# ScopeFace: Digital Stethoscope

SHAYAN BOGHANI, University of California, San Diego, USA
AMRITA MOTURI, University of California, San Diego, USA
SHELBY MYRMAN, University of California, San Diego, USA
SIYA RAJPAL, University of California, San Diego, USA

Traditional stethoscopes are widely used in clinical settings but face limitations in noisy, fast-paced environments where faint cardiac or respiratory sounds may be missed, compromising diagnostic accuracy and continuity of care, increasing uncertainty for clinicians and risk for patients. To address these challenges, we present ScopeFace, a cloud-enabled digital stethoscope designed to amplify, record, and replay auscultatory sounds with clarity. By preserving high-fidelity audio and enabling in-app playback and session storage, ScopeFace enhances diagnostic confidence, supports collaboration, and promotes more informed, timely decision-making. Our system aims to improve both clinical workflow and patient outcomes through accessible, technology-enhanced auscultation.

**ACM Reference Format:**
Shayan Boghani, Amrita Moturi, Shelby Myrman, and Siya Rajpal. 2025. ScopeFace: Digital Stethoscope. 1, 1 (June 2025), 18 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

## 1 Introduction

### 1.1 Background and Motivation

Auscultation, or the act of listening to internal body sounds, remains one of the most fundamental diagnostic techniques in medicine. For over two centuries, the acoustic stethoscope has served as the primary tool for this task, prized for its simplicity and portability. However, in practice, clinicians often work in loud, chaotic environments like emergency rooms or inpatient wards where ambient noise, patient movement, and competing demands can make it difficult to detect subtle but clinically significant sounds. These limitations can result in missed diagnoses, delayed care, and added anxiety for both clinicians and patients.

Recent advancements in digital health technology offer opportunities to address these challenges by augmenting traditional tools with capabilities like amplification, recording, and cloud connectivity. Digital stethoscopes exist in the market, but many fall short in usability, data integration, cost, and accessibility. Most options do not provide seamless workflows that allow care teams to replay sounds clearly or collaborate asynchronously on complex cases. As telemedicine and interdisciplinary care become increasingly important, the ability to share high-fidelity auscultation data in real time or retrospectively is a necessity.

Authors' Contact Information: Shayan Boghani, University of California, San Diego, La Jolla, California, USA, sboghani@ucsd.edu; Amrita Moturi, University of California, San Diