clock scaling and DMA buffering to further optimize data throughput and minimize latency, enabling high-speed sensing without overloading the ESP32’s processing cores.

3.13.3. Infrared Thermal Sensor (MLX90640)

The Melexis MLX90640 infrared thermal array interfaces with the ESP32-WROOM-32 microcontroller via the I2C protocol, transmitting temperature data as a 32x24-pixel matrix. The I2C bus was chosen for this sensor because of its simplicity, low pin count, and compatibility with the ESP32’s hardware I2C controllers. These controllers support clock speeds up to 1 MHz (Fast Mode

Plus), which is more than enough to handle the MLX90640's maximum frame rate of approximately 16 Hz. The firmware configures the I2C interface with proper pull-up resistors and address mapping to ensure stable communication, particularly given the sensor's relatively high data payload per frame. During initialization, the microcontroller retrieves factory calibration constants and emissivity parameters stored in the sensor's onboard EEPROM, then applies Melexis's proprietary compensation algorithms to convert 16-bit raw pixel readings into precise temperature values in degrees Celsius. This compensation corrects for sensor offset drift, gain variation, and thermal gradients across the sensor die, which is critical for maintaining reliable human-heat detection accuracy in varying ambient conditions.

To maintain data integrity and mitigate frame-to-frame noise, the firmware implements a temporal averaging filter, which blends consecutive frames using a weighted moving average to stabilize pixel readings while preserving response time. A two-dimensional Gaussian smoothing kernel is then applied spatially across the thermal image. This process minimizes pixel level noise and sensor jitter while preserving sharp transitions at the boundaries of heat sources. These sharp transitions are crucial for distinguishing human-shaped signatures from background warmth. Each processed frame is analyzed for temperature gradients and intensity clustering to identify heat zones indicative of a human presence. A maximum intensity detection algorithm scans the filtered frame for temperature anomalies that exceed a dynamic threshold above the ambient baseline, adapting automatically to environmental conditions.

The ESP32 firmware manages this process through a dedicated I2C polling task running within the FreeRTOS environment, ensuring the sensor is sampled at consistent intervals synchronized with the radar subsystem's frame rate. This synchronization allows direct temporal alignment between radar detections of motion or breathing and corresponding thermal readings of body heat, significantly increasing the reliability of life detection outcomes. To prevent data collisions on the I2C bus, especially when other sensors such as the M5Stack M135 INS module share the same bus, the firmware schedules sequential read operations with prioritized timing slots. Error handling routines detect communication interruptions and automatically reinitialize the I2C connection if a bus hang or NACK condition occurs, preserving long term operational stability. Finally, the processed thermal data, which includes the centroid coordinates of the warmest regions, regional temperature averages, and maximum temperature gradients, is normalized and transferred to the data fusion subsystem. There, it is time stamped and cross correlated with radar and UWB sensor outputs. This multi-sensor integration allows the firmware to confirm the presence of a heat-emitting body in conjunction with motion or micro-vibration data, dramatically reducing the probability of false positives. Future firmware enhancements may incorporate adaptive temporal filters that adjust averaging strength based on motion activity or ambient temperature changes, improving response precision in dynamic search-and-rescue conditions.

3.13.4. Inertial Navigation System (M5Stack M135)

The M5Stack M135 INS module provides essential orientation, motion, and environmental data to stabilize and contextualize sensor readings across the 2HAD-RAD platform. The module communicates with the ESP32-WROOM-32 microcontroller via the I2C protocol, sharing the same bus as the MLX90640 infrared array. The I2C interface was chosen for its efficiency and ability to multiplex multiple sensors on a common two-wire connection, minimizing the number of pins used while maintaining sufficient communication speed. The M135 integrates multiple sensing elements, including the Bosch BMI270 six-axis IMU, BMM150 three-axis magnetometer, and BMP280 barometric pressure sensor, which collectively provide acceleration, angular velocity, magnetic heading, and altitude data. Each component within the module has its own I2C address, allowing the ESP32 to address and poll data sequentially within a dedicated FreeRTOS managed I2C task. To ensure reliability on a shared bus, the firmware employs staggered read scheduling and bus arbitration control, preventing timing conflicts between the INS and the thermal sensor.

During initialization, the firmware configures the BMI270 IMU's output data rate, dynamic range, and digital filtering modes, while also calibrating the magnetometer's offset and scale factors. Continuous I2C polling retrieves accelerometer and gyroscope data, which are inherently noisy due to mechanical vibration and user handling. To mitigate this, the firmware implements digital low-pass filters to attenuate high frequency noise and bias compensation routines to correct sensor drift over time. A complementary filter then fuses accelerometer and gyroscope outputs to produce smooth and stable orientation estimates by weighting the gyroscope's short-term accuracy against the accelerometer's long-term stability. This provides a precise estimate of the device's roll and pitch angles even during motion. The magnetometer provides yaw (heading) correction, ensuring that orientation data remains stable despite drift in the gyroscope's integration over time. For enhanced accuracy and robustness, the firmware architecture allows the future integration of a Kalman filter, which can perform dynamic state estimation using covariance weighting, particularly beneficial for scenarios involving significant motion or vibration.

The BMP280 barometric pressure sensor measures ambient pressure and computes relative altitude, giving the system context for vertical positioning and environmental changes. The firmware periodically polls this data and averages multiple readings to suppress pressure fluctuations caused by transient air movement. Together, these sensors enable the MCU to characterize how the handheld device is being moved or tilted during operation, providing the reference frame needed to interpret radar and thermal sensor readings accurately. For instance, the INS data can identify whether a shift in radar return is due to external motion or internal device movement caused by the operator. The processed INS data, which includes filtered linear acceleration, angular velocity, orientation quaternions, and altitude, undergoes normalization and time stamping before being transmitted to the data fusion layer. There, it serves as a stability reference, allowing the radar and thermal algorithms to compensate for user-induced motion or vibration. This correction is critical in search and rescue environments where the operator may be scanning uneven terrain or moving rapidly between positions. INS integration thus enhances detection confidence and minimizes false positives caused by motion artifacts. Future firmware enhancements may include real-time motion compensation, where radar or UWB readings are automatically adjusted for device tilt, and gesture-based controls, allowing operators to trigger mode changes through motion patterns rather than physical buttons.

3.13.5. User Interface and Feedback

The UI firmware serves as the operator’s primary point of interaction with the 2HAD-RAD system, enabling visual, tactile, and auditory feedback under varying environmental conditions. The user interface comprises a 4.3-inch EastRising RA8875 TFT LCD display, three OSRAM R976 indicator LEDs (red, yellow, and orange), three Apem IPR1SAD2 pushbuttons, a CMI-1295-03TH piezoelectric buzzer, and a PUI AS01508AO-WR-R capsule speaker powered by a PAM8302A audio amplifier. Together, these components ensure clear and redundant signaling in low-visibility or high-noise rescue environments. The firmware running on the ESP32-WROOM-32 coordinates all interface elements using a combination of SPI, GPIO, and PWM communication channels, allowing independent yet synchronized control across subsystems.

The RA8875-based LCD module communicates with the MCU through the SPI bus, selected for its high-speed throughput and efficient full-duplex data transfer. SPI allows rapid updating of graphical elements such as sensor readings, alert indicators, and mode selections without burdening the CPU. The firmware initializes the LCD controller with appropriate clock polarity, phase configuration, and frame buffer settings, then handles graphical rendering through pre-defined routines optimized for FreeRTOS task scheduling. A dedicated LCD task manages periodic refresh cycles and prioritizes critical display updates, such as life-detection confirmations or power warnings, while maintaining consistent frame rates. To minimize latency, direct memory access transfers are enabled where possible, offloading data movement from the main processor.

The indicator LEDs are connected to dedicated GPIO pins configured for digital output. The firmware employs simple yet dependable LED control logic, utilizing timer interrupts or task delays to generate blinking and pulsing patterns that correspond to specific system states, such as “active scan,” “detection confirmed,” or “low battery.” Each LED’s brightness and duty cycle are adjusted through Pulse Width Modulation control, providing visual modulation under different lighting conditions. In addition, the LEDs may be interfaced to the outer enclosure using light pipes, ensuring that optical feedback remains visible even when the electronics are enclosed within a ruggedized housing.

The pushbuttons are configured as GPIO inputs with internal pull-up resistors and are monitored via interrupt service routines (ISRs) to detect both short and long presses. Debouncing algorithms, implemented in software using millisecond-level sampling and hysteresis, ensure accurate button press recognition even in the presence of vibrations or environmental noise. Each button corresponds to a distinct control function, such as initiating a scan, toggling display modes, or acknowledging alerts. The firmware assigns each ISR to a FreeRTOS queue, allowing clean separation between input detection and main system logic. This architecture prevents input lag and ensures deterministic response times even when other tasks are active.

The buzzer and speaker form the auditory feedback system. The buzzer is controlled via PWM modulation on a GPIO pin, generating short alert tones for immediate notifications such as successful scans, user errors, or power warnings. The PAM8302A amplifier and capsule speaker provide broader-band audio output for higher-level feedback, such as voice alerts or alarm tones. The firmware includes an audio task that stores tone sequences in a lookup table and drives the amplifier using PWM waveforms at specific frequencies and durations. To prevent audio distortion and reduce power consumption, dynamic volume control is implemented based on the system's current state. For instance, the speaker is muted during silent modes or when user commands suppress audio feedback.

All UI firmware tasks operate under FreeRTOS to maintain real-time responsiveness and concurrency with sensor data processing. Priority assignment guarantees that user interactions, such as button presses or detection alerts, are promptly addressed without disrupting radar or thermal sampling tasks. The firmware also includes failsafe mechanisms, such as visual LED patterns or audible alerts triggered automatically during system faults or communication failures. These features ensure that the operator receives continuous system feedback, even if one sensory channel becomes ineffective due to environmental conditions.

Collectively, the UI firmware integrates all user-facing components into a cohesive, low latency control interface that enhances operational awareness and safety. The combination of SPI for high-speed display control, GPIO and PWM for discrete signaling, and interrupt driven input handling ensures the 2HAD-RAD system remains intuitive, responsive, and dependable in demanding field operations.

3.13.6. Power Management and System Control

The power management and protection firmware governs energy distribution, safety, and efficiency across the 2HAD-RAD system, ensuring stable operation under the dynamic loads imposed by high-frequency radar bursts, thermal frame acquisition, and display refresh cycles. The device's hardware power subsystem consists of a BQ24074 Li-ion charger, TPS61236P 5 V boost regulator, TPS563200 3.3 V buck regulator, and DW01A plus FS8205A-based 1S BMS board, all coordinated through firmware control routines managed by the ESP32-WROOM-32. This microcontroller acts as the supervisory unit, dynamically monitoring voltage rails, managing enable pins for individual regulators, and transitioning the system between performance and low power states as mission conditions require.

At startup, the firmware performs a power on self-test routine that verifies the availability of all regulated rails. The 5 V boost converter powers the radar and UWB modules, while the 3.3 V rail supplies the MCU, LCD display, and low-voltage peripherals. The ESP32 samples rail voltages through onboard ADC channels, using resistor dividers to scale measurements within safe limits. If voltage levels deviate outside predefined thresholds, firmware triggers an interrupt-driven protection response that disables non-critical loads and alerts the user via LED and buzzer indicators.

The BQ24074 charger IC interfaces indirectly with the ESP32 via status pins, allowing the firmware to monitor battery charging status and USB power availability. The firmware periodically reads these logic levels and interprets them to update system indicators. For instance, it illuminates a charging LED when external power is connected and signals the completion of charging once the battery reaches full capacity. The firmware also monitors battery voltage through an ADC channel connected to the BMS output, enabling real-time estimation of state of charge. To further enhance energy tracking, the system design optionally supports integration of a MAX17048 fuel gauge over I2C, allowing accurate coulomb counting and low battery prediction without imposing significant processing overhead.

To maximize operational runtime in the field, the power firmware employs dynamic power scaling and task-based sleep management. During idle periods, such as between radar chirps or thermal frame captures, the ESP32 enters a light sleep mode. This mode allows the device to maintain memory and peripheral states while disabling high-frequency clocks. Peripheral enable lines, controlled by GPIOs, selectively power down non-essential modules such as the display backlight, speaker amplifier, or radar subsystem when inactive. These transitions are orchestrated by a FreeRTOS power management task that tracks active sensor states and user input events, ensuring minimal power consumption without compromising responsiveness.

Protection mechanisms are implemented in both hardware and firmware to safeguard against overcurrent, overtemperature, and short circuit conditions. The DW01A/FS8205A BMS provides cell-level protection at the hardware level, while the firmware complements it by enforcing software current limiting and temperature-based derating. If the MCU detects a sustained high current draw or rapid temperature rise, it triggers an orderly shutdown sequence that disables power rails and notifies the operator through visual and audible alerts. Temperature readings can be optionally acquired through internal ESP32 sensors or the BMP280 barometric sensor on the INS module to provide environmental context for derating decisions. The firmware also includes fault logging and recovery routines to assist in post event diagnostics. Whenever a voltage, temperature, or current anomaly is detected, the system stores event metadata, including the timestamp and subsystem affected and measured parameters, in non-volatile memory. Upon restart, the firmware retrieves this log and displays the fault summary on the LCD, allowing operators or developers to diagnose the source of power disruptions.

Collectively, the power management and protection firmware integrates hardware regulation, real-time monitoring, and intelligent fault response into a cohesive energy control system. By combining ADC based voltage sensing, status pin monitoring, task level sleep coordination, and event-driven fault recovery, the 2HAD-RAD maintains stable, safe, and energy-efficient operation throughout extended search and rescue deployments.

3.13.7. Data Fusion and Detection Algorithm

At the core of the 2HAD-RAD’s processing architecture lies the data fusion firmware, which synthesizes inputs from the radar, UWB, thermal, and INS subsystems into a single, coherent life detection output. Each sensor operates at a unique sampling rate and communication bandwidth. To ensure consistent output, the ESP32-WROOM-32’s firmware employs time stamping, normalization, and frame alignment procedures. Every data stream, such as radar Doppler data, UWB channel impulse response, thermal frame, or INS motion vector, is assigned a synchronized timestamp. This timestamp is derived from the MCU’s high-precision hardware timer. These timestamps allow the fusion algorithm to correlate simultaneous signal events across different modalities, compensating for latency differences between sensor acquisition intervals.

The fusion routine runs as a dedicated FreeRTOS task with the highest processing priority, ensuring uninterrupted execution even under heavy peripheral activity. Data is shared between tasks through thread-safe message queues, preventing frame loss or overlap during concurrent sensor operations. The firmware performs multi-sensor correlation, weighting radar and UWB data more heavily for motion-based detection while using thermal imagery for biological confirmation. Radar and UWB outputs are first evaluated through cross correlation functions to determine temporal and amplitude overlap in their reflected signal profiles. A multi-threshold decision logic is then applied, requiring both sustained radar motion signatures and corresponding UWB energy peaks that persist within a defined time window. Once this temporal condition is satisfied, the algorithm references the most recent thermal frame to confirm the presence of a localized heat source above ambient temperature threshold.
The INS data is simultaneously processed to characterize device orientation and motion. By analyzing acceleration and angular velocity vectors, the firmware

identifies user-induced movement or tilt that could generate false radar or UWB signals. When motion beyond a defined threshold is detected, the fusion algorithm temporarily suppresses life-detection classification until the device stabilizes. This approach guarantees that environmental vibrations or operator repositioning won't cause false positives, thereby maintaining a high level of detection accuracy in the field.

Filtering and decision making occur in stages: radar and UWB data undergo temporal smoothing and energy envelope extraction, thermal data is analyzed using intensity clustering, and INS data is filtered through a complementary filter to isolate genuine device motion. The resulting features are fed into a rule-based decision matrix, which classifies detections as either "transient," "suspected," or "confirmed." The firmware elevates the classification to "life detected" only when multiple modalities consistently reinforce each other over several frames. For instance, persistent radar oscillations synchronized with UWB amplitude fluctuations, and a stable thermal hotspot are required.
Once a detection is confirmed, the fusion output is routed to the user interface subsystem, where both visual and auditory alerts are triggered in real time. The LCD displays an alert banner and highlights detection confidence levels, while LEDs and the buzzer provide redundant cues for high-noise environments. The firmware records each detection event with timestamp, confidence score, and sensor contribution weight in non-volatile memory, allowing later review or algorithm refinement.

Looking ahead, the rule based fusion framework is intentionally modular to support transition toward a machine-learning-based classifier. Future iterations will incorporate lightweight embedded AI models, such as decision trees or neural networks, that are trained on multi-sensor datasets collected during field testing. These models dynamically adjust detection thresholds based on environmental context and noise patterns, enabling adaptive fusion intelligence that can identify subtle human signatures in complex and cluttered conditions.
Together, the data fusion firmware transforms the raw, asynchronous outputs of individual sensors into a unified, reliable life-detection decision engine. By combining temporal synchronization, sensor-specific filtering, and correlation-based validation, it provides the computational foundation for the 2HAD-RAD's mission: rapid, accurate, and field-ready human presence detection through obstacles and debris.

3.13.8. Communication Summary

Each subsystem's firmware is developed in C/C++ using either the Arduino framework or the Espressif IoT Development Framework (ESP-IDF), which are the two primary environments supported by the ESP32-WROOM-32. The Arduino environment offers a simplified interface and extensive library support, making it an excellent choice for rapid prototyping and hardware validation. On the other hand, the ESP-IDF provides advanced real-time task control, memory optimization, and low-level peripheral access through FreeRTOS, enabling the development of robust, production-grade firmware. The choice between these environments provides flexibility for team members to work within their preferred ecosystem while ensuring interoperability and code consistency across modules. To ensure efficient and reliable communication between the MCU and all connected subsystems, optimal data protocols were selected based on bandwidth, timing, and synchronization requirements. The IWR6843AOP radar module, which produces large volumes of motion and Doppler data, would benefit most from SPI communication due to its high throughput and low latency, ensuring fast frame transfers and minimal signal delay during real-time detection. However, since the radar's evaluation module (EVM) is configured by default for UART output, the firmware will initially use UART communication for simplicity. Later design revisions will transition to SPI for improved performance. The DWM3000 UWB module can also operate over SPI or UART. However, SPI is preferred in the final design to maintain synchronization with radar sampling intervals and minimize overhead during data bursts. The MLX90640 infrared thermal array and M5Stack M135 INS module both operate optimally over I2C, which supports multiple peripheral devices on a shared bus while maintaining sufficient speed for moderate-rate sensor polling. Finally, GPIO and PWM interfaces are reserved for UI components such as LEDs, pushbuttons, the buzzer, and the speaker, where deterministic control and minimal latency are required for real-time operator feedback.

Regardless of the selected environment, the firmware architecture ensures consistent timing, communication reliability, and modular expandability. Filtering algorithms, such as temporal averaging, Gaussian smoothing, and complementary orientation estimation, ensure signal fidelity even in noisy operating conditions. Collectively, these firmware modules constitute a flexible and scalable software backbone for the 2HAD-RAD system, enabling future enhancements like BLE connectivity, machine-learning-based detection, and advanced data visualization capabilities.

3.14. Communication Frameworks for Smartphone App Development

When designing a smartphone-connected embedded device, there are a handful of options for which communication framework to use. The choice affects power consumption, responsiveness, user experience, and scalability.

3.14.1. Bluetooth Low Energy (BLE)

BLE is part of the Bluetooth 4.0+ specification and is specifically tuned for low-power, short-range communication. It's designed around brief data exchanges rather than continuous streams, which makes it ideal for mobile-connected sensing devices, health monitors, and wearables, which is precisely our design case. The ESP32 family (WROOM-32, S3, C6) supports BLE peripheral mode, enabling our device to advertise itself to smartphone apps.

Bluetooth Low Energy (BLE) communication has some drawbacks, including low data throughput (theoretically around 1-2 Mbps, usually 200-300 kbps) and a limited range. However, due to the nature of our data, not much throughput is required, and our use-case is for the smartphone to be relatively close to the device (no more than 20 meters).
BLE aligns well because we want real-time updates with minimal power, the ESP32 line implements complete BLE GATT APIs, and both OS ecosystems handle BLE securely without internet.

3.14.2. Wi-Fi (HTTP / WebSocket)

Wi-Fi connectivity leverages standard TCP/IP for data transfer. Our ESP32 can act as either a station (connected to a router or access point), or a soft access point (AP) (hosting its own network). Once on the same network, the MCU and smartphone communicate via HTTP REST endpoints (pull-based or push via POST) and WebSockets (persistent real-time TCP connection). This model suits dashboards, configuration pages, or cloud-connected scenarios.
Benefits of this setup include high bandwidth, compatibility with different MCUs, and much higher distance than BLE. However, some drawbacks include a higher power consumption, setup friction (user must join network and manage credentials), and potential latency issues.

Wi-Fi is well-suited for stationary setups (such as a lab or a smart home product) with power availability, cases where we'd want to use a web dashboard or local API access over the LAN, or scenarios where multiple mobile clients connect simultaneously. These are not really scenarios of use for our device, which is handheld, prefers minimal setup friction, and doesn't need any APIs.

3.14.3. MQTT (Message Queuing Telemetry Transport)

MQTT is a lightweight publish/subscribe protocol built over TCP/IP, engineered for IoT systems that span devices and networks. Rather than a smartphone connecting directly to the MCU, both connect to a broker (cloud or local server). Messages consist of simple topics (e.g., device123/co2, device123/heartbeat). This decouples data producers (MCUs) and consumers (apps) entirely, enabling robust asynchronous systems.

Some benefits of MQTT include its scalability to multiple devices and users, power efficiency, and cloud integration. Some drawbacks are that it requires an internet connection and a broker setup, it is less less latent than BLE, it requires cloud management, and it increases firmware complexity.
MQTT is excellent if we envision our device evolving toward a fleet of cloud-connected monitoring (multiple deployed units), real-time data dashboards from anywhere, and multi-user collaboration or server-side analytics. However, for localized personal sensing, direct BLE or Wi-Fi is much simpler and fit better into the scope of our project, so we are not going with MQTT.

3.14.4. USB (Wired Communication)

USB is the traditional physical data link between embedded devices and host computers or smartphones. Only some MCUs (like ESP32-S3) natively support USB 2.0 OTG, enabling them to act either as a serial device or even HID (Human Interface Device)/MSC (Mass Storage Class) device. With appropriate adapters, such as a USB-C to micro-USB cable, smartphones can directly interface for data transfer or configuration using native serial drivers.

Benefits of this method of communication include reliability, high throughput, security, instant pairing and setup. Some drawbacks include the fact that it requires physical tethering, and that it adds unnecessary mechanical complexity. USB is most useful for development and debugging purposes, such as serial logging and fast firmware updates. It's also suitable for calibration or maintenance connections where a stable wired link is crucial. However, it's not a viable option as the primary user interface in portable or consumer devices. Since USB is not suitable for our device usage conditions, such as the inconvenience of always being connected to a smartphone, we opted not to use this method for smartphone connection.

Communication Protocol	Typical Use Case	Pros	Cons
BLE (Bluetooth Low Energy)	Direct local link; common in ESP32/S3/C6	Power-efficient, works offline	Lower bandwidth; requires phone proximity
Wi-Fi (HTTP / WebSocket)	Local LAN or MCU AP mode	Simple TCP/IP; easy browser integration	MCU must run local web server; more power use
MQTT over Internet	Cloud integration	Reliable, standard IoT protocol; good for multi-device systems	Requires broker setup; needs network connectivity
USB (if MCU supports it like ESP32-S3)	Debugging, wired link	Fast, direct	Limited user convenience, adds physical complexity

Table 24: Smart Phone Application Communication Comparison

3.15. Smartphone Application Development

3.15.1. Development Infrastructure Comparison

An advanced goal of our project is to develop a smartphone app to display output data from the device, which is data the device MCU receives from the sensors

(radar, IR, UWB, etc). For app development, there are two main options: Android or iOS.
Some factors that determine whether to develop for iOS or Android include ease of development, integration with the MCU low setup friction, and support for data visualization (UI). Due to such factors, we believe native Android development is the most pragmatic path and not iOS. Android provides flexible access to Bluetooth, Wi-Fi, and USB APIs. We can run the APK directly on any Android phone or even an emulator, and there will be no provisioning or Apple developer account friction. Within the Android development space, there are a handful of approaches to take, varying in programming languages, UI implementation, novelty, community support and documentation, et cetera.

3.15.2. Android Technology Comparison

3.15.2.1. Android Native (Kotlin + Jetpack Compose)

This is Google's official modern pathway for Android. We would use Kotlin, with Android Studio. UI will be built using Jetpack Compose, which is a declarative toolkit like React for native components.
Benefits of this approach are a modern, clean syntax and reduced boilerplate versus Java, excellent Android documentation and tooling. Jetpack Compose simplifies the rapid iteration of user interfaces for dashboards and gauges, which is essentially our objective. This works seamlessly with any ESP32 BLE example. In addition, this approach gives us full access to system permissions, background services, and notifications. One drawback of this approach is its relatively steep learning curve, as we are new to Kotlin. Jetpack Compose is still evolving, which means some advanced UI behaviors are not well-documented. Additionally, Android Studio can be resource-intensive on modest computers. Overall, this is our choice because it's modern and thus relevant to apps built today, well supported documentation, good UI tools, and total freedom in terms of the customization of our app.

3.15.2.2. Android Native (Java + XML Layouts)

This approach is the 'classic' Android stack (Java + XML). It is still fully supported and many BLE and HTTP tutorials exist in Java, including Espressif's legacy Android templates.
Benefits of this approach include the fact that it has huge existing codebase examples, it offers easier adoption for anyone with experience in Java or classic OOP patterns, and it is stable as almost any Android version supports it. Drawbacks of this approach include the fact that it has verbose syntax and more boilerplate than Kotlin/Compose. Its UI iteration is slower as it must recompile layout XML, it is an older approach and thus less future-proof.
We wanted to go for a more modern development approach and avoid Java, and the UI tools in this approach seem less attractive and flexible. Hence, we did not choose this development path.

3.15.2.3. Android + MIT App Inventor / Kodular (Visual Builder)

MIT App Inventor, Kodular, and similar web-based tools allow graphical block programming. Devs drag/drop UI components (buttons, labels, charts) and assign logic visually. The benefits of this approach are that it boasts a very fast development cycle, no deep programming expertise required. It's great for visual demonstrations or classroom prototypes and offers a free platform (MIT App Inventor) with built-in emulation.
One downside of this approach is that it offers limited customization beyond what blocks allow, which may result in a less professional-looking GUI compared to native approaches. In addition, it cannot implement complex background threads or advanced parsing easily, and the BLE stack is less reliable on recent Android versions. This approach is not our choice due to limited UI customization and a non-traditional development pattern.

3.15.2.4. Android with Flutter (Dart)

Flutter is a high-level way to build Android apps in Dart (a programming language developed by Google). It offers UI deployed via native rendering. In terms of complexity, this approach sits between native based implementation (building platform-specific apps using each platform's official languages and tools) and web-based implementation (building apps that run in a web-based application, such as HTML/CSS/JavaScript frameworks).
Benefits of this technology are that it offers a fast reload and instant preview of UI changes, structured UI that is visually impressive, and easier charting and data visualization using libraries. Drawbacks include the fact that there's extra start-up overhead (installing Flutter SDK), and it's not as low-level as native Android if we need fine control. We did not go with this choice because the learning curve of Flutter and Dart seems too steep for the relatively simple application we are trying to develop.

3.15.2.5. Android Native (Java + XML Layouts)

This approach is the 'classic' Android stack (Java + XML). It is still fully supported and many BLE and HTTP tutorials exist in Java, including Espressif's legacy Android templates.
Benefits of this approach include the fact that it has huge existing codebase examples, it offers easier adoption for anyone with experience in Java or classic OOP patterns, and it is stable as almost any Android version supports it.
Drawbacks of this approach include the fact that it has verbose syntax and more boilerplate than Kotlin/Compose. Its UI iteration is slower as it must recompile layout XML, it is an older approach and thus less future-proof. We wanted to go for a more modern development approach and avoid Java, and the UI tools in this approach seem less attractive and flexible. Hence, we did not choose this development path.

3.15.2.6. Android + MIT App Inventor / Kodular (Visual Builder)

MIT App Inventor, Kodular, and similar web-based tools allow graphical block programming. Devs drag/drop UI components (buttons, labels, charts) and

assign logic visually. The benefits of this approach are that it boasts a very fast development cycle, no deep programming expertise required. It's great for visual demonstrations or classroom prototypes and offers a free platform (MIT App Inventor) with built-in emulation.

One downside of this approach is that it offers limited customization beyond what blocks allow, which may result in a less professional-looking GUI compared to native approaches. In addition, it cannot implement complex background threads or advanced parsing easily, and the BLE stack is less reliable on recent Android versions. This approach is not our choice due to limited UI customization and a non-traditional development pattern.

3.15.2.7. Android with Flutter (Dart)

Flutter is a high-level way to build Android apps in Dart (a programming language developed by Google). It offers UI deployed via native rendering. In terms of complexity, this approach sits between native based implementation (building platform-specific apps using each platform's official languages and tools) and web-based implementation (building apps that run in a web-based application, such as HTML/CSS/JavaScript frameworks).

Benefits of this technology are that it offers a fast reload and instant preview of UI changes, structured UI that is visually impressive, and easier charting and data visualization using libraries. Drawbacks include the fact that there's extra start-up overhead (installing Flutter SDK), and it's not as low-level as native Android if we need fine control. We did not go with this choice because the learning curve of Flutter and Dart seems too steep for the relatively simple application we are trying to develop.

Approach	Language / Framework	Communication Libraries / SDKs	UI Development Style	Best For
Android Native (Kotlin + Jetpack Compose)	Kotlin (Android Studio)	Android BluetoothGatt, Retrofit (HTTP), MQTT Client	Modern declarative (Compose)	Clean modern prototype, BLE-based MCU comms
Android Native (Java + XML Layouts)	Java (Android Studio)	Same as above + legacy libs (OkHttp, Paho MQTT)	XML & View system	Stable legacy base, extensive examples
Android + MIT App Inventor / Kodular	Block-based visual programming	Built-in BLE and Web APIs	Drag-and-drop GUI	Ultra-fast demos, limited customization

Approach	Language / Framework	Communication Libraries / SDKs	UI Development Style	Best For
Android with Flutter (Dart)	Dart (Flutter SDK)	flutter_blue_plus (BLE), http/mqtt_client	Declarative widget tree	Quick modern UI with BLE, single codebase optional

Table 25: Android Technology Comparison

CHAPTER 4

4.1. Introduction to Standards and Design Constraints

This chapter defines the rules and limits that govern the design of the 2HAD-RAD device. Since the system integrates multiple radio technologies, 60 GHz FMCW radar, 2.4 GHz BLE, and UWB ranging, it can't be treated as a simple embedded project; it must comply with U.S. FCC Part 15 requirements for both intentional and unintentional radiators to operate legally in unlicensed spectrum and to avoid causing interference to other services. At the same time, several industry and international standards (IEEE 802.15.4a/z for UWB, GNSS ICDs for multi-constellation positioning, and I²C/SPI serial bus specifications) dictate how subsystems communicate, how signals must be structured, and what levels of accuracy or interoperability are expected. Beyond standards, the device is also constrained by practical engineering factors, available power (≤ 10 W), passive thermal dissipation inside a small enclosure, handheld mechanical form factor, limited computational resources on the ESP32, and the cost ceiling typical of a university prototype. Finally, the project is constrained by ethical and safety considerations: RF exposure must remain within established limits, collected data must remain on-device, and the intended use must align with life-detection and rescue scenarios. The sections that follow translate these regulatory, technical, and practical requirements into concrete design decisions for RF emissions, communication buses, enclosure layout, power distribution, and verification/documentation practices.

4.1.1. FCC Part 15 – Radio Frequency Devices

2HAD-RAD both transmits and receives radio frequency (RF) energy through its 60 GHz radar module, and it communicates with external devices using Bluetooth Low Energy (BLE). Therefore, our device is classified by the U.S. Federal Communications Commission (FCC) as an intentional radiator under Title 47 CFR Part 15, which means it is a "device that intentionally generates and emits radio frequency energy by radiation or induction that may be operated without an individual license," [1]. Thus, our team must ensure that we adhere to the standards for intentional radiators. These standards include strict limits on output power, spectral emissions, and spurious radiation to prevent interference with licensed communication services.

Furthermore, under Part 15 Subpart B, all digital electronics within the device, such as the microcontroller, analog-to-digital converters, and switching regulators, etc., must comply with radiated and conducted emission limits for unintentional radiators, which the FCC defines as "a device that by design uses digital logic, or electrical signals operating at radio frequencies for use within the product, or sends radio frequency signals by conduction to associated equipment via connecting wiring, but is not intended to emit RF energy wirelessly by radiation or induction." To meet these requirements, the PCB layout will use continuous ground planes, short return paths, and decoupling capacitors on every supply rail. The switching regulators are enclosed within grounded copper pours, and all high-speed clock traces are impedance-matched to reduce harmonic radiation.

The BLE radio, operating in the 2.4 GHz ISM band, is regulated by Part 15 Subpart C, which regulated low-power intentional radiators. The module transmits below +10 dBm EIRP, occupies the 2400-2483.5 MHz band, and employs frequency-hopping to minimize interference. The microcontroller, the ESP32-WROOM, is already FCC-certified, so integration only requires maintaining its approved antenna configuration and keeping trace lengths short to maintain pre-designed radiated-emission characteristics.
Additionally, the 60 GHz radar module, the TI IWR6843AOP EVM, operates under Part 15 Subpart C, Section 15.255, which governs unlicensed use of the 57–71 GHz ISM band. This section specifies the maximum allowable equivalent isotropic radiated power (EIRP), power density, and out-of-band emissions to prevent interference with licensed millimeter-wave services such as point-to-point backhaul and satellite downlinks. The rule limits continuous-wave emissions to +40 dBm EIRP peak and requires that any emissions outside the 57–71 GHz range be attenuated by at least 20 dB below the fundamental. Our radar's antenna-on-package (AoP) design confines energy within this band, and the measured EIRP is well below the regulatory maximum, allowing for full compliance without additional attenuation circuitry.
Through these design measures, the 2HAD-RAD is compliant with the FCC Part 15 Subparts B and C, covering both intentional and unintentional radiators, thereby ensuring lawful operation in unlicensed bands and preventing harmful interference to other users of the electromagnetic spectrum.

4.1.2. IEEE 802.15.4 UWB Standard (802.15.4a / 802.15.4z)

The 2HAD-RAD system utilizes an Ultra-Wideband (UWB) transceiver, the DWM3000EVB, that operates under the IEEE 802.15.4a/z physical layer standard for high-precision ranging and localization. The IEEE 802.15.4 framework defines low-rate wireless personal area networks (LR-WPANs), but its UWB extensions (802.15.4a and 802.15.4z) specify impulse-radio transmission techniques optimized for time-of-flight (ToF) measurements and secure ranging. These standards are crucial to our project because they establish how short-duration pulses are generated, modulated, and time-synchronized to allow for centimeter-level accuracy while not interfering with other wireless systems in the 3.1–10.6 GHz spectrum.

Under IEEE 802.15.4a, UWB signals are transmitted as nanosecond-scale pulses spread across a wide bandwidth, typically exceeding 500 MHz, to achieve extremely fine temporal resolution. This spread-spectrum approach produces a low power spectral density (below –41.3 dBm/MHz), ensuring that UWB transmissions do not cause harmful interference to narrowband systems such as Wi-Fi, Bluetooth, or LTE. The 802.15.4z revision, introduced later, enhanced security and improved ranging robustness by defining Secure Ranging Exchange (SRE) protocols, Scrambled Timestamp Sequences (STS), and improved pulse-position modulation schemes. These updates allow UWB modules like the DWM3000EVB to authenticate distance measurements and resist spoofing or relay attacks, which is an important feature when multiple devices are operating in proximity during search-and-rescue operations.

In the 2HAD-RAD design, the UWB module supports time-of-flight ranging for the handheld device. The system operates within the 6.5–8.0 GHz band, fully aligned with FCC Part 15 Subpart F (for UWB transmitters) and the IEEE spectral mask requirements. Transmission duty cycles, burst durations, and pulse repetition frequencies are configured to maintain compliance with both the average and peak power limits defined in the standard. The DWM3000EVB's hardware implementation employs time-hopping and preamble codes specified in the 802.15.4a/z PHY to minimize cross-interference and support multi-user operation.

The UWB's PCB incorporates controlled impedance microstrip lines for the antenna feed and maintains a 50 Ω characteristic impedance across the RF path. Ground-via stitching surrounds the antenna region to confine the radiation field and reduce coupling into adjacent circuits.
Through adherence to these standards and verification practices, the 2HAD-RAD ensures proper UWB communication in accordance with both IEEE 802.15.4a/z and FCC Part 15 Subpart F.

4.1.3. GNSS System Standards (GPS, Galileo, BeiDou ICD compliance)

The 2HAD-RAD system utilizes a Global Navigation Satellite System (GNSS) receiver as part of its inertial navigation subsystem to obtain absolute position in outdoor environments and to synchronize relative motion data collected indoors. The selected module, the M5Stack M135, integrates a u-blox NEO-M9N GNSS receiver, which is compliant with the Interface Control Documents (ICDs) for major satellite constellations including the United States Global Positioning System (GPS), Europe's Galileo, China's BeiDou, Russia's GLONASS, and Japan's QZSS systems. Each constellation defines specific navigation message formats, ranging codes, and signal structures through its ICD, which manufacturers must adhere to ensure interoperability, accuracy, and reliability across global systems.

Under these ICD standards, such as IS-GPS-200H (for GPS L1 C/A), Galileo OS SIS ICD, and BeiDou B1I ICD, GNSS receivers must decode standardized ephemeris and almanac data, pseudorange measurements, and timing signals to compute 3D coordinates and velocity. The M135's multi-constellation support ensures compliance with these defined parameters by implementing the appropriate signal-tracking loops, correlator spacing, and message decoding protocols required for each constellation. The receiver's firmware follows the ICD-defined navigation data structures to interpret satellite orbital elements, health status, and time offsets, thereby producing position estimates consistent with global navigation standards.

In addition to the ICD signal compliance, the GNSS subsystem adheres to international performance and accuracy standards such as RTCA DO-229D and IEC 61108-1, which specify minimum navigation accuracy, time-to-first-fix (TTFF), and integrity parameters for civil receivers. The NEO-M9N achieves horizontal position accuracy of approximately 1.5 m (CEP) under open-sky conditions and supports update rates up to 25 Hz, meeting or exceeding these defined thresholds. The receiver also conforms to ITU-R Recommendation M.1477, which governs the spectral characteristics and bandwidth limits for GNSS emissions, ensuring that its reception is compatible with globally allocated navigation frequency bands (L1/E1 ≈ 1575 MHz).

To maintain signal integrity and electromagnetic compatibility within the 2HAD-RAD system, the GNSS antenna is mounted on a grounded RF plane with a clear line of sight to the sky. The PCB layout routes the low-noise amplifier (LNA) input with controlled impedance and isolates it from digital switching lines to minimize coupling and spurious noise pickup. Ferrite bead filters and RC networks are placed near the antenna feed and power supply to further suppress conducted interference.

The combination of ICD adherence, multi-constellation compliance, and robust EMC design ensures that the 2HAD-RAD's GNSS subsystem delivers accurate and reliable absolute positioning in full accordance with GPS, Galileo, and BeiDou system standards.

4.1.4. I2C / SPI Interface Standards (Serial Bus Communication)

The 2HAD-RAD system employs industry-standard serial communication protocols, Inter-Integrated Circuit (I2C) and Serial Peripheral Interface (SPI), to ensure reliable and standardized data exchange between the microcontroller and its peripheral sensors. The I²C bus, originally defined by Philips (now NXP), provides a two-wire communication interface consisting of a data (SDA) and clock (SCL) line shared among multiple devices. It supports multi-master, multi-slave configurations and uses open drain signaling with pull-up resistors to maintain proper logic levels. Within 2HAD-RAD, the I2C interface operates at 3.3 V in Fast Mode (400 kHz) to connect low-bandwidth devices such as the barometric sensor (BMP280), magnetometer (BMM150), and infrared thermal sensor. Trace lengths and pull-up resistor values are selected in accordance with the I2C specification to minimize capacitive loading and preserve waveform integrity.

For higher-speed data exchange, the system utilizes the SPI standard, first established by Motorola, which defines a four-wire interface consisting of a master clock (SCLK), chip select (CS), master-out–slave-in (MOSI), and master-in–slave-out (MISO) lines. SPI enables full-duplex communication at several megahertz and is used for bandwidth-intensive modules such as the UWB transceiver and radar interface. Following the SPI electrical and timing requirements, all high-speed lines are impedance-controlled and routed with short, direct connections to reduce crosstalk and electromagnetic interference. By adhering to these standardized serial bus protocols, the 2HAD-RAD ensures interoperability, data consistency, and timing synchronization between all sensing subsystems, supporting robust and deterministic communication across the device.

4.1.5. IEC Safety and EMC Standards (ESD, EMI mitigation)

The 2HAD-RAD system is designed in accordance with International Electrotechnical Commission (IEC) safety and electromagnetic compatibility (EMC) standards to ensure reliable and interference-free operation in mixed electronic environments. Specifically, it aligns with IEC 61000-4 series for EMC immunity testing and IEC 61326 for electrical equipment operating in industrial and scientific environments. These standards define procedures and performance limits for exposure to electrostatic discharge (ESD), electromagnetic interference (EMI), and radiated or conducted disturbances, ensuring that the device maintains normal function without causing or suffering from disruptive emissions.

To comply with IEC 61000-4-2 (ESD immunity), the design incorporates transient voltage suppression diodes on external I/O lines, grounded shielding, and a copper pour around connectors and exposed interfaces. These measures protect sensitive circuits such as the radar front end and GNSS receiver from electrostatic discharges up to ±8 kV (contact) and ±15 kV (air), meeting the IEC test levels for portable electronics. For EMI mitigation under IEC 61000-4-3 and -4-6, the PCB layout uses continuous ground planes, ferrite bead filtering on power inputs, and differential signal routing for high-speed lines like SPI and UART. The radar module is further enclosed within a metal EMI shield to contain radiated emissions at 60 GHz and prevent coupling into digital subsystems. By following these IEC safety and EMC practices, the 2HAD-RAD minimizes susceptibility to electrical disturbances, prevents interference with nearby equipment, and ensures user safety and operational stability across diverse electromagnetic environments.

4.1.6. ISO / IEC Quality and Testing Standards

While the 2HAD-RAD prototype will not undergo formal certification or laboratory testing, its design and development process reference key ISO and IEC quality and testing standards to promote consistency, reliability, and traceability. Standards such as ISO 9001:2015 for quality management and IEC 60068 for environmental durability provide a framework that guides the team's approach to documentation, component selection, and design verification. Although full qualification testing will not be performed, the principles of these standards, structured reviews, defined acceptance criteria, and iterative validation, are applied throughout the development cycle to maintain engineering discipline and repeatability.

For sensor calibration, the team follows best practices derived from ISO/IEC 17025, which defines general requirements for testing and calibration accuracy. Instead of formal laboratory procedures, sensors such as the barometer, magnetometer, and IMU will be evaluated through controlled functional tests and manufacturer datasheet validation to ensure consistent readings under typical operating conditions.

By referencing these ISO and IEC frameworks conceptually rather than implementing full compliance testing, the 2HAD-RAD project maintains a realistic balance between academic prototyping and professional engineering standards, ensuring design credibility, organized workflow, and defensible technical documentation without exceeding the scope of a student research prototype.

4.1.7. Power Constraints

The 2HAD-RAD is a portable, battery-powered life detection and localization device designed for use in field environments where access to external power is limited. Therefore, strict adherence to power efficiency and distribution constraints is critical to maintain system performance, reliability, and operating time. The overall system is designed to operate within a 10 W total power budget, which encompasses all sensing modules, the microcontroller, display, and user interface components. To meet this requirement, each subsystem has been selected and configured to balance performance with low power consumption while maintaining sufficient processing speed and sensing accuracy for real-time operation.

The most power-intensive component in the system is the TI IWR6843AOP radar module, which consumes approximately 2.0–2.5 W during continuous operation. This radar serves as the primary sensing element, performing frequency-modulated continuous wave (FMCW) transmissions and on-chip digital signal processing. Its internal design integrates both analog front-end and baseband processors, minimizing the need for external computation but necessitating dedicated power regulation and thermal management. To ensure stable operation, the radar is supplied through a dedicated 3.3 V DC-DC buck converter with input filtering capacitors and ferrite beads to suppress switching noise that could otherwise interfere with nearby sensors.

Supporting sensors such as the M5Stack M135 INS module, barometric pressure sensor (BMP280), and magnetometer (BMM150) together contribute less than 0.2 W combined. These low-power devices operate intermittently, polling data at controlled intervals rather than continuously, reducing average consumption. The DWM3000EVB UWB module consumes approximately 100–150 mW during active ranging, and its duty cycle is managed through firmware scheduling to avoid simultaneous high-current draws from multiple RF subsystems. Similarly, the ESP32-WROOM microcontroller, which handles data fusion, display control, and wireless communication, operates around 500–600 mW under dual-core processing, with deep-sleep modes utilized when radar or ranging tasks are inactive.

The system's power distribution architecture employs multiple low-dropout (LDO) and switching regulators to provide isolated rails for analog, digital, and RF sections, minimizing crosstalk and ground noise. The PCB layout segregates high-current traces and uses wide copper pours to limit voltage drop and heat generation. Each subsystem includes local decoupling capacitors to mitigate transient current spikes during RF transmission or processor wake events. These measures ensure that the device maintains stable voltage regulation and electromagnetic cleanliness even under dynamic load conditions.
Battery selection is based on a 7.4 V lithium-ion pack with a nominal capacity of 3000–4000 mAh, regulated down to subsystem voltage levels. Assuming an average system draw of approximately 6–7 W during continuous use, the device can sustain operation for 45–60 minutes per charge, suitable for prototype testing and short deployment scenarios. In future iterations, implementing power gating, adaptive sensor polling, and radar duty cycling could extend runtime substantially without sacrificing functionality.

Through these architectural and design decisions, the 2HAD-RAD system remains within its defined 10 W constraint while ensuring safe, efficient, and thermally stable operation. The deliberate separation of high-power RF components from low-power sensor subsystems, coupled with robust power management and regulation strategies, ensures that the device delivers reliable performance within realistic energy limitations for a field-deployable search-and-rescue prototype.

4.1.8. Thermal Constraints

The 2HAD-RAD system must maintain stable operation across a wide range of environmental conditions while preventing excessive internal heat buildup from its high-power electronic components. Because it integrates both low-power sensors and high-frequency radio subsystems, attention to thermal design and heat dissipation is essential for reliability, user safety, and sensor accuracy. The device is engineered to operate within a 0 °C to +70 °C ambient range, consistent with typical electronic-grade operating limits, while individual components such as the radar and power regulators are rated for up to +85 °C junction temperature.

The most thermally demanding component in the system is the TI IWR6843AOP radar module, which dissipates approximately 2.0–2.5 W of power during active transmission and onboard signal processing. This heat is concentrated within a compact 9 × 6 cm board area containing both RF front-end and DSP cores. To manage this thermal load, the radar PCB section is mounted to a copper ground plane with thermal vias that transfer heat into internal copper layers and the aluminum mounting plate of the enclosure. This conduction path distributes heat away from the sensitive RF components and reduces localized temperature rise. The enclosure itself serves as a passive heatsink, enabling the device to operate continuously without active cooling while maintaining surface temperatures below safe handling limits.

Additional heat contributors include the ESP32-WROOM microcontroller (~0.5 W) and the DC-DC converters powering the radar and UWB subsystems. These regulators are positioned near ventilation cutouts within the enclosure to promote natural convection. The PCB layout isolates high-dissipation components from temperature-sensitive sensors such as the barometer and magnetometer, which can experience accuracy drift due to thermal gradients. By maintaining a minimum spacing of 20 mm and using local copper shielding, thermal coupling between these subsystems is minimized.

To ensure safe operation under load, the firmware continuously monitors internal temperature sensors located on the radar module and the microcontroller. If the measured temperature exceeds predetermined thresholds (typically +80 °C for the radar and +70 °C for the MCU), the system can automatically enter a reduced-power mode by lowering radar duty cycle or suspending nonessential tasks. This thermal throttling mechanism prevents overheating and extends component lifespan.

Through a combination of passive conduction, strategic component placement, and firmware-based temperature management, the 2HAD-RAD maintains compliance with its thermal design constraints. This ensures continuous, safe, and reliable operation even during extended use in confined or elevated-temperature environments, critical for field-deployable search-and-rescue applications.

4.1.9. Mechanical / Size Constraints

The 2HAD-RAD is designed as a portable, handheld search-and-rescue device, which requires strict mechanical and dimensional constraints on the internal layout, enclosure design, and component integration. To remain practical for field use, the total system, including the radar, microcontroller, sensors, and power source, must fit within a compact, lightweight housing that can be comfortably held, mounted, or carried by a single operator. The enclosure dimensions are targeted at approximately 160 × 100 × 60 mm, balancing internal volume for components and airflow with ergonomic handling and manufacturability using 3D printing or CNC machining.

Mechanical constraints are dictated primarily by the TI IWR6843AOP radar module, which requires unobstructed line-of-sight for 60 GHz transmission and reflection detection. This necessitates precise placement of the radar aperture relative to the front face of the enclosure, avoiding any metallic surfaces or thick dielectric materials that could distort the electromagnetic beam. The radar PCB is aligned flush with a non-metallic window, typically acrylic or polycarbonate, chosen for its low dielectric loss at millimeter-wave frequencies. The enclosure front wall thickness is maintained below 1.5 mm in the radar's emission zone to minimize attenuation, while the remaining walls use reinforced ribs for rigidity. The mechanical mounting of internal components follows a layered architecture. The main PCB, housing the ESP32 microcontroller, power regulators, and sensor connectors, serves as the structural backbone. Above it, the user interface PCB, containing LEDs, push buttons, and the buzzer, is mounted orthogonally to align with the front panel openings. Both boards are secured with nylon standoffs and M2.5 screws, ensuring vibration resistance while avoiding conductive shorting paths. Critical connectors, such as the 60-pin Samtec header between the main board and radar module, are positioned to maintain mechanical stability and consistent impedance alignment without strain or torsion.

The design also accounts for environmental mechanical stressors. The enclosure includes integrated gaskets and sealed interfaces to protect against dust and moisture during field operation. Internal cable routing is minimized and bundled using short flexible jumpers to reduce wear and mechanical noise coupling. To mitigate the effects of handling shock or drops, heavy components like the lithium-ion battery pack and radar board are mounted close to the enclosure's center of mass and cushioned with thin neoprene pads that absorb impact energy without restricting heat dissipation.

Mechanical tolerances are maintained within ±0.2 mm for 3D-printed prototype assemblies to ensure proper alignment of mating surfaces and connectors. The final design must also allow for ease of assembly, maintenance, and component replacement, particularly for sensors or modules that may require recalibration or reconfiguration.

Through adherence to these mechanical and spatial constraints, the 2HAD-RAD achieves a compact, structurally sound design that supports functional alignment of RF, power, and sensing subsystems. The result is a mechanically robust prototype that balances durability, manufacturability, and ergonomic usability while maintaining the precise physical alignment necessary for reliable radar and sensor performance.

4.1.10. Cost Constraints

The 2HAD-RAD is designed under a strict academic budget, with a total allowable cost of approximately $560, including all sensing, computing, and power components. The largest cost driver is the TI IWR6843AOP radar module, priced around $175, which represents a large portion of the total budget but is essential for achieving real-time vital-sign detection capabilities. Remaining funds are allocated toward the ESP32-WROOM microcontroller, UWB module, sensors, battery, and 3D-printed enclosure. Design choices prioritize integration of multifunctional modules over discrete components to reduce both cost and complexity. The team also minimizes PCB fabrication and assembly expenses by consolidating subsystems onto shared boards. These trade-offs ensure a cost-effective yet fully functional prototype that meets design requirements within the financial limits of a university project.

4.1.11. Environmental Constraints

Since the 2HAD-RAD is intended for field operation in search-and-rescue environments, it must withstand variable temperature, humidity, dust, and mechanical vibration. The system is designed for reliable operation between 0 °C and +70 °C, consistent with commercial-grade electronics, and uses environmentally resistant materials such as ABS or PETG plastic for the enclosure. All exposed interfaces are sealed with gaskets or silicone to prevent dust ingress, and sensor ports are protected by mesh vents to allow airflow while maintaining particulate resistance. The lithium-ion battery includes overcurrent and overtemperature protection to prevent degradation in hot environments. Moisture-sensitive components, including the barometric and radar modules, are coated with a thin conformal layer to resist condensation. While the device is not waterproof, it is designed for short-term outdoor exposure, ensuring dependable performance in unpredictable field conditions without compromising portability or safety.

4.1.12. Computational Constraints

The 2HAD-RAD must execute multiple concurrent tasks, such as, radar signal processing, sensor fusion, display management, and wireless communication, on a limited embedded platform. The ESP32-WROOM, with a dual-core 240 MHz processor and 520 KB of SRAM, serves as the central controller but provides only moderate computational resources compared to desktop or FPGA-based systems. To operate within these limits, data processing is optimized through time-sliced task scheduling, lightweight algorithms, and efficient use of DMA channels for SPI and UART communication. Radar preprocessing is largely offloaded to the IWR6843AOP's on-chip DSP, reducing CPU burden. Nonessential tasks such as display updates or BLE communication are throttled during radar operation to preserve real-time responsiveness. These architectural and firmware-level optimizations allow the system to maintain real-time performance within its finite memory and processing bandwidth, demonstrating efficient embedded computation under constrained resources.

4.1.13. Ethical / Safety Constraints

The 2HAD-RAD design incorporates safety considerations consistent with IEEE and FCC standards to ensure that all transmitted electromagnetic energy remains within non-hazardous exposure limits. Operating in unlicensed frequency bands, the radar and UWB modules emit at power levels well below thresholds defined for human exposure to electromagnetic fields. From an ethical standpoint, the device is intended solely for life detection and rescue applications, avoiding any intrusive or privacy-violating use of radar sensing on individuals. All sensing data are processed locally, with no cloud transmission, preserving confidentiality and aligning with responsible data handling practices. Electrical safety is addressed through reverse-polarity protection, current limiting, and fused battery inputs to prevent overheating or short circuits. By adhering to recognized electrical and ethical standards, the 2HAD-RAD prioritizes human safety, responsible use of technology, and compliance with societal expectations for ethical engineering conduct.

CHAPTER 5

5.1. Overview

Since its initial release in November 2022, ChatGPT has evolved into one of the most widely adopted AI assistants for academic and professional applications. With the development of GPT-5, the model introduced enhanced reasoning, contextual understanding, and technical depth, making it a powerful companion for engineering research. Going into the early stages of Senior Design I, I anticipated that ChatGPT would become one of my primary tools for researching hardware components, comparing specifications, and assisting in documentation. While it certainly proved helpful, my experience revealed both its strengths and limitations in technical design contexts. When using ChatGPT for research, I typically began my prompts by clearly outlining my goals and expected outputs, such as tables comparing component specifications, summaries of datasheets, or explanations of circuit behavior.

Across our team, every member approached AI assistance differently. Some of us relied on it for literature review and formatting, while others focused on code optimization, simulation setup, or part selection. To better understand how these models performed under similar conditions, we compared ChatGPT with Claude, developed by Anthropic, and Perplexity AI, an emerging research focused assistant. Claude excelled in long form reasoning and clarity, often producing detailed technical explanations ideal for report writing. Perplexity AI, on the other hand, was notably effective for real-time data retrieval and source verification, making it valuable for up-to-date product comparisons and component availability checks. By analyzing their outputs using identical prompts, our team evaluated each model's accuracy, organization, and relevance. The following sections summarize how each member incorporated these AI tools into their workflow and how the insights gained from them contributed to the efficiency and technical depth of our overall project.

5.2. Advantages and Limitations

5.2.1. Advantages

ChatGPT offers several advantages for our team. It helps quickly summarize academic papers and datasheets on radar, infrared, and inertial sensors, providing a solid foundation for component selection and system design. Additionally, ChatGPT streamlines drafting the abstract, introduction, and system descriptions by paraphrasing and formalizing text while maintaining technical accuracy. It also improves the quality of reports by identifying unclear phrasing, redundant sections, and inconsistencies, enhancing both readability and professionalism. During the early design phases, ChatGPT serves as a technical tutor by offering simplified explanations of complex concepts like Doppler radar theory, UWB pulse generation, and signal conditioning, which reduces the time spent cross-referencing multiple textbooks.

5.2.2. Limitations

While ChatGPT proved valuable throughout the research and design process, several limitations were identified during its use. The most prominent issue involved fact verification, as the model occasionally returned inaccurate technical specifications such as component power ratings, MCU clock speeds, or frequency ranges. These discrepancies required consistent cross-checking against official datasheets to ensure design accuracy. Additionally, the model lacks real-time data or simulation capability, meaning that while it can explain expected circuit behavior conceptually, empirical validation still had to be performed manually in LTspice and MATLAB. The quality of responses also depended heavily on prompting style, for example, detailed questions often yielded highly relevant answers, whereas vague or generalized prompts resulted in off topic or incomplete information. Another noted drawback was the risk of overreliance, as its convenience for writing and summarization could discourage deeper analytical reasoning if not balanced with sound engineering judgment. Finally, the use of AI-generated text introduced copyright and authorship concerns, since large language models draw from broad data sources that may not always be properly attributed. As a result, any material obtained from ChatGPT was treated as an assistive reference rather than an original or citable source, ensuring academic integrity throughout the project.

5.3. Case Study 1: Power Management Optimization

Please refer to Appendix E, [1a], [1b], and [1c] for the Case Study 1 given prompt and responses from ChatGPT and Claude.

This case study was designed to test both models' abilities to generate technical recommendations for optimizing power efficiency in an embedded system using dual-rail Li-ion regulation. ChatGPT provided a concise and practical analysis, emphasizing synchronous buck-boost converters, low quiescent current operation, and firmware strategies such as duty cycling and deep-sleep scheduling. Its response demonstrated strong alignment with the power management goals of the project, offering clear design-level insights that could be directly implemented.

Claude, in contrast, produced a broader systems-level response, prioritizing long-term energy balance, dynamic voltage scaling, and intelligent radar activation based on motion triggers. While both models arrived at similar conclusions regarding converter efficiency and power gating, Claude's answer read more like a white paper, highlighting lifecycle monitoring and hardware-level optimization. ChatGPT's structured, implementation-driven style was more immediately actionable for our design phase, whereas Claude's added context for future optimization potential.

5.4. Case Study 2: Signal Filtering in Noisy Environments

Please refer to Appendix E, [2a], [2b], and [2c] for the Case Study 2 prompt and responses from ChatGPT and Claude.

This question evaluated the models' understanding of digital signal processing within the context of a radar-based life-detection system. ChatGPT's response was grounded in classical filter design, suggesting Butterworth and Chebyshev filters for noise suppression, as well as Kalman filtering and FFT-based noise analysis for dynamic signal isolation. Its focus on the respiration frequency band (0.1–0.4 Hz) aligned well with real physiological data, making it a strong technical reference for our filtering implementation.

Claude's response was far more detailed and research-oriented, offering a layered approach that included clutter removal, adaptive background modeling, range gating, and the incorporation of machine learning classifiers to distinguish genuine human signatures from environmental noise. While ChatGPT excelled in explaining practical filter architectures, Claude provided more advanced and multidisciplinary techniques that could be applied in future versions of the system. Together, they reinforced both foundational and modern approaches to radar signal processing.

5.5. Case Study 3: Communication Protocol Selection

Please refer to Appendix E, [3a], [3b], and [3c] for the Case Study 3 prompt and responses from ChatGPT and Claude.

This prompt aimed to compare each model's reasoning in selecting an optimal communication interface between the radar baseboard and the ESP32 microcontroller. ChatGPT identified SPI as the most balanced solution for speed, reliability, and ease of wiring, noting its superior data rate and deterministic timing compared to I^2C. It also suggested UART as a secondary channel for debugging and low-speed control tasks.

Claude supported this recommendation but elaborated further on potential scalability, suggesting that CAN bus or RS-485 could be explored in later iterations for environments requiring longer cable runs or improved EMI resistance. Its focus on redundancy and flexibility highlighted a more modular communication approach. Both models ultimately aligned on SPI as the best option for the prototype stage, validating the design direction taken by the team.

5.6. Case Study 4: Enclosure Material Selection

Please refer to Appendix E, [4a], [4b], and [4c] for the Case Study 4 prompt and responses from ChatGPT and Perplexity.

The fourth case study assessed each model's material science reasoning as it relates to radar transparency and environmental durability. ChatGPT recommended polycarbonate-ABS (PC-ABS) as the optimal enclosure material, emphasizing its low dielectric constant, high impact strength, and ease of machining. It also considered IP-rated sealing and conformal coating for additional protection, aligning closely with real-world mechanical design principles.

Perplexity offered a more data-driven response, comparing several dielectric-friendly plastics and citing radar attenuation levels in decibels. It recommended ABS, PETG, and fiberglass-reinforced nylon as viable materials depending on mechanical and environmental demands. Its explanation of composite-layer shielding and UV stability provided additional depth to the discussion. In this case, both AI models effectively validated the selected PC-ABS material, reinforcing confidence in its balance between RF transparency, strength, and weight.

5.7. Case Study 5: Data Visualization on Embedded Displays

Please refer to Appendix E, [5a], [5b], and [5c] for the Case Study 5 prompt and responses from ChatGPT and Claude.

This case study explored how each model could propose visualization strategies for displaying radar-based motion and respiration data on a compact 4.3-inch TFT screen. ChatGPT's output was practical and implementation-oriented, proposing waveform and bar graph combinations, adaptive brightness control, and color-coded motion indicators. Its suggestions closely matched embedded GUI design best practices and emphasized readability and responsiveness under real-time constraints.

Claude provided a more human-centered response, describing a heat-map style interface with trend graphs, signal confidence indicators, and color-coded alerts to support quick interpretation by field users. It also discussed optimizing refresh rates to match radar update cycles and minimizing redraws to conserve power. Collectively, both responses demonstrated strong design intuition, with ChatGPT emphasizing functionality and Claude prioritizing usability and situational awareness.

CHAPTER 6

6.1. Hardware Block Diagram

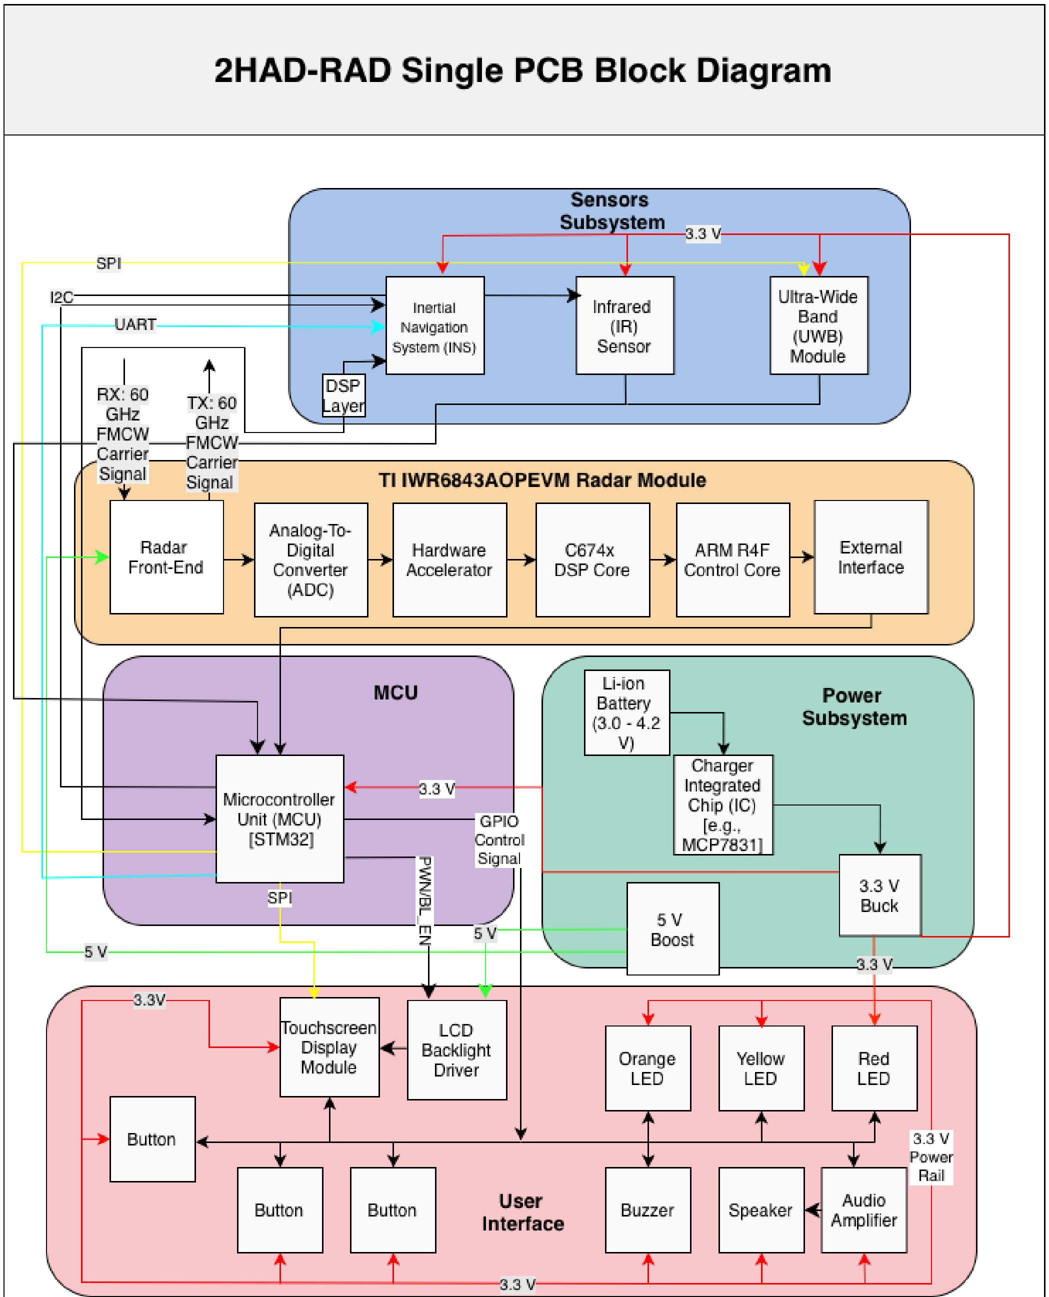

Division of Labor

Subsystem	Primary Lead	Secondary Support
Power Subsystem	Bailey Hitt (EE)	Shannon Hankinson (EE)
Sensors Subsystem	Ryan Andersen (EE)	Bailey Hitt (EE)
Radar Module	Shannon Hankinson (EE)	Ryan Andersen (EE)
MCU	Joshua D'Souza (CpE)	Bailey Hitt (EE)
User Interface	Joshua D'Souza (CpE)	Ryan Andersen (EE)
Radio	Shannon Hankinson (EE)	Ryan Andersen (EE)

Figure 6: Hardware Block Diagram

6.2. Sensor Subsystem

In figures 7-10, the sensor subsystem comprises three primary sensor modules: the M5Stack M135 inertial navigation module, the Qorvo DWM3000 UWB transceiver, and the Adafruit MLX90640 infrared (IR) array sensor. Each module directly interfaces with the ESP32-WROOM-32E microcontroller, which acts as the central processing and communication unit of the system. The sensor boards are mounted on a single PCB with power and signal routing designed to minimize noise, maintain consistent voltage regulation, and ensure synchronization across all data lines.

6.2.1. Radar Module

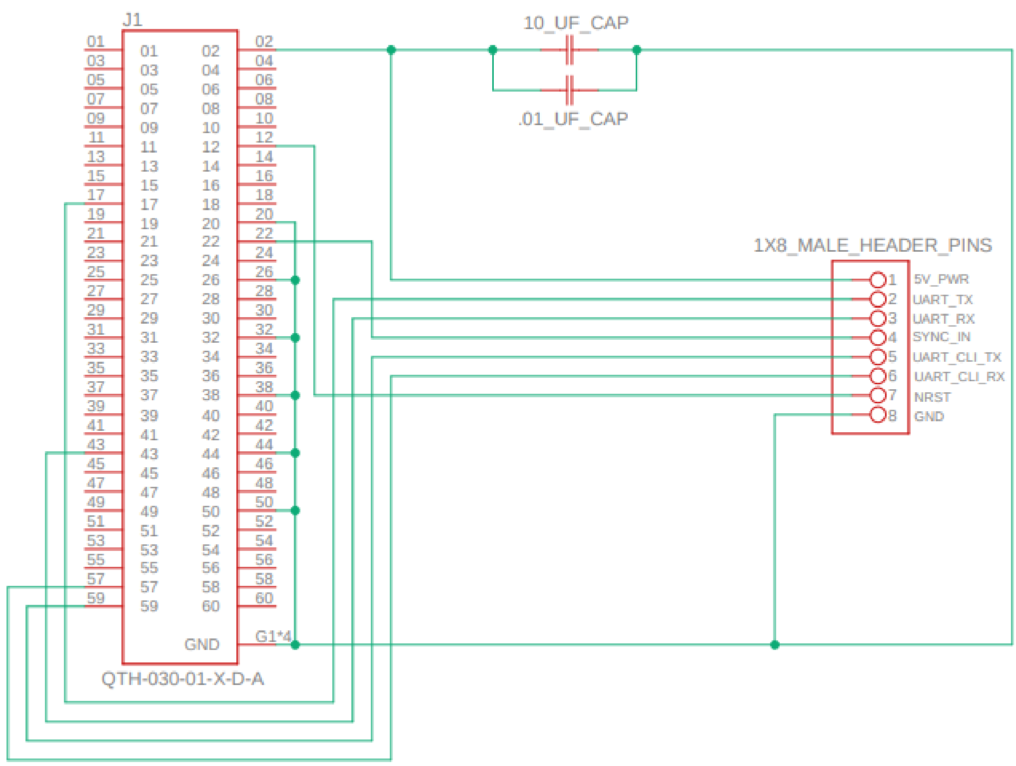

Figure 7: Radar Module PCB Schematic.

6.2.2. INS Module

The M5Stack M135 module integrates a u-blox NEO-M9N GNSS receiver, a Bosch BMI270 IMU, a BMM150 magnetometer, and a BMP280 barometer, providing a complete 9 axis inertial navigation solution. The module communicates with the ESP32 over the I2C bus for sensor data (SDA on GPIO21, SCL on GPIO22) and uses the UART for GNSS data streaming. A PPS (Pulse-Per-Second) line from the M135 is optionally routed for precise timing synchronization between GPS and UWB measurements.

Series resistors are included on all data lines to limit ringing and crosstalk. Local 10 µF and 0.1 µF decoupling capacitors are placed adjacent to the M135's 3.3 V power input to maintain signal integrity.

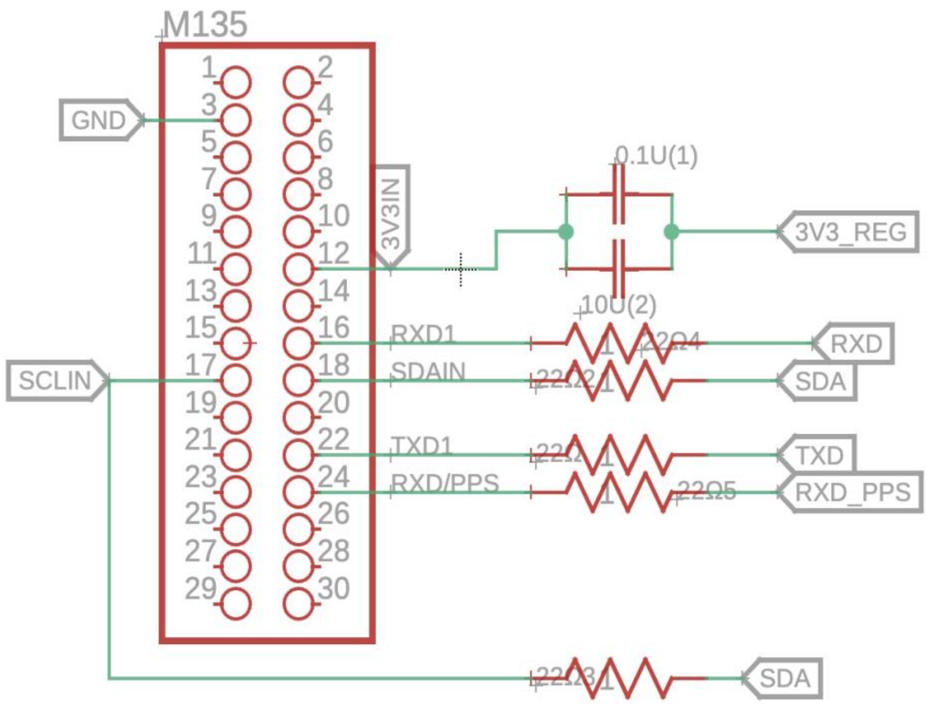

Figure 8: M135 module schematic

6.2.3. IR Sensor

The Adafruit MLX90640 breakout provides thermal imaging capability via the same I2C bus used by the M135. The sensor's SCL and SDA lines are connected in parallel with the INS, while the pull-up resistors located on the MCU maintain proper bus biasing. Power is provided from the 3.3 V rail with local decoupling. This sensor provides a thermal grid, which will be processed alongside radar and UWB data for motion and human presence detection.

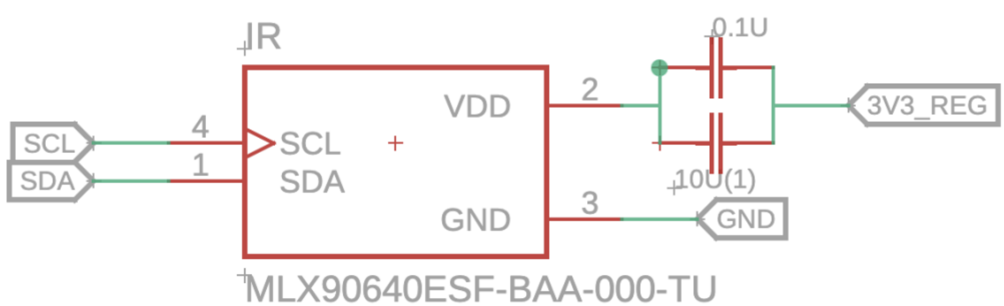

Figure 9: IR sensor schematic.

6.2.4. UWB Sensor

The Qorvo DWM3000 module operates using the SPI interface to communicate with the ESP32. Control signals include IRQ and RST with an optional wake up line. Each SPI line incorporates a small series resistor to reduce reflection along high-speed traces, and local bypass capacitors are used to suppress transient voltage spikes near the UWB module. The DWM3000 enables precision ranging and time-of-flight measurements that will be fused with the INS and IR sensor data to detect and locate victims. Due to the susceptibility of UWB communication to noise, the DWM3000 and MCU share a clean, isolated 3.3 V rail, which is filtered through a network of bulk capacitance.

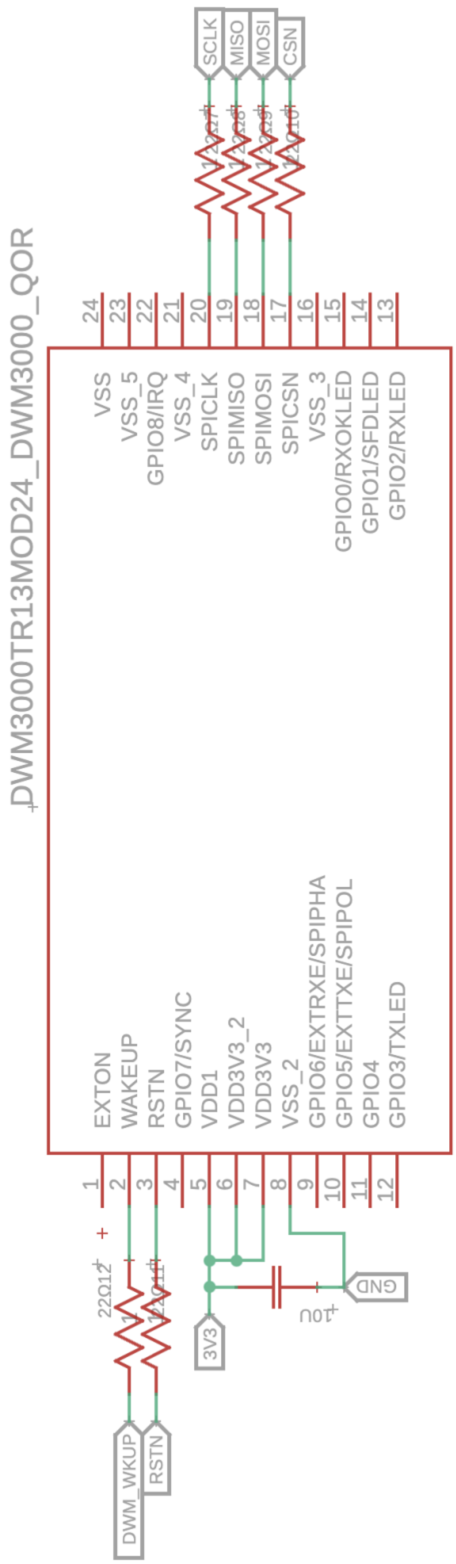

Figure 10: UWB sensor schematic

6.3. MCU Configuration

The ESP32-WROOM-32E module in figure 11 is powered through a dedicated 3.3 V linear regulated supply provided by the LM1117 regulator. The regulator accepts input a 3.7 V Li-ion battery through a charging module. Decoupling capacitors of 10 µF and 0.1 µF are placed close to the ESP32's power pins to stabilize the voltage rail and suppress transient switching noise. Communication with the microcontroller is facilitated via UART, SPI, and I2C buses, which were selected according to the data rate and protocol requirements of each sensor. The FT231XS USB-to-UART converter is included for programming and debugging during development, while the system can operate fully on battery power in field deployment.

To protect sensitive components from voltage spikes, transient surges, and static discharge, several protection devices are integrated throughout the design. The USB data lines (USBD+ and USBD–) include a bidirectional transient voltage suppression (TVS) diode array (CDSOT23-SLVU2.8) placed close to the micro-USB connector. This device provides ESD protection up to +/-15 kV and clamps any high-energy transients that could otherwise damage the FT231XS USB-to-UART converter or propagate into the 3.3 V logic domain. Each data line is routed through the diode to ground, ensuring that surges are safely dissipated without affecting signal integrity.

To protect signal integrity between the USB-UART interface and the microcontroller, small-signal diodes are placed on the DTR and RTS control lines. These diodes isolate the auto-programming network, ensuring that toggling one line does not inject noise or reverse current into the other or into the ESP32's enable (EN) and boot (IO0) pins. In addition to providing logic isolation, the diodes act as transient suppressors for voltage spikes generated during USB connection events or rapid serial toggling. No series protection diodes are used on the main 5 V rail since the charging module internally handles source selection and reverse-current protection, and the LM1117 LDO regulator provide its own internal reverse bias safeguards. The LM1117 LDO regulator does include local decoupling capacitors in our design that act as additional low pass filters, reducing conducted noise and preventing supply transients from coupling into sensitive sensor lines. In addition, series resistors are used on high-speed communication lines such as SPI and UART to damp ringing and limit inrush current during logic transitions. This protects both the ESP32 and peripheral modules from overshoot and improves signal quality at higher data rates. Together, these protection and conditioning elements improve the electrical robustness of the PCB, extend component lifespan, and ensure stable operation under varying load and environmental conditions.

6.3.1. ESP32-WROOM-32E

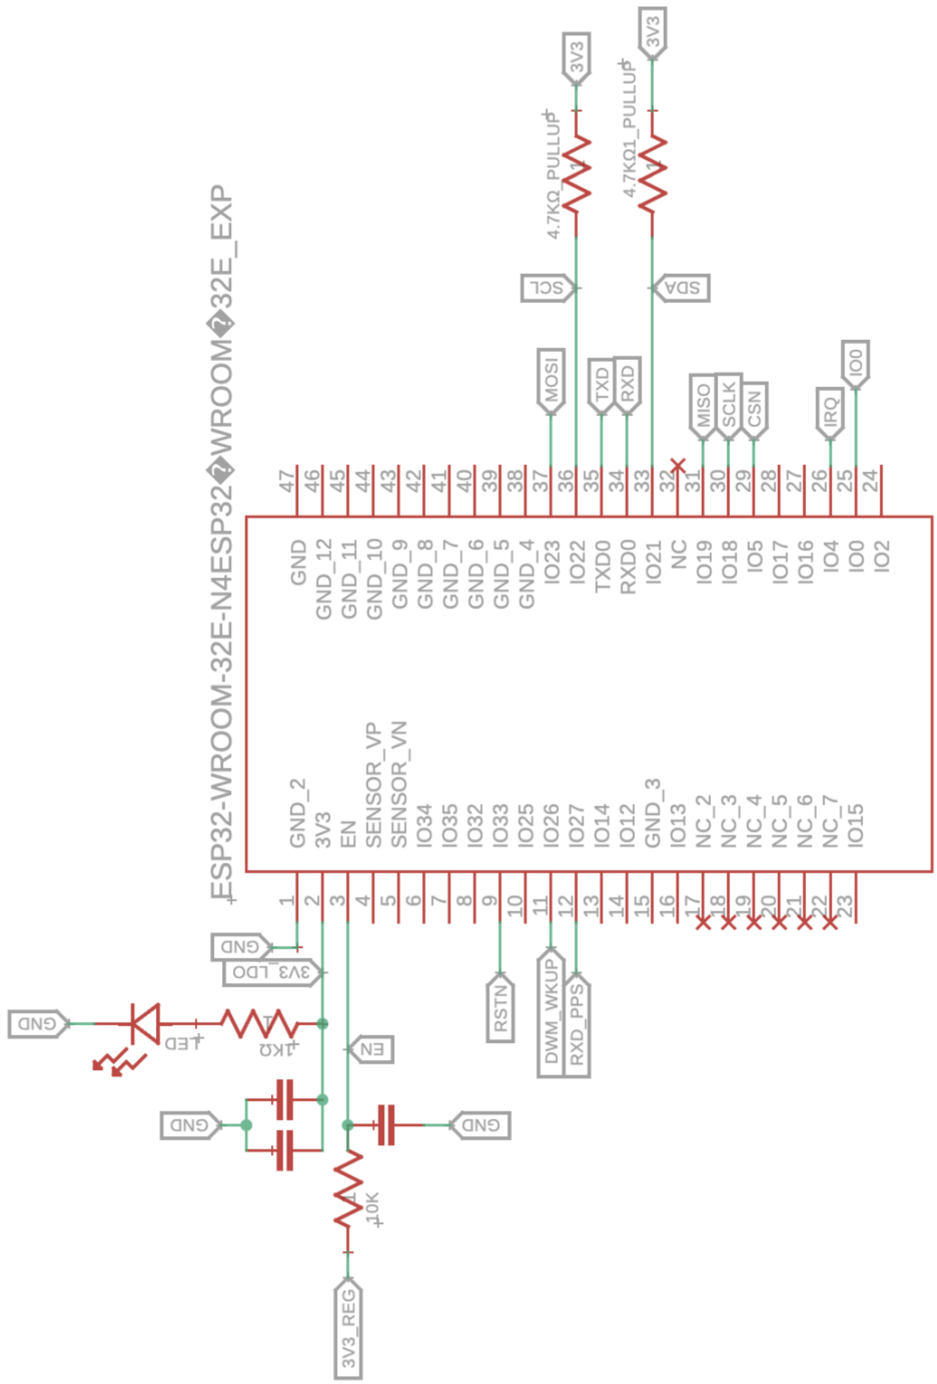

Figure 11: ESP32-WROOM-32E schematic

6.3.2. USB to UART

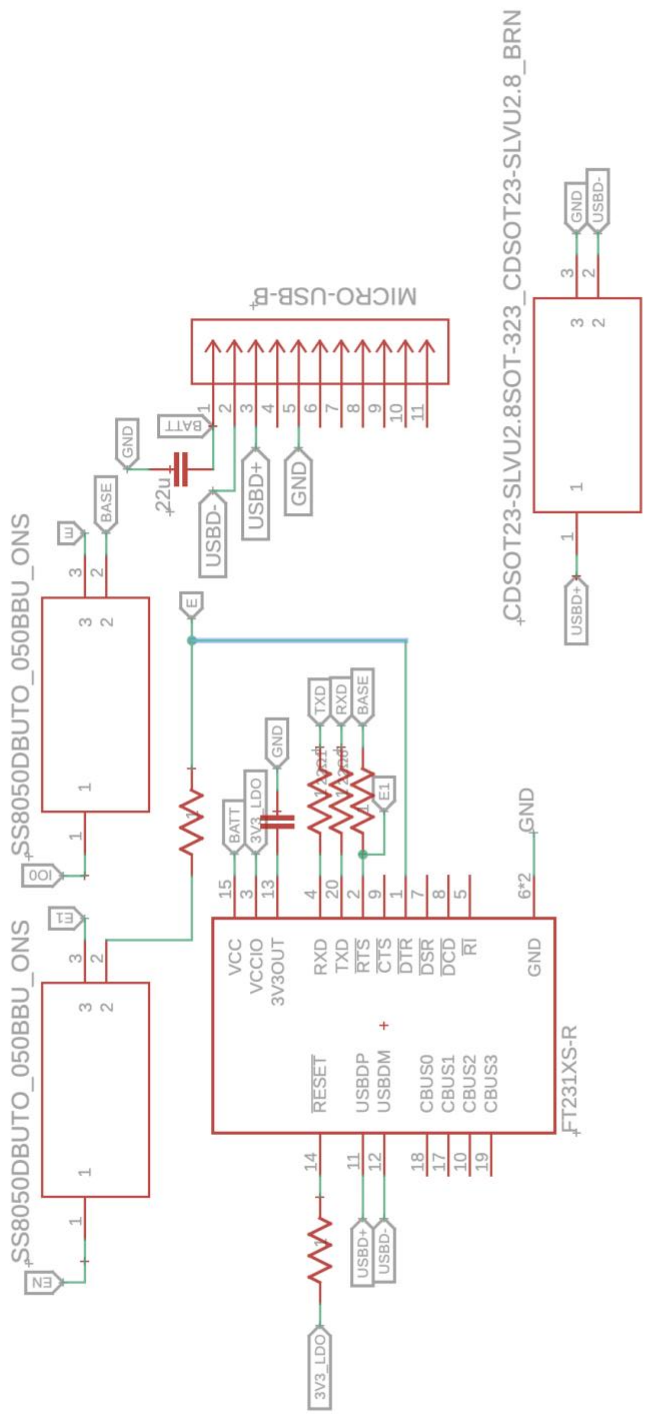

Figure 12: USB to UART schematic.

6.3.3. LDO Regulator

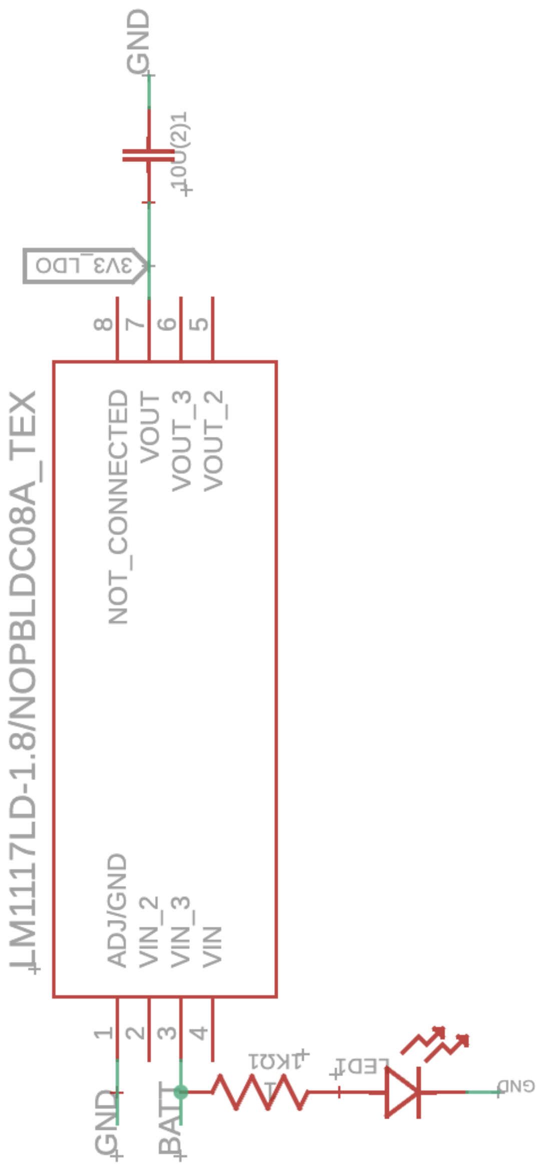

Figure 13: LDO regulator schematic for MCU.

6.4. User Interface

The user-interface elements shown in Figure 14 integrate three LED status indicators and a single digital pushbutton input directly with the ESP32-WROOM-

32E module to provide operator feedback and system control. Each LED is driven from an individual ESP32 GPIO pin using a dedicated series resistor sized for a nominal 2–5 mA indicator current. The LEDs are powered from the regulated 3.3 V rail to ensure GPIO compatibility and prevent back-feeding into higher-voltage domains. Placing each resistor in series with its corresponding LED limits inrush current during switching transitions, reduces stress on the microcontroller pins, and ensures consistent brightness regardless of supply variation. The pushbutton interface uses a momentary, normally open SPST switch that connects the selected GPIO pin to ground when actuated. This pin is configured with the ESP32's internal pull up resistor so the input remains in a stable logic-high state when the button is not pressed. To improve signal integrity and suppress contact bounce, a 100 Ω series resistor and a 0.1 μF shunt capacitor are included at the switch node, forming a compact RC debounce filter placed close to the connector footprint. This filter prevents false triggering and minimizes the effect of mechanical chatter during transitions.

Additional protection features are included to ensure reliable operation of the user interface under electrically noisy conditions. Each LED control line incorporates the series resistor as both a current-limiting and transient-damping element, restricting surge currents that may be generated during rapid GPIO toggling. The pushbutton input is protected from ESD and cable-injected noise through its RC filter and through the ESP32's internal ESD structures, which remain within safe limits due to the controlled impedance of the interface. Since the pushbutton is panel-mounted and connected to the main PCB via a two-wire harness, the inclusion of the series resistor at the microcontroller side prevents long cable inductances from coupling fast edges back into the GPIO input path. No external diode protection is required for the LEDs or pushbutton interface, as neither element interfaces with external high-voltage sources, and all signals remain within the 3.3 V logic domain managed by the ESP32. Collectively, these design choices provide a clean, noise-tolerant, and electrically robust user-interface subsystem that maintains predictable behavior across varying operating conditions and ensures long-term reliability in both development and field environments.

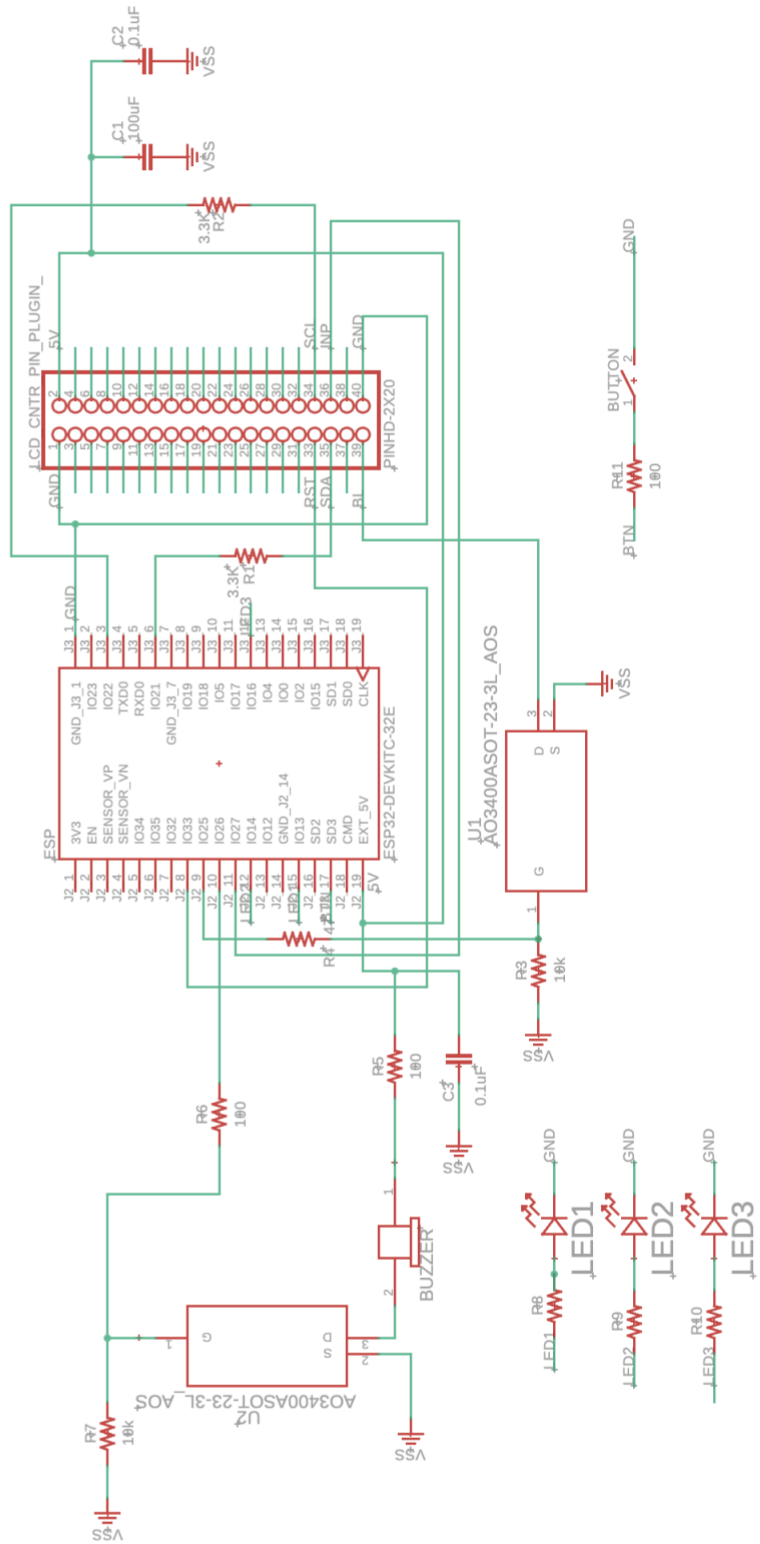

Figure 14: Schematic for User Interface: Buttons, Indicators, and Display Connections

6.5. Power Subsystem

The schematic in Figure 15 implements the complete power-management subsystem for the 2HAD-RAD handheld life-detection device. The circuit allows the system to operate from either a USB Type-C input or a single-cell Li-ion battery, with seamless transition between external and battery power. The BQ24074RGTR charge-management IC regulates power flow from the USB-C input, protected by a polyfuse and local decoupling (0.1 µF / 10 µF), into the Li-ion cell connected through the JST-PH2 jack. Its internal power-path management supports load sharing, maintaining system operation while charging the battery. The configuration network (R3–R7) defines the fast-charge current, termination thresholds, and safety limits, while the BAT and OUT pins are bypassed with low-ESR capacitors to minimize transients during connection or USB insertion events.

Downstream, two high-efficiency DC-DC controllers, LM5176 and LM5155, generate the regulated rails required by the ESP32-WROOM microcontroller, radar baseboard, and peripheral sensors. The LM5176 operates in a synchronous buck-boost configuration, maintaining a stable 5 V system bus across the full battery range. Its timing, slope-compensation, and feedback networks (R18–R22, C7–C13) establish proper current-mode stability, while bulk capacitors at VIN and VOUT ensure low-noise operation. The LM5155, located in the auxiliary section, forms a dedicated boost converter for higher-voltage peripherals when required. Its power stage includes an external N-MOSFET, a 0.05 Ω sense resistor, and a 68 µF output capacitor, with all control pins (RT, COMP, SS, FB, CS) terminated through tuned RC networks to produce a realistic, stable topology.

All analog, digital, and power grounds are unified under a single GND net for consistent reference potential and minimal ground-bounce. The subsystem includes appropriate filtering, compensation, and protection components at every interface, forming a robust and self-contained power architecture capable of USB-powered operation, intelligent battery charging, and multi-rail regulation for the 2HAD-RAD embedded platform.

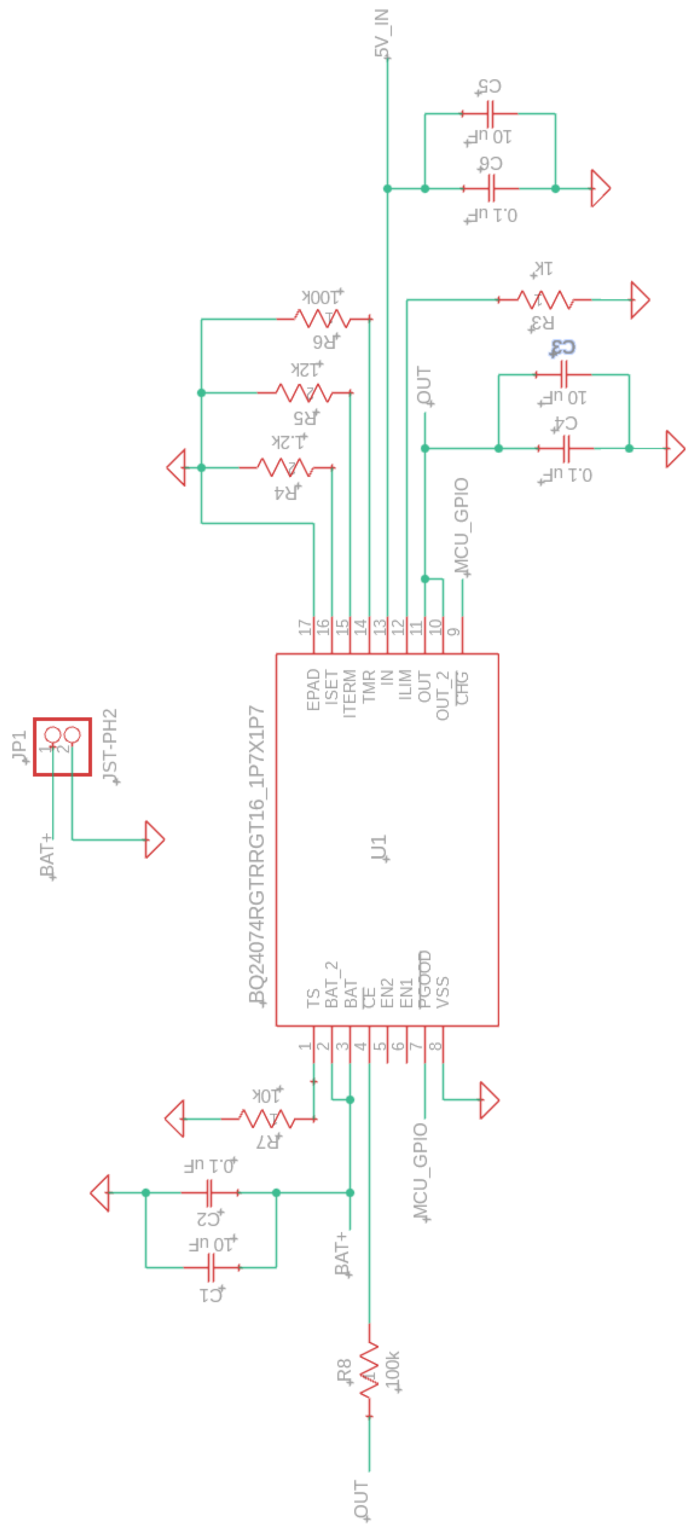

Figure 15: Charger IC and battery cable connector

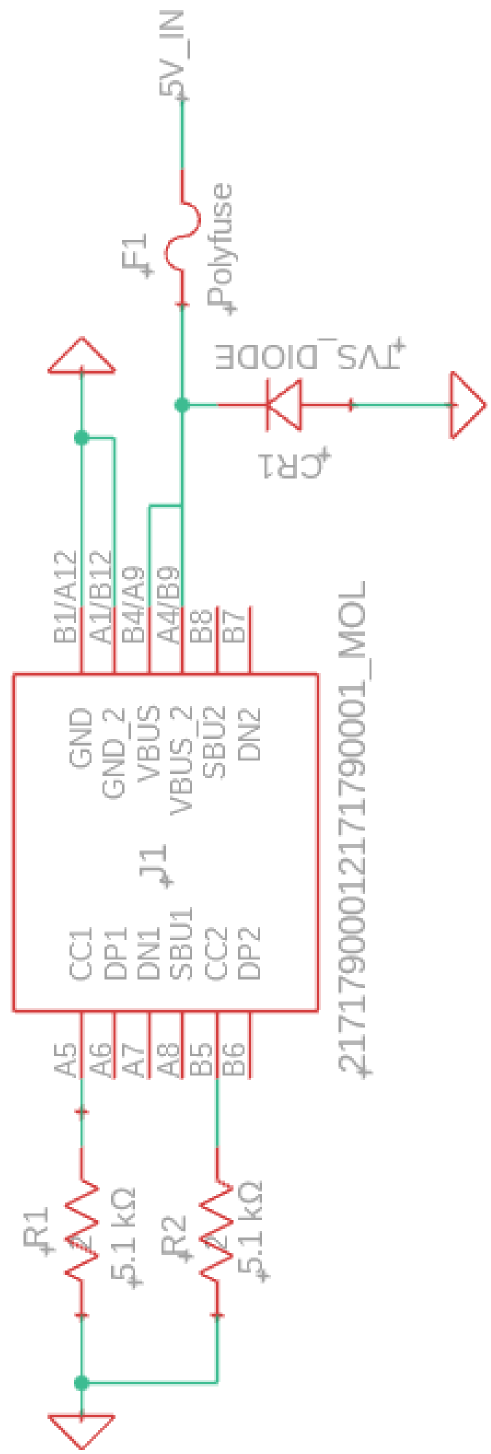

Figure 16: USB-C power input

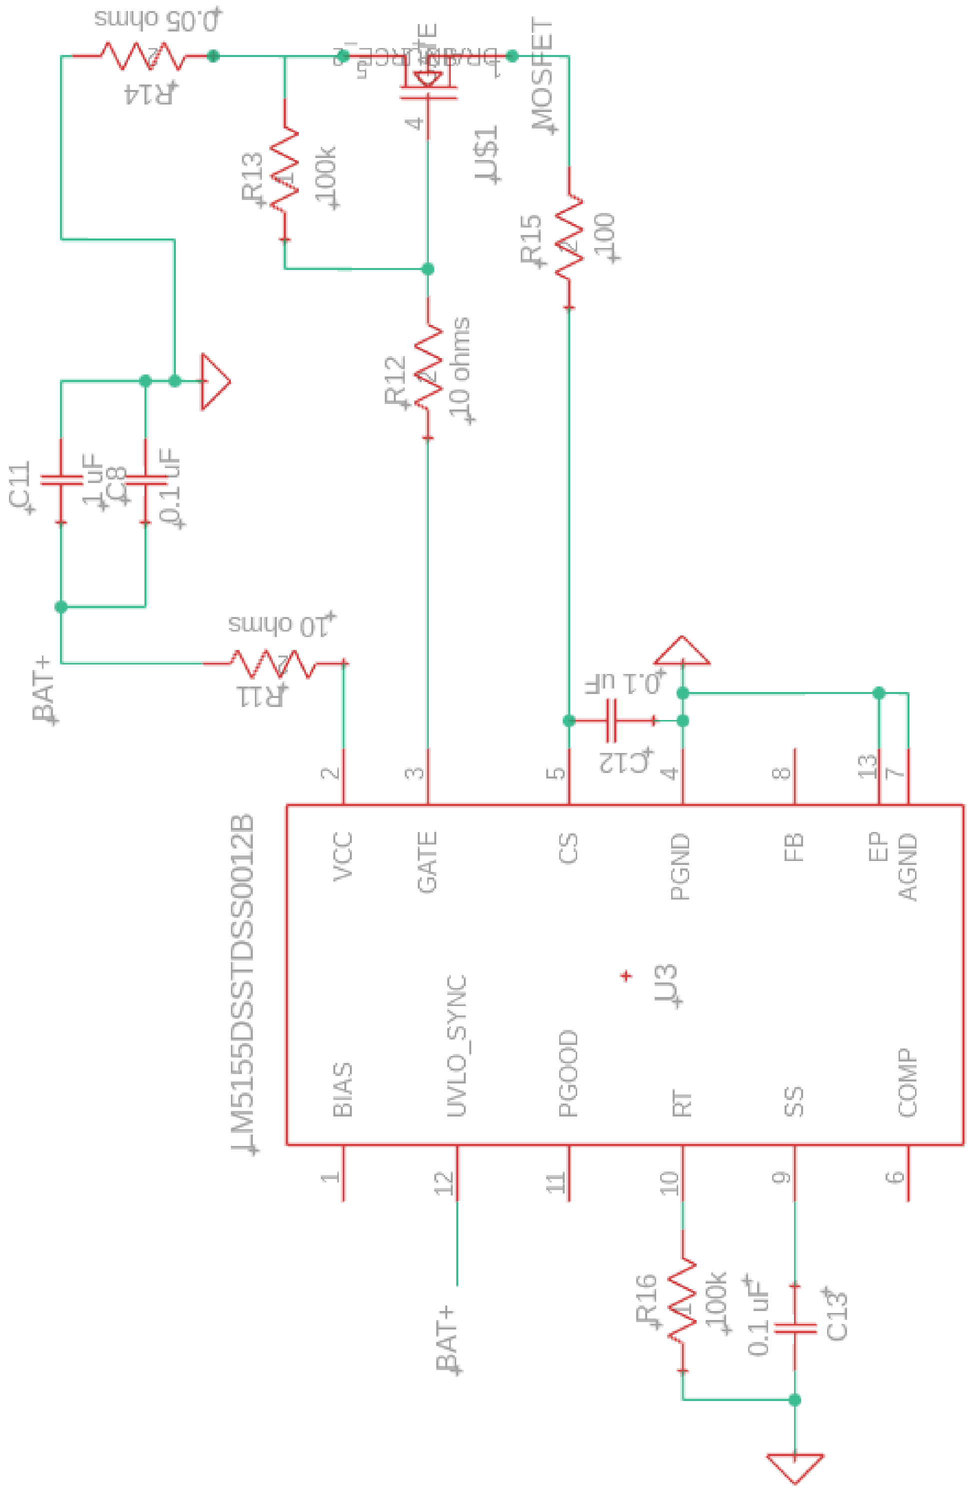

Figure 17: LM5155 Boost Regulator

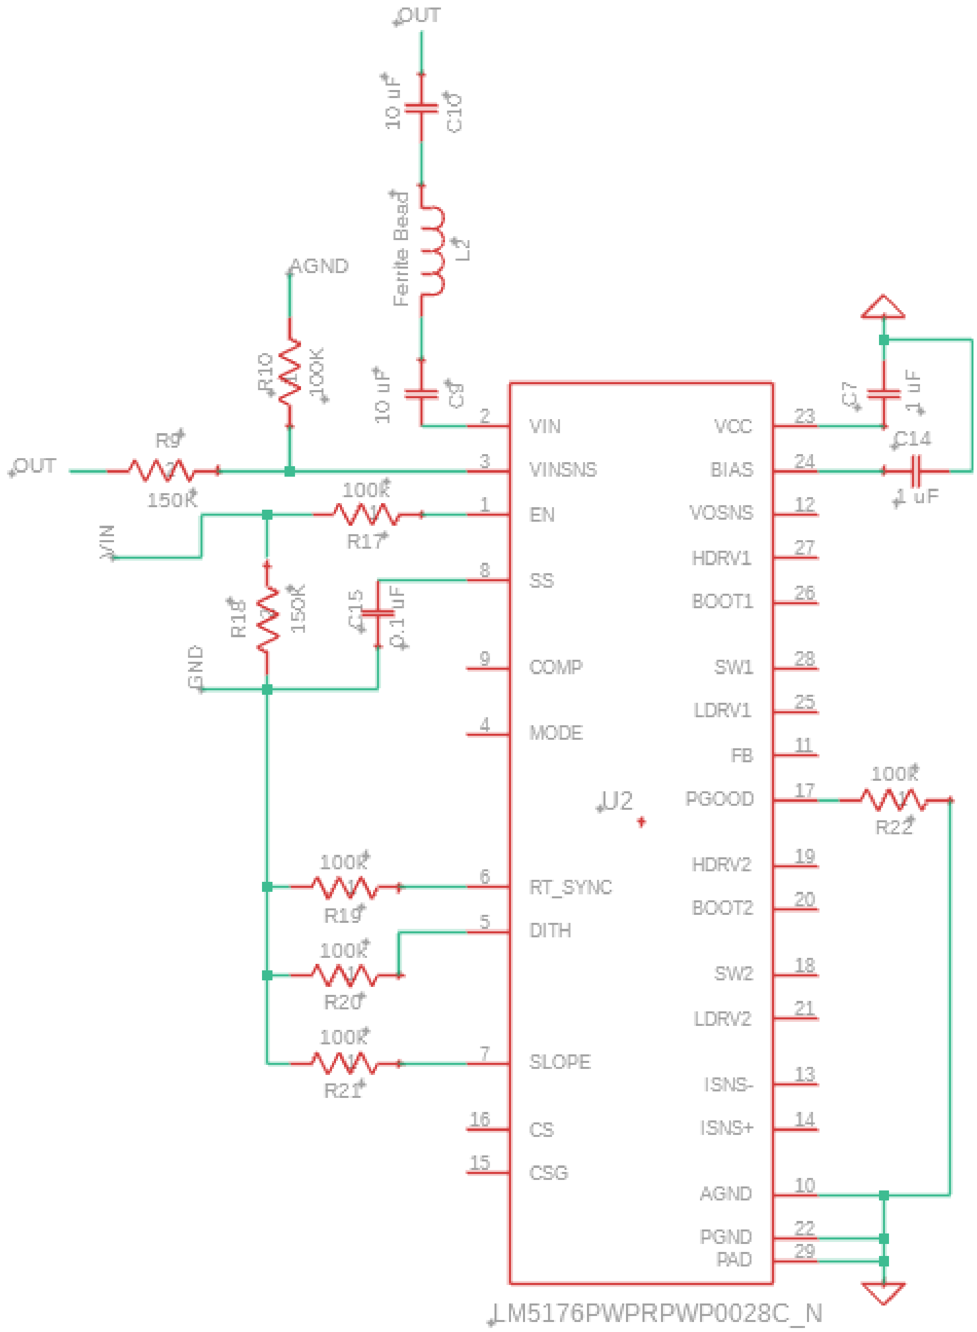

Figure 18: LM5176 Buck-Boost Regulator

CHAPTER 7

This project involves the creation of two major software components. One is the base software uploaded on the central ESP32 which handles data from the various sensors and displays it onto the device screen. The second, which is an advanced goal, is the smartphone application that communicates with the MCU through Bluetooth low energy (BLE) to display data to the user.

7.1. ESP32 Software Design

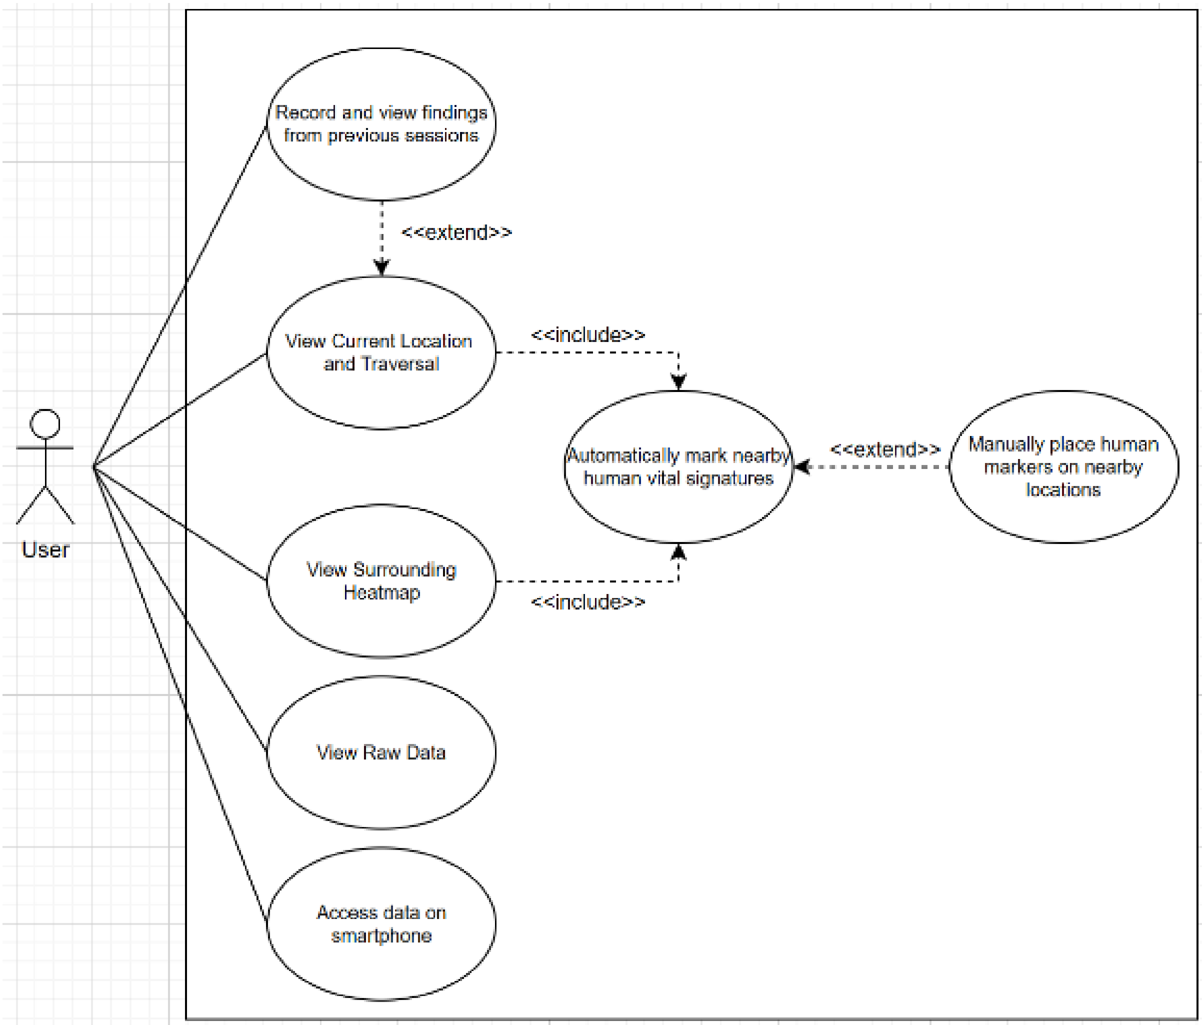

Figure 19: Use Case Diagram

The ESP32 receives signals from the Radar (TI IWR6843AOPEVM), Ultra-wideband (Qorvo DWM3000EVB), Infrared sensor (Adafruit MLX90640), and Inertial Navigation System (M5Stack M135). The ESP32 configures radar parameters so that signals from the radar can be used to detect human frequencies around the device, such as breathing and heartbeat. Signals from the ultra-wideband are used to detect human presence in ways the radar may not be able to, such as through walls. The infrared sensor sends temperature data to the ESP32 for an additional layer of human detection. The accelerometer, gyroscope, magnetometer, and barometric sensor contained in the inertial navigation system provide data to the ESP32 for location tracking of the device. An interactive map presents the data from the sensors, displaying the device’s position and path. It also includes markers indicating potential human activity and the confidence level of whether a human is present. This is presented through the user interface using the 4.3-inch EastRising RA8875 TFT LCD display. A

CMI-1295-03TH piezoelectric buzzer compliments the device notification system, providing noises for the user to hear (especially useful in harsh environments).

The device houses pushbuttons corresponding to distinct control functions, such as turning the device on/off, initiating a scanning session, toggling display modes, and acknowledging alerts. A scanning session is a period in which the interactive map places the device location as its start point and tracks the movement of the device. During this period, the sensors are continuously recording data, and a path is drawn on the map to depict the user's traversal over time. Display modes refer to different views that the user interface provides, discussed in subsequent sections in this chapter. Alerts on the display screen are used to notify the user that a human has been detected, an error has occurred, or when battery gets low. The device houses 3 LEDs controlled using interrupts and delays that leverage the ESP32 internal clock. A blue LED blinks slowly while a scan session is underway. A green LED blinks rapidly for 15 seconds when a new human has been detected with a sufficient confidence score. A red LED turns on when the device reaches a low battery level. The device houses a CMI-1295-03TH piezoelectric buzzer generates short alert tones for immediate notifications for when a scanning session begins and ends, when a human is detected, when user makes an error with the UI, and when battery gets low.

7.1.1. User Interface

The user interface is organized into three main display modes: raw data view, a coordinate map, and a heat map. You can toggle between these display modes using either of the pushbuttons or the touchscreen. Regardless of display mode, device settings can be accessed through a button with a cog icon. Settings contain the initial date and time setup which leverages the ESP32's internal RTC. It also provides controls for adjusting the blinking speed of LEDs, the speed of disappearance of low-confidence markers, the refresh rate of data, the representation of distances in miles or kilometers, and other similar settings.

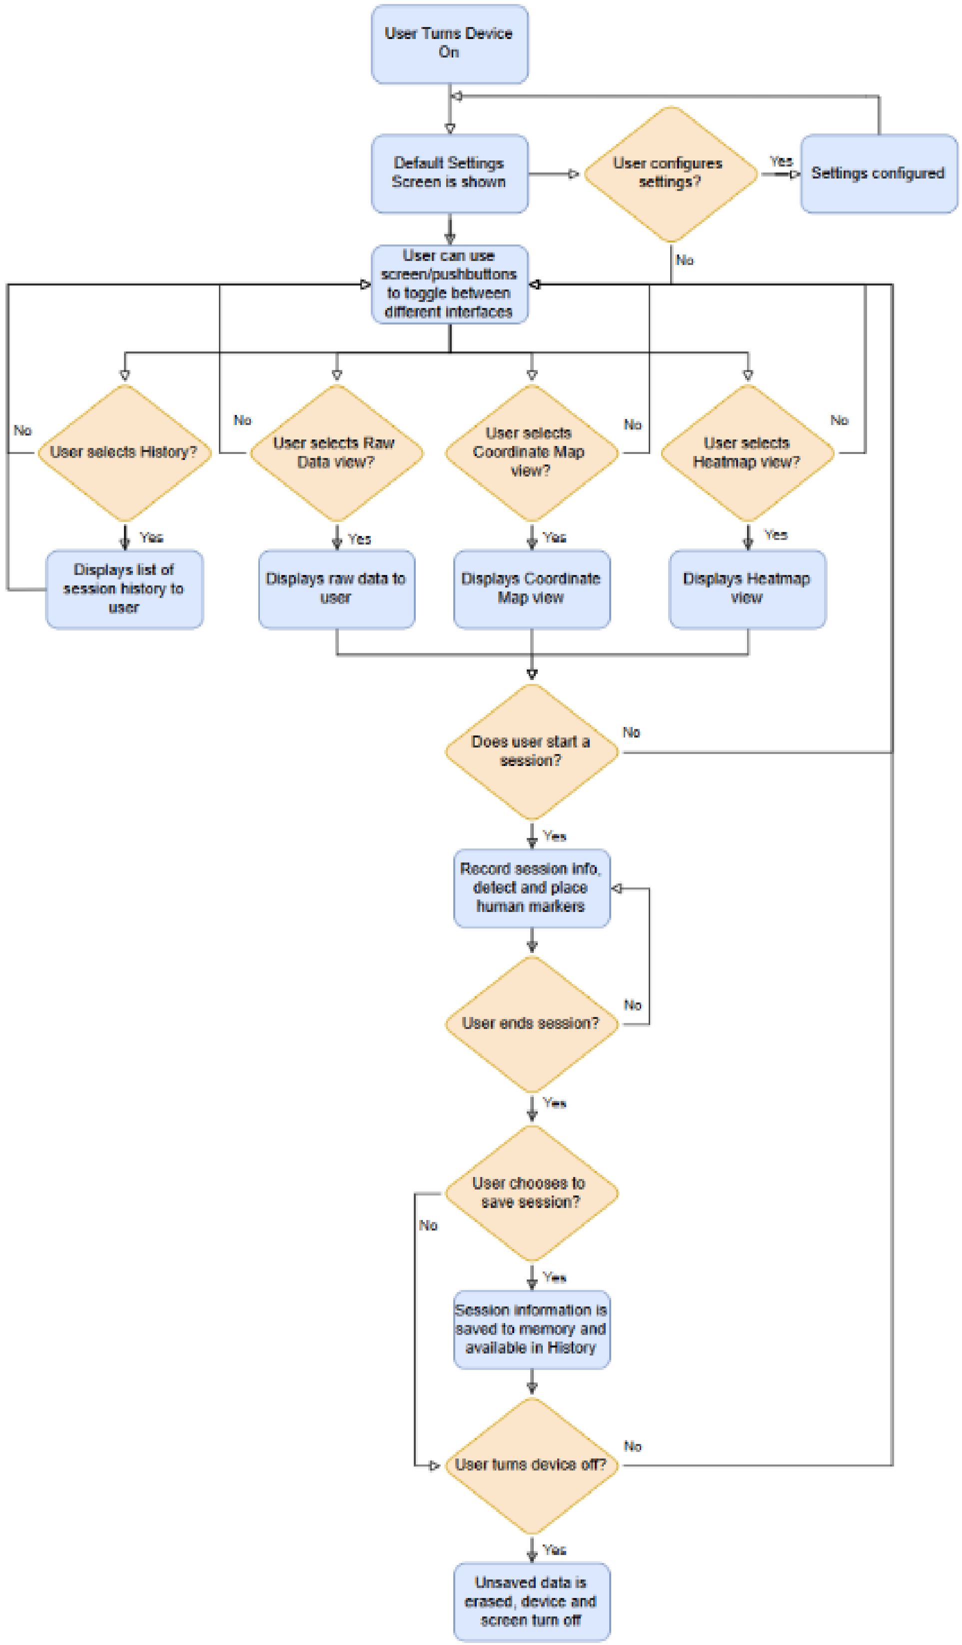

Figure 20: Overall User Interface Flowchart

7.1.2. Raw Data View

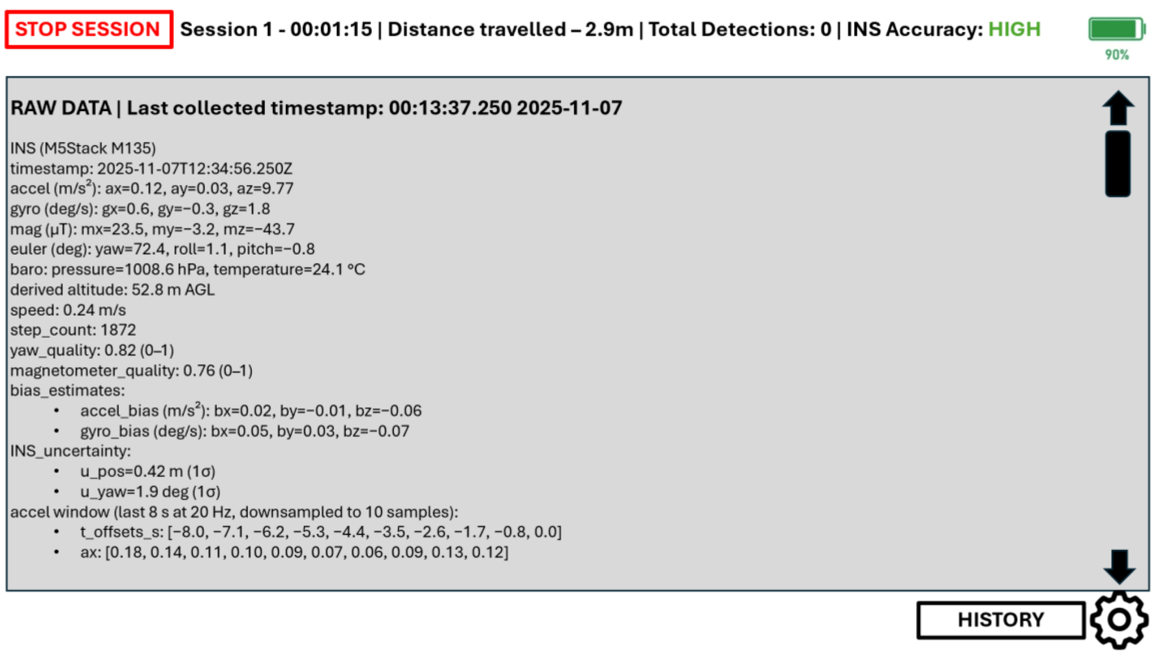

Figure 21: Sketch of Raw Data View

The raw data view is vital for calibration and debugging. It will show data coming from the radar, UWB, INS, and IR sensors in a list view that the user can scroll through vertically. The top of the section will show a timestamp of when the last data was collected. It will refresh at 10-15 Hz max to keep the screen responsive and updated with incoming data.

The INS section will provide a numeric stream of acceleration (ax, ay, az in m/s²), gyroscope (gx, gy, gz in deg/s), magnetometer (mx, my, mz in µT) and yaw/roll/pitch, barometric pressure, derived altitude, speed, step count, yaw quality. A mini plot will also show 3× accel and 3× gyroscope over the last 5-10s. Another mini plot will show the yaw with magnetometer corrections marked. Due to the inherent inevitability of INS accuracy waning, the raw data view will also show the magnetometer quality bar, bias estimates (from gyroscope/accelerometer), and INS uncertainty estimates (u_pos, u_yaw).

The radar raw data section will provide a summary per frame, including frame rate, number of points/tracks. It will also show profile parameters (range resolution, max range, FoV). A point list of the top 10-20 real frames will be included, featuring their range r, angle θ, SNR, velocity. A 2D polar mini view will show a polar plot of detections in radar frames. Diagnostics such as FFT peaks snapshot counts (potentially using a histogram) and presence/breathing/heartbeat flags will be included as well.

The UWB raw data will provide information about the CIR peak strength, which represents the amplitude of the strongest reflected path(s). It also includes the peak-to-noise ratio (PNR), multipath spread, which measures the width or RMS delay spread of significant taps, stability over time, and a presence score. The presence score is a heuristic that combines the above factors (such as the change in peak strength compared to the baseline and the spatial consistency) to indicate human presence through walls.

The IR raw data will contain statistics such as the maximum and mean temperatures, the index of the hottest pixel, and the status of temperature thresholds. Additionally, it will feature a 32×24 or 16×12 downsampled heatmap rendered as a thumbnail within a small area. Furthermore, it will include a list of hotspots, which provides the pixel coordinates (i, j), their temperature, and the inferred bearing.

7.1.3. Coordinate Map View

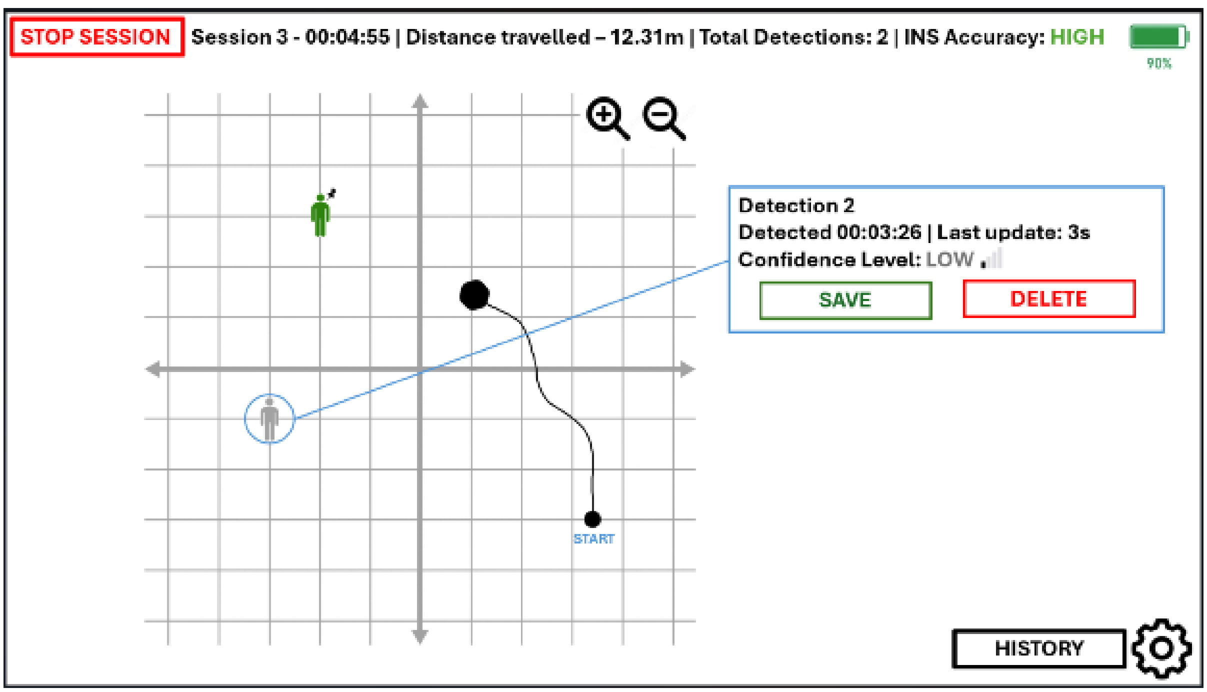

Figure 22: Sketch of Coordinate Map View

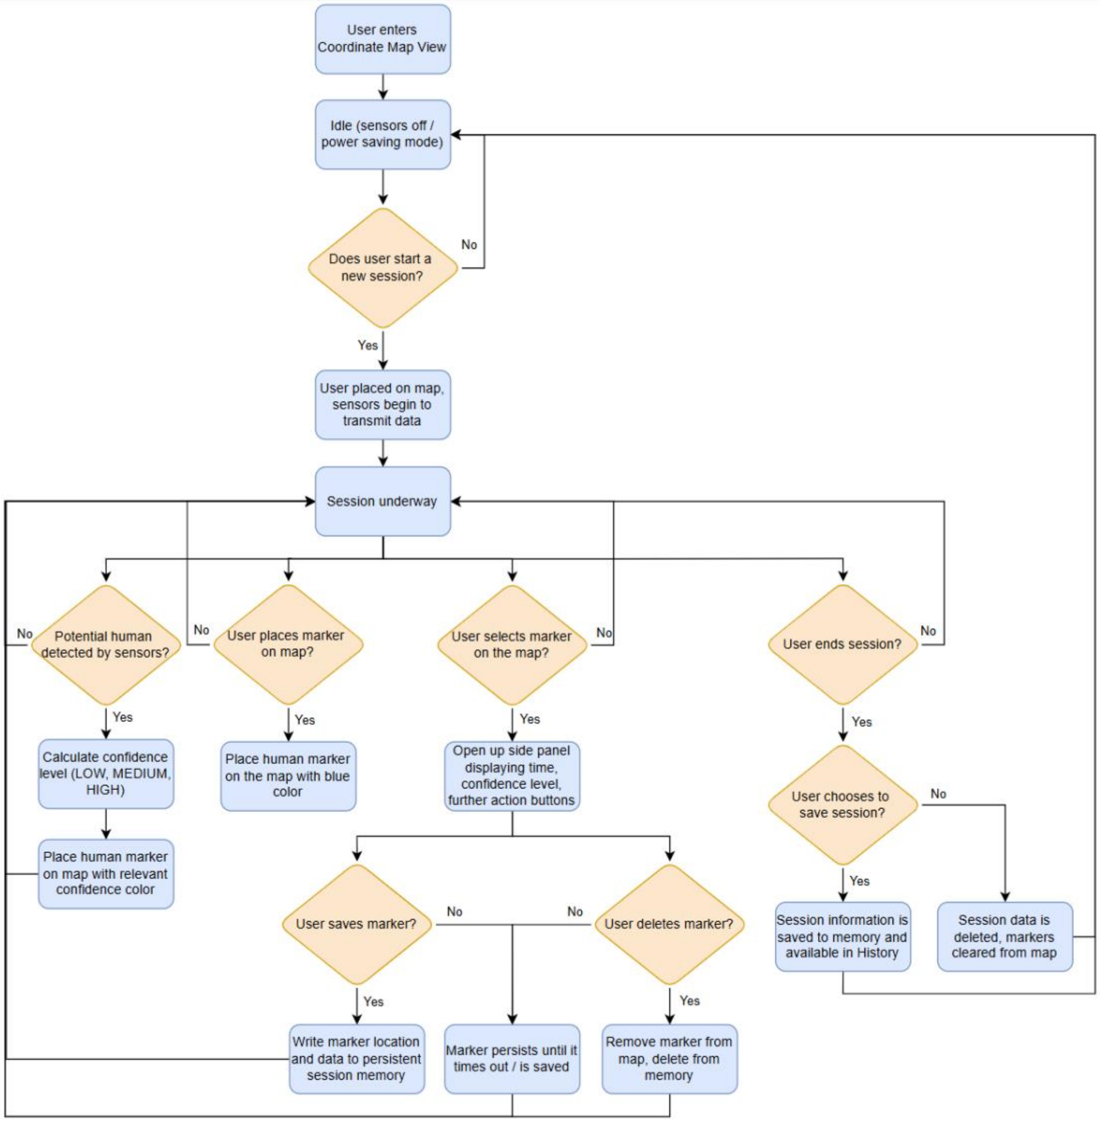

Figure 23: Coordinate Map View Flowchart

The coordinate map is the primary view for situational awareness during a scanning session. It presents a clean metric grid where the device's position, motion, and detected human activity are drawn in real time. The session origin is established when a new recording session begins, and all subsequent coordinates are plotted relative to this origin. The visual design favors clarity on the 4.3inch EastRising RA8875 TFT LCD, with a neutral background and high contrast elements for the device dot, path, and markers.

At the top of the screen, a status bar provides the essential context for the session. When the user presses the onscreen "Start New Session" button or the physical pushbutton, the session timer begins counting in hours, minutes, and seconds, the current device position is captured as the origin, and the blue LED begins its slow blink to indicate an active scan. The status bar shows the running timer, the cumulative distance traveled along the path, the number of currently visible candidate detections, the battery percentage, and a compact indicator of inertial navigation accuracy. The timer pauses and resumes with the corresponding controls and resets when the session is ended. Ending a session stops the timer, returns the map to an idle state, and prompts the user to save or discard a snapshot of the session for later review. Previous sessions can be reviewed through a history menu.

The center of the display is the map viewport. The device's current location is drawn as a solid dot with a short tick indicating heading. As the user moves, the map draws a continuous path from the origin to the current position. The path uses a color that remains legible over the grid. The distance traveled readout in the status bar updates as the path length grows. The map automatically manages scale, so the user does not need to constantly adjust zoom while walking. As the covered area expands, the viewport zooms out just enough to keep the current device position and most of the path comfortably in view with a small margin. To preserve readability on the RA8875, the autozoom includes a hard lower bound on scale; the map will not zoom out beyond a level where the full screen width represents roughly 90-100 meters. When this limit is reached, a brief onscreen hint clarifies that the maximum zoom out has been reached, and the viewport will then scroll instead of shrinking further. Zooming in past a practical upper bound is also constrained so that the user always retains sufficient context around the device. Magnifying glass buttons with '+' and '-' icons are available to manually zoom in and out. The touchscreen also affords direct manipulation of the view. Dragging a finger across the map pans the viewport, suspending auto centering. An unobtrusive "map unlocked" indicator appears to remind the user that the view is no longer following the device. A dedicated "recenter" button appears during this panning to allow the user to re-center the map on the current device position, thus resuming auto centering. Detections from the radar, ultrawideband, and infrared sensors are projected onto the map as human markers. These markers appear automatically at the estimated coordinates as new evidence arrives during the session. Each marker's icon and color communicate confidence in a consistent hierarchy: high confidence detections appear as solid green human shaped icons, medium confidence detections as amber icons, and low confidence candidates as gray icons. This allows the user to understand briefly which detections are likely to be reliable, and which require more corroboration or scrutiny. To reduce clutter, low-confidence candidate markers that are not refreshed by new sensor evidence disappear from the view after a period, customized in the settings. If the user wants to place a custom marker, they can do so by selecting the "place marker" button and tapping an area on the map. Doing so will place a human marker at that position, with a blue color to signify that it was a custom marker and not one detected by the device. Tapping any marker opens a concise detail panel on the right side of the screen and emphasizes the selected marker with a blue halo. The panel shows the current confidence level, when the marker was detected, and the timestamp of the last update. The panel also provides actions to save or dismiss the marker. Saving a marker promotes it to a persistent object on the map that will remain visible and stored in device memory until explicitly deleted or

a new session starts. Saved markers are visually distinguished with a small "pin" icon so they are easy to recognize. Tapping a saved marker reopens its detail panel with a delete option, allowing the user to remove it from the screen.
For practical reasons and to match the expected accuracy of the sensing modalities, the coordinate map enforces a strict rule that no two markers may exist within 2 meters of each other. This rule is applied consistently to reduce duplicate reporting of the same person and to keep the view uncluttered. If a new candidate appears within 2 meters of an existing saved marker, the system updates the existing marker's "last seen" time and may adjust its confidence upward rather than creating a new marker. If two temporary candidates are proposed within 2 meters of one another, only the highest confidence candidate is drawn, and the others are absorbed to avoid redundancy. If the user attempts to save a temporary marker that lies within 2 meters of an already saved marker, the save action is disabled, and a short message explains that the marker is too close to an existing one. This behavior gives predictable results and prevents accidental duplication in dense areas.

The coordinate map remains responsive without sacrificing stability. The UI updates at a capped refresh rate of approximately 10 to 15 updates per second, which is sufficient for smooth motion while allowing the ESP32 to prioritize sensor processing. Only the changed portions of the map are redrawn, when possible, such as appending the newest path segment, moving the device dot, or adding, updating, and removing markers. Sensor inputs are buffered so that UI drawing does not block data acquisition.

Session data is handled in a way that respects the user's expectations. Each session starts with a clear origin, a fresh path, and a zeroed timer. When the session ends, the user can save a snapshot of the path geometry, the session's timestamp, and any saved markers, which are then available for later review through a "History" button. The History menu provides a vertical scrollable list of entries that can be selected to view more information on them. Each entry is a numbered session with its total duration and total detections, and the real-world date and time the session was recorded on (using the ESP32's RTC initialized in the device settings page). Once any session is selected, the coordinate map from that session is displayed, which includes the start and end points and travel path, along with all saved

Altogether, the coordinate map gives the user a focused view of movement and detections in real time. It provides a dedicated button to start a new recording session and begin the timer, continuously tracks and displays the device location as a dot while drawing the traversed path, automatically scales the view as distance increases while honoring a sensible zoom-out limit, supports touch-based panning with a straightforward control to re-center on the device, automatically places detection markers that can be selected, saved, and deleted, and saves session data on persistent memory to be reviewed and accessed at any time.

7.1.4. Heatmap View

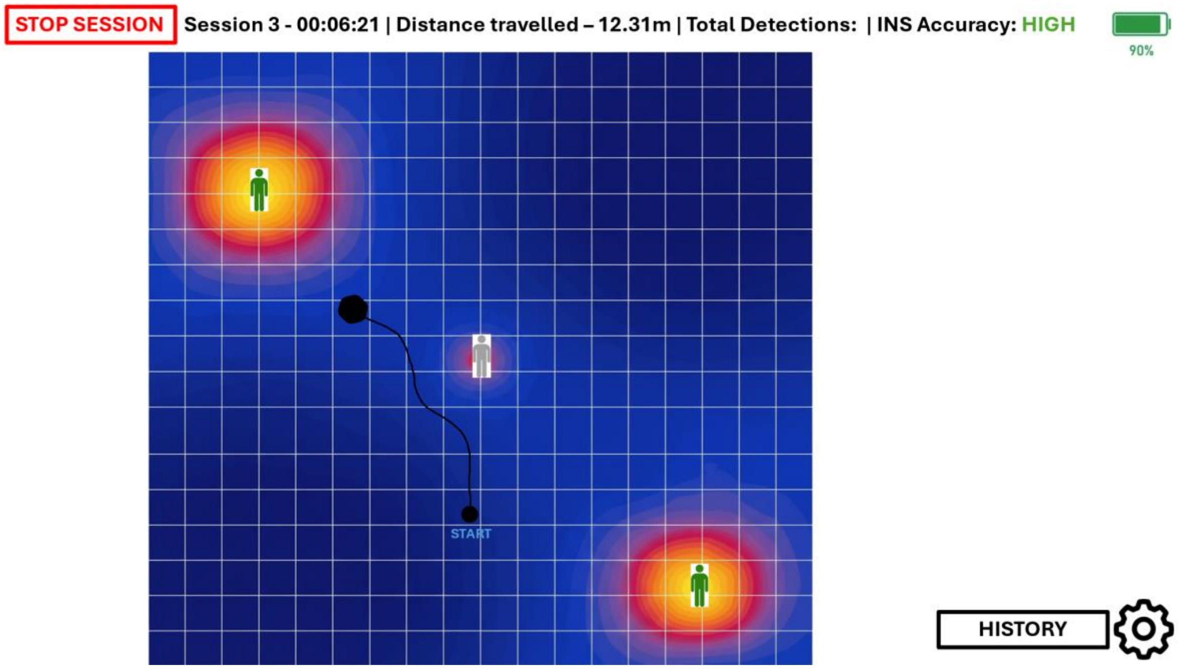

Figure 24: Sketch of Heatmap View

The heatmap view presents a top-down, gridded intensity field that highlights likely human presence through temperature and motion cues. The map shares the same session origin and coordinate frame as the coordinate map, but instead of discrete markers it renders a continuous surface where color encodes thermal prominence. Detection markers from the coordinate map, saved or unsaved, also appear on the heatmap. A neutral background grid remains faint, while a calibrated color ramp ranging from deep blue to dark red fills cells according to their score. The device's current position is shown as a small dot, and a dynamic scale bar helps interpret distances. As the user walks, the screen updates at a steady 10-15 Hz, smoothly fading older cells to emphasize the most recent, better-aligned observations while preventing the display from becoming saturated by stale data.

IR data from the MLX90640 drives the thermal component of the heatmap. Each frame's 32×24 temperature array is filtered to remove sensor noise and normalized against a rolling ambient baseline. A legend indicates the color scale and current ambient reference. Repeated passes accumulate into cells, increasing intensity where multiple observations agree. To reduce false positives, the thermal contribution is gated by corroboration from radar micro-Doppler or UWB presence cues; when both temperature elevation and presence evidence align spatially, the heatmap saturates more quickly.

Interaction is similar to the coordinate map interface. Dynamic zoom occurs according to distance travelled. Tapping any region opens a small panel with the peak temperature above ambient and last-seen time. The user can promote a region to a saved marker, which then appears in the coordinate map as well.

7.1.5. Smartphone App Software Design

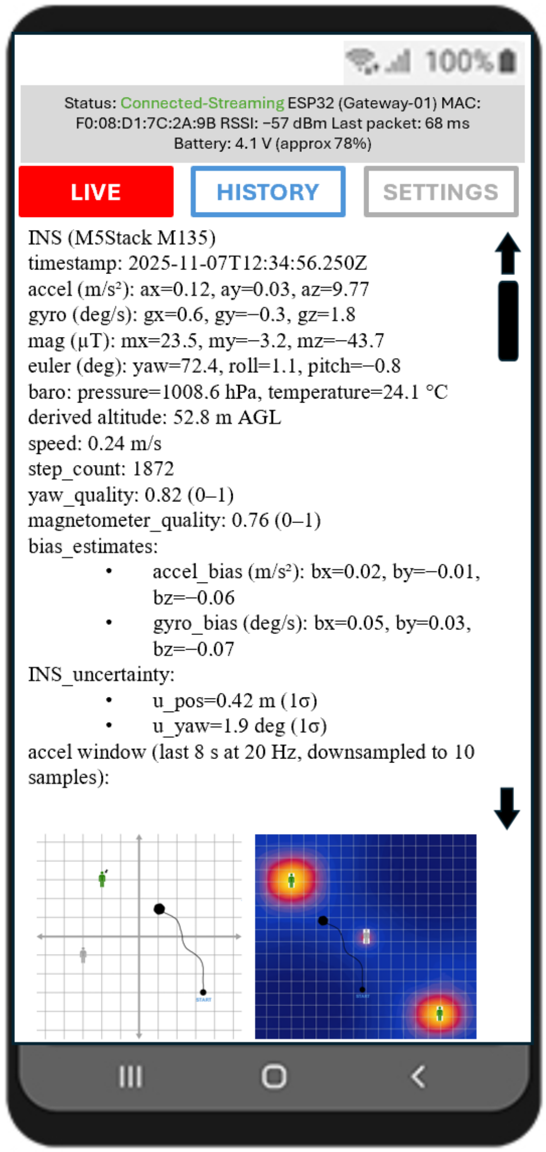

Figure 25: Sketch of Smartphone App User Interface

The Android app will be developed using Android Native (Kotlin and Jetpack Compose) and will mimic the Raw Data View of the device. The app opens to a clean "Live" dashboard that auto-attempts a BLE connection to the nearby device and shows a compact connection banner at the top (status, device name/ID, RSSI, last packet age). Below it, a simple tab strip lets the user switch between Live, History, and Settings. The Live tab presents readable, latency-aware tiles grouped by sensor domain: INS, Radar, UWB, and IR. Each tile uses large typography for headline metrics and terse labels, with unit suffixes always shown.

The INS tile leads, showing yaw/roll/pitch, speed, altitude, barometric pressure, step count, and yaw quality. The last 5-10 seconds of 3-axis accelerometer and gyroscope data is displayed along with bias estimates and INS uncertainties (u_pos, u_yaw). A mini-map card sits just under the INS tile. If a session is underway, it renders a scaled-down version of the device's coordinate map, including any markers on the map. The Radar tile summarizes the current frame rate, detections, and tracks. A compact 2D polar thumbnail shows the latest detections and tapping it reveals a sheet with the top 10–20 point list (r, θ, SNR, velocity) in a sortable table. The profile parameters (range resolution, max range, FoV) are shown as muted caption text, while diagnostics are accessible via a "Diagnostics" sub view that includes an FFT-peak histogram and breathing/heartbeat flags. The UWB tile shows a presence score and stability indicator. The strongest CIR peak strength and PNR are shown with clear units, and a small badge communicates temporal stability. A mini-trend line for peak strength/PNR occupies a single row. Tapping opens a detail sheet with multipath RMS delay spread. If the presence score crosses a threshold, the tile animates minimally (fade/pulse). The IR tile reports ambient, max, and mean temperatures and names the hottest pixel (i, j) with its inferred bearing. A tiny down sampled heatmap thumbnail is present. Tapping it opens a full screen viewer with a simple temperature legend and hotspot list. Hotspots are tappable chips that pan the heatmap to the selected pixel and show a callout with temperature and bearing, a toggle overlays contour lines. If a high/low threshold is tripped, the tile shows a prominent but unobtrusive banner with the threshold label and current max/min.

Using the app is straightforward: on opening, the user confirms the BLE connection and glances at the Live tiles, which will be constantly updating the data received from the device if a session is underway. If they're walking, they'll likely lock the map heading and watch the breadcrumb trail while monitoring speed and yaw. To check presence detection, they'll scan the Radar and UWB tiles. Any presence/breathing flags prompt them to tap for detail, check the polar plot clusters, and optionally pin a presence marker to compare against the map position. For thermal checks, a quick tap on the IR tile brings up the heatmap and hotspot list, useful for aiming or verifying a heat source direction that complements the radar detections.

The History tab offers a simple timeline of sessions with date, duration, and rough coverage distance. Tapping a session shows summaries for INS drift over time, radar detection counts, UWB presence score trends, and IR max temperature ranges, each as lightweight sparklines and key stats. A "Compare" toggle can overlay two metrics on a shared relative axis for quick correlation without complex charting.

The Settings tab focuses on clarity and safety: a device picker with remembered devices, a sampling-rate display (read-only), units (metric/imperial for derived speed and altitude, °C/°F for IR), and notification toggles (presence detected, high temperature). Accessibility options include larger text and color-blind-friendly palettes that adjust the polar plot and heatmap gradients. A “Data Glossary” sheet explains terms like PNR, SNR, yaw quality, and u_pos with one-line definitions tied to the tiles’ info icons, ensuring the UI remains minimalist while discoverable.

Small touches to UI enhance user experience. Values animate smoothly when they change to reduce flicker, stale data greys out with an age counter, and each tile shows a compact “last update” timestamp. The overall experience is efficient and easy to navigate, with detailed information just a few taps away whenever needed.

CHAPTER 8

CHAPTER 9

CHAPTER 10

10.1. Budget

Since this project is entirely self-funded, maintaining affordability without compromising functionality was a primary consideration throughout the component selection process. The primary objective was to maintain the total cost close to $500, ensuring a harmonious balance between performance, sensor accuracy, and system reliability. The total cost of the current bill of materials is approximately $505, which encompasses all major modules, including the radar, inertial navigation system, UWB transceiver, LCD display, and supporting power management circuitry. This estimate includes an estimated cost for the PCB fabrication and assembly, which will be updated once the final design is confirmed. The team prioritized using cost-effective yet proven components, such as the ESP32 microcontroller and off-the-shelf sensor modules, to minimize prototyping expenses while ensuring future expandability. Considering the intricate nature of the system and the specialized design of the radar and thermal imaging sensors, maintaining this proximity to the target budget showcases efficient component selection and effective resource management.

Sub-System	Budget
Sensing Modules (Radar, UWB, INS, IR Thermal Sensor)	$396.80
Processing & Control (MCU, Power Regulation, MOSFETs, Inductors)	$21.00

Sub-System	Budget
User Interface (LCD, LEDs, Pushbuttons, Buzzer, Light Pipes)	$46.00
Power & Protection (Battery, Charger/BMS, USB, TVS, Diodes, Connectors)	$16.00
PCB Fabrication & Assembly (Estimated)	$50.00
Passive Components (Resistors, Capacitors, Misc. Hardware)	$10.00
Total Estimated Cost	$539.00

Table 26: Budget allocation table for sub-systems.

The table above outlines the budget according to each subsystem, providing a clear distribution of how project funds were allocated. The sensing modules, which encompass the radar, UWB, INS, and thermal imaging sensor, constitute the largest portion of the budget. These components are crucial for the device's primary life-detection function. The user interface and processing systems follow, encompassing the LCD display, pushbuttons, LEDs, and the ESP32 microcontroller that coordinates all sensor operations. Power management and protection circuits, including the battery, voltage regulators, and safety components, were designed to remain efficient and affordable. Passive elements such as capacitors, resistors, and inductors add minimal cost but are essential for stable and low noise performance. Overall, this breakdown showcases a well-rounded approach that strikes a harmonious balance between functionality, durability, and cost effectiveness, all while adhering to the self-funded constraints of the project.

Component	Model	Price	Quantity
Radar Module	TI IWR6843AOP EVM	$224.35	1
INS	M5Stack M135	$50.00	1
UWB	DWM3000EVB	$42.50	1
MCU	ESP32-32E	$11.72	1
LCD	EastRising 4.3″ IPS RA8875 TFT	$25.00	1
Buzzer	CMI-1295-03TH	$1.14	1
Micro-USB-B	105017-0001	$0.86	
Red LED	OSRAM LS R976	$0.18	1
Yellow LED	OSRAM LY R976	$0.18	1

Component	Model	Price	Quantity
Orange LED	OSRAM LG R976	$0.18	1
Pushbutton	Apem IPR1SAD2	$5.50	3
IR Thermal Sensor	Adafruit MLX90640	$79.95	1
Radar Module Breakout Board	Custom PCB	$5.51	1
Battery	18650 Lithium-ion	$9.00	1
Light Pipe	515-1020-805F	$3.00	1
3.3V Regulator	TPS563200	$1.10	1
5V Regulator	TPS61236P	$2.40	1
Diode	SS8050 NPN Schottky	$0.29	2
Diode USB	Diode USB CDSOT23	$1.36	1
USB-UART	FT231XS-R	$2.34	1
1S BMS Module	DW01A + FS8205A-based board	$1.50	1
Inductor (5V Boost)	Bourns SRP4020TA-2R2M	$1.85	1
Inductor (3.3V)	Bourns SRP5030TA-4R7M	$1.95	1
Backlight LED Driver	MPS MP3302	$1.75	1
MOSFET	AO3400A	$0.10	2
Capacitors	Assorted	$3.50	25
Resistors	Assorted	$3.00	25
Total		$480.21	

Table 27: Bill of Materials

10.2. Division of Work

The table below provides a clear breakdown of the responsibilities assigned to each member of the 2HAD-RAD development team. It outlines their primary and secondary roles within the project's various subsystems. This structure ensures a balanced distribution of workloads and fosters cross-disciplinary collaboration between electrical and computer engineering disciplines. Each subsystem was assigned a primary lead responsible for its design, implementation, and testing. A secondary team member provided support for verification, integration, and documentation. This collaborative approach enables the sharing of overlapping expertise, particularly in power management, sensing, and data acquisition. This ensures the system's reliability and continuity throughout the design, prototyping, and testing phases.

Subsystem	Primary Lead	Secondary Support
Power Subsystem	Bailey Hitt (EE) – Designs and integrates the Li-ion battery system, BMS, and DC/DC regulation circuits for 3.3 V and 5 V rails.	Shannon Hankinson (EE) – Assists with component sourcing, protection circuitry, and power efficiency testing.
Sensors Subsystem	Ryan Andersen (EE) – Leads integration of UWB, INS, and IR sensors, managing I2C/SPI communication with the MCU.	Bailey Hitt (EE) – Supports sensor calibration and system validation under field conditions.
Radar Module	Shannon Hankinson (EE) – Implements the TI IWR6843AOP radar, handling configuration, mounting, and communication with the main controller.	Ryan Andersen (EE) – Provides firmware assistance and performance tuning for radar data fusion.
Microcontroller (MCU)	Joshua D'Souza (CpE) – Develops firmware, SPI/UART interfacing, and data acquisition timing control.	Bailey Hitt (EE) – Supports hardware verification and debugging of MCU-based communication.

Subsystem	Primary Lead	Secondary Support
User Interface (UI)	Joshua D'Souza (CpE) – Designs and codes LCD interactions, LED logic, button handling, and buzzer control.	Ryan Andersen (EE) – Assists with signal routing, interface layout, and UI testing.

Table 28: Division of Work

10.3. Milestones

Milestone	Start Date	Anticipated End Date
Project Brainstorming	8/19/25	8/26/25
Project Scope Finalized	8/26/25	9/6/25
Individual Research	8/26/25	9/6/25
Initial Proposal Submission	8/26/25	9/6/25
30-Page Milestone	9/20/25	9/27/25
60-Page Milestone	9/6/25	9/27/25
80-Page Milestone	9/27/25	10/31/25
90-Page Milestone	10/31/25	11/7/25
120-Page Milestone	11/7/25	11/25/25
Final Draft	11/25/25	12/3/25

Table 29: Senior Design 1 documentation milestones

Milestone	Start Date	Anticipated End Date
System Design	11/1/25	11/15/25
System Testing	11/16/25	11/22/25
Prototype	11/23/25	11/29/25
PCB	11/20/25	12/6/25

Table 30: Senior Design 1 design/prototype milestones

Milestone	Start Date	Anticipated End Date
PCB Testing	1/12/26	1/31/26
System Integration and Testing	2/1/26	3/14/26
Practice Demo	3/15/26	3/28/26
Finalize Documentation	3/29/26	4/11/26
Practice Final Presentation	4/12/26	4/25/26
Final Presentation	TBD	TBD

Table 31: Estimated Senior Design 2 milestones

The overall project plan outlines a well-defined timeline for the development of the 2HAD-RAD system, providing a roadmap for the team from the initial concept phase to the complete integration and testing of the system. The milestones are strategically spaced to provide sufficient time for research, design validation, and hardware development, ensuring clear deliverables at each stage. While the Senior Design 1 schedule offers a comprehensive roadmap, including detailed documentation and preliminary prototyping, the timeline for Senior Design 2 is still an initial estimate. As the project progresses and more design details are finalized, this schedule will be revisited and refined to reflect the actual progress, component lead times, and testing outcomes. This ensures that the completion plan is realistic and achievable.

Appendix A

[1a] Prompt: “What are the most effective ways to optimize power consumption in a dual-rail Li-ion–powered radar sensing system?”

[1b] ChatGPT Outcome: To optimize power consumption in a dual-rail Li-ion–powered radar sensing system, it’s important to consider both hardware efficiency and firmware-level power management. On the hardware side, using synchronous buck-boost converters with low quiescent current (IQ < 50 µA) significantly reduces idle losses. Separate regulation for the 3.3 V and 5 V rails ensures that conversion efficiency is maintained across load variations. Additionally, implementing load switches or MOSFET-based power gating on subsystems like the display, radar, and communication modules allows selective power delivery only when necessary.

On the firmware side, duty-cycling the radar—activating it only during scan intervals—can reduce average current by more than 60%. The ESP32’s deep-sleep mode should be leveraged between radar bursts, and periodic wake-ups scheduled via timers or interrupts. Adjusting radar sampling frequency dynamically based on signal activity can further improve efficiency. Real-time current monitoring via a shunt-based sensor (e.g., INA219) can help profile load behavior and validate design assumptions. Together, these techniques establish a layered strategy that improves both runtime and thermal performance without compromising system responsiveness.

[1c] Claude Outcome: Optimizing power consumption in a dual-rail Li-ion-powered radar sensing system requires focusing on a few key areas. The single most impactful strategy is maximizing sleep time between radar measurements, since the RF transmitter typically dominates power consumption during active operation. Use intelligent duty cycling with motion sensors or other low-power triggers to activate the radar only when needed, and implement deep sleep modes that draw only microamperes when idle.

On the hardware side, use high-efficiency DC-DC converters (>90% efficiency) with low quiescent current for both power rails, and consider power gating to completely disconnect unused sections between measurements. Dynamic voltage scaling can adjust rail voltages based on real-time performance requirements. Select radar transceivers and microcontrollers specifically designed for low-power operation with fast sleep-to-active transitions, and use hardware accelerators for signal processing tasks like FFTs rather than relying on general-purpose processors.

For radar operation itself, optimize chirp parameters to balance performance with energy per measurement, reduce sampling rates when no targets are detected, and implement adaptive frame rates based on scene activity. Accurate battery monitoring through fuel gauge ICs enables intelligent power management decisions throughout the system's operational life. With proper optimization across these areas, you can achieve 100-1000x power reduction compared to continuous operation, dramatically extending battery life.

[2a] Prompt: "What filtering techniques can be applied to reduce false positives in a short-range radar-based life-detection system?"

[2b] ChatGPT Outcome: In short-range radar-based life-detection systems, false positives often arise from motion artifacts, environmental reflections, and electrical noise. Applying digital signal filtering at both the hardware and software stages is crucial. A Butterworth low-pass filter effectively removes high-frequency noise while maintaining signal shape, whereas Chebyshev Type II filters provide steeper roll-off for environments with dense interference. For dynamic background suppression, adaptive Kalman filtering can continuously update prediction models based on signal variance, improving accuracy over time.

Additionally, moving-average filters help stabilize respiration or heartbeat signals, particularly when integrated with real-time peak detection algorithms. When computational resources allow, a combination of FFT-based noise analysisand band-pass filtering centered on the human respiration frequency range (0.1–0.4 Hz) yields optimal discrimination between genuine motion and background noise. These methods can be implemented in MATLAB or directly within the radar's embedded DSP pipeline.

[2c] Claude Outcome: Reducing false positives in a short-range radar-based life-detection system requires combining multiple filtering techniques across the signal processing chain. Start with robust clutter removal using high-pass filtering or background subtraction to eliminate static objects, which are often the primary source of false alarms. Implement adaptive clutter cancellation that continuously updates the background model to handle slowly changing environments like swaying vegetation or temperature-related drift.

For vital sign detection, apply narrow bandpass filtering around expected heartbeat (0.8-2 Hz) and respiration (0.1-0.5 Hz) frequencies to isolate physiological signatures while rejecting out-of-band noise. Use phase-based processing techniques rather than just amplitude, as phase information provides more reliable vital sign extraction and is less susceptible to multipath interference. Implement range gating to focus on specific distance bins where a person is likely located, which dramatically reduces the search space and false alarm potential. Apply temporal consistency checks by requiring life signs to persist for multiple consecutive frames before declaring a detection, filtering out transient disturbances like passing vehicles or moving objects. Use machine learning classifiers trained on real vital sign data versus common false alarm sources to distinguish genuine human signatures from mimics like fans, machinery vibration, or animal movement. Cross-validation techniques like requiring simultaneous detection of both heartbeat and respiration patterns, with physiologically plausible phase relationships between them, significantly improve specificity. Finally, implement adaptive thresholding that adjusts detection sensitivity based on environmental noise levels and clutter statistics, preventing both missed detections in challenging conditions and excessive false alarms in benign environments.

[3a] Prompt: “What communication protocol offers the best balance between speed, reliability, and wiring simplicity for interfacing a radar baseboard with an ESP32 microcontroller?”

[3b] ChatGPT Outcome: For interfacing a radar baseboard with an ESP32 microcontroller, SPI (Serial Peripheral Interface) provides the optimal balance of speed, reliability, and design simplicity. SPI offers data rates exceeding 10 Mbps, ideal for transferring radar frame data with minimal latency. Its full-duplex communication mode and support for multiple slaves make it suitable for multi-device setups involving a display or secondary sensors.

ChatGPT also discussed signal integrity considerations, recommending short traces and proper grounding to minimize noise coupling. In systems requiring asynchronous logging or debug data, UART serves as a complementary channel. The model also highlighted that SPI’s lower overhead and deterministic timing outperform I^2C for time-critical radar data streams, making it the preferred protocol for the main data path.

[3c] Claude Outcome: For interfacing a radar baseboard with an ESP32 microcontroller, SPI (Serial Peripheral Interface) typically offers the best balance between speed, reliability, and wiring simplicity. SPI provides high data rates of 10-80 MHz on the ESP32, which is more than adequate for transferring radar data frames, and requires only four wires (MOSI, MISO, SCLK, and CS). The protocol is straightforward to implement with well-supported libraries, minimal CPU overhead through DMA support, and deterministic timing that's crucial for real-time radar data acquisition.

For most radar applications requiring raw or minimally processed data transfer, SPI's combination of high throughput, low latency, hardware acceleration support on the ESP32, and relatively simple four-wire connection makes it the optimal choice. The ESP32 has multiple SPI peripherals available, allowing you to dedicate one exclusively to radar communication while using others for displays or SD cards if needed.

[4a] Prompt: “Which enclosure materials are best suited for housing a radar-based life-detection device that must minimize interference and remain lightweight?”

[4b] ChatGPT Outcome: For this type of radar-based system, I recommend using polycarbonate-ABS (PC-ABS) due to its excellent mechanical durability, low dielectric constant, and compatibility with RF applications. PC-ABS maintains radar signal transparency and provides sufficient impact protection for field use. I

would avoid metallic or carbon-filled enclosures since they can reflect or absorb radar energy, reducing signal quality.

To improve environmental performance, I suggest an IP54 or higher enclosure rating, achieved by using silicone gaskets and minimal seam lines. A conformal coating on internal PCBs further protects against humidity and dust. Thermal management can be achieved with discreet vents or internal heat spreaders while maintaining RF clarity. Overall, PC-ABS offers the best balance between strength, low weight, and radar transparency.

[4c] Perplexity Outcome: For housing a radar-based life-detection device, the best enclosure materials balance RF transparency, mechanical protection, and weight. Plastics are generally the optimal choice, with ABS, polycarbonate, and PETG offering excellent radar transparency across common frequencies (24 GHz, 60 GHz, 77 GHz bands) while providing good impact resistance and being lightweight. These materials introduce minimal signal attenuation, typically less than 1-2 dB, which preserves both transmit power and receive sensitivity critical for life detection.

Composite materials like fiberglass or carbon fiber offer excellent strength-to-weight ratios, but carbon fiber is electrically conductive and will block radar signals, so it should only be used in non-critical areas away from the antenna aperture. If you need EMI shielding for the electronics section to prevent interference with the radar receiver, use internal metal partitions or conductive coatings in compartments separate from the antenna, creating isolated RF zones. For outdoor applications, ensure whatever material you choose has good UV stability and moisture resistance, as environmental degradation can change dielectric properties over time and affect radar performance.

[5a] Prompt: “What are effective methods for visualizing radar-based motion and breathing data on a 4.3-inch TFT display?”

[5b] ChatGPT Outcome: For visualizing radar data in real time, I recommend using waveform and bar graph visualizations to represent both motion amplitude and respiration rate. A scrolling waveform offers a clear view of breathing patterns, while a color-coded bar indicator can show motion intensity. I suggest implementing double-buffered rendering and partial screen updates to prevent flicker or tearing, especially when displaying continuous data streams.

In addition, I would apply adaptive brightness control based on ambient light, along with high-contrast color schemes for readability in bright conditions. Icons or status indicators could display connection state, detection status, or error messages. Overall, the goal is to keep the display simple, responsive, and visually intuitive for users operating in high-stress or outdoor environments.

[5c] Claude Outcome: For visualizing radar-based motion and breathing data on a 4.3-inch TFT display, focus on real-time waveforms and simple visual

indicators that convey information at a glance. Display breathing and heartbeat as scrolling time-domain waveforms, similar to medical monitors, with the past few seconds of data visible so users can quickly assess regularity and presence of vital signs. Use color coding to indicate signal quality or detection confidence, such as green for strong signals, yellow for marginal, and red for no detection or poor quality.

For motion detection, use a simple zone-based display showing which range bins or spatial regions contain detected movement, represented as a top-down or side-view map with color intensity indicating motion magnitude. A waterfall or spectrogram display can effectively show frequency content over time, useful for debugging and detailed analysis, though this may be too technical for end users. Keep the interface uncluttered given the small screen size—prioritize the most actionable information and use tabbed views or button presses to access detailed diagnostics. Finally, ensure your refresh rate matches your radar update rate but limit unnecessary screen redraws to minimize flicker and power consumption, particularly important if battery-powered.

REFERENCES

[1] O. B. Lubecke, P.-W. Ong, and V. M. Lubecke, “10 GHz Doppler radar sensing of respiration and heart movement,” IEEE Xplore, [Online]. Available: https://ieeexplore.ieee.org/document/999462

[2] Infineon Technologies AG, BGT60TR13C Radar Sensor Datasheet, Rev. 1.0, 2020. [Online]. Available: https://www.infineon.com/dgdl/Infineon-BGT60TR13C-DS-v01_00-EN.pdf

[3] Federal Communications Commission, Title 47 CFR Part 15 – Radio Frequency Devices, U.S. Government Publishing Office, 2023. [Online]. Available: https://www.ecfr.gov/current/title-47/chapter-I/subchapter-A/part-15

[4] International Electrotechnical Commission, IEC 62133-2: Secondary Cells and Batteries… for Portable Applications, 2nd ed., Geneva, Switzerland, 2017.

[5] Underwriters Laboratories, UL 2054: Household and Commercial Batteries, Northbrook, IL, USA, 2011.

[6] IEEE Standards Association, IEEE 802.15.4™-2020 – Low-Rate Wireless Personal Area Networks (LR-WPANs), IEEE, New York, NY, USA, 2020.

[7] Department of Homeland Security, Portable Signs of Life Identifier: Technology Scouting Research Summary, Aug. 2019. [Online]. Available: https://www.dhs.gov/sites/default/files/publications/portable_signs_of_life_updated-508c_v3.pdf

[8] Department of Homeland Security, Detection of Presence of Life Through Walls, S&T Directorate, Dec. 2022. [Online].
Available: https://www.dhs.gov/sites/default/files/2022-12/22_1214_st_Detection_Presence_Life%20throughWalls_122022.pdf

[9] G. Charvat, Small and Short-Range Radar Systems, CRC Press, 2014.

[10] H. Mostafa et al., "Sensor Fusion for Non-Contact Human Presence Detection," IEEE Sensors Journal, vol. 21, no. 10, pp. 11938–11947, 2021.

[11] Camero-Tech Ltd., Xaver™ 100 and 400 Product Overview, 2023.

[12] Vayyar Imaging, Vayyar Sensor Technology Overview, Technical Datasheet, 2023.

[13] S. Wang et al., "Combined UWB Radar and Infrared Sensing for Human Vital Sign Detection," Sensors, vol. 23, no. 5, 2023.

[14] A. Kharbouch et al., "Hybrid Radar–Thermal Sensor System for Disaster Victim Localization," IEEE Trans. Geosci. Remote Sens., vol. 60, pp. 1–14, 2022.

[15] Texas Instruments, IWR6843AOP Single-Chip 60–64 GHz mmWave Sensor (AoP) – Datasheet, Rev. B, 2023. [Online].
Available: https://www.ti.com/lit/ds/symlink/iwr6843aop.pdf

[16] Qorvo Inc., DWM3000 Ultra-Wideband Module Data Sheet, Rev. B, May 2021. [Online]. Available: https://www.qorvo.com/products/p/DWM3000

[17] Melexis N.V., MLX90640 Infrared Array (32×24) Datasheet, Rev. 14, Mar. 2024. [Online]. Available: https://www.melexis.com/-/media/files/documents/datasheets/mlx90640-datasheet-melexis.pdf

[18] Espressif Systems, ESP32-WROOM-32 Datasheet, Rev. 3.9, Feb. 2024. [Online].
Available: https://www.espressif.com/sites/default/files/documentation/esp32-wroom-32_datasheet_en.pdf

[19] Espressif Systems, ESP-IDF Programming Guide (Firmware & FreeRTOS API), 2024. [Online]. Available: https://docs.espressif.com/projects/esp-idf/en/latest/esp32/

[20] NXP Semiconductors, UM10204 – I²C-Bus Specification and User Manual, Rev. 7, Oct. 2021. [Online]. Available: https://www.nxp.com/docs/en/user-guide/UM10204.pdf

[21] RAiO Technology Inc., RA8875 TFT LCD Controller Datasheet, Jan. 10, 2012.

[22] EastRising Technology Co., Ltd., ER-TFTM043A1-7S TFT LCD Module Datasheet, Apr. 2, 2021. [Online]. Available: https://www.buydisplay.com/download/manual/ER-TFTM043A1-7S_Datasheet.pdf

[23] Bosch Sensortec, BMI270 IMU Datasheet, BST-BMI270-DS000-07, Mar. 13, 2023.

[24] Bosch Sensortec, BMM150 Magnetometer Datasheet, 2023.

[25] Bosch Sensortec, BMP390 – Digital Pressure Sensor Datasheet, Rev. 1.7, Mar. 2021.

[26] u-blox AG, ZED-F9R High Precision Sensor Fusion GNSS Receiver (ZED-F9R-03B) Data Sheet, Doc. UBX-22024085, Rev. R04, May 3, 2024.

[27] u-blox AG, NEO-M9V Multi-Mode Dead Reckoning Module Data Sheet (NEO-M9V-20B), Doc. UBX-21029781, Rev. R07, Jun. 2, 2025.

[28] M5Stack, Module GNSS – SKU M135 (Product Wiki and Datasheet), Aug. 2025.

[29] Sercel, SM-24 Geophone Element Brochure, Carquefou, France, 2012. [Online]. Available: https://cdn.sparkfun.com/datasheets/Sensors/Accelerometers/SM-24%20Brochure.pdf

[30] Seis Tech Co., Ltd., Data Sheet of SM-24 10 Hz Equivalent Geophone, Shenzhen, China, 2024. [Online]. Available: https://www.seis-tech.com/wp-content/uploads/2024/08/data-sheet-of-sm24-10hz-equivalent-geophone.pdf

[31] D. H. Titterton and J. L. Weston, Strapdown Inertial Navigation Technology, 2nd ed., Stevenage, U.K.: IET, 2004.

[32] P. D. Groves, Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems, 2nd ed., Boston, MA: Artech House, 2013.

[33] O. J. Woodman, “An Introduction to Inertial Navigation,” Univ. of Cambridge, Tech. Rep. UCAM-CL-TR-696, Aug. 2007.

[34] E. Foxlin, “Inertial Head-Tracker Sensor Fusion by a Complementary Separate-Bias Kalman Filter,” Proc. IEEE Virtual Reality Symp., 1996, pp. 185–194.

[35] R. Zetik, J. Sachs, and R. Thomä, “UWB Localization—Advantages and Challenges,” Proc. IEEE Int. Conf. Ultra-Wideband (ICUWB), Sep. 2004, pp. 771–778.

[36] F. Montorsi, M. Shubair, and R. Fantacci, “Cooperative Localization Techniques for UWB Wireless Sensor Networks,” IEEE Commun. Mag., vol. 53, no. 6, pp. 124–131, Jun. 2015.

[37] S. Lazaro, D. Girbau, and R. Villarino, “Analysis of Vital Signs Monitoring Using an IR-UWB Radar,” Progress In Electromagnetics Research, vol. 100, pp. 265–284, 2010.

[38] J. Sachs, R. Zetik, and R. Thomä, “Through-Wall Imaging with UWB Radar,” Proc. IEEE European Radar Conf., Oct. 2005, pp. 561–564.

[39] International Telecommunication Union (ITU-R), Recommendation M.1477 – Technical and Performance Requirements for GNSS Receivers, Geneva, Switzerland, 2017.

[40] U.S. Department of Defense, IS-GPS-200H: Navstar GPS Space Segment/Navigation User Interfaces, El Segundo, CA, USA, 2019.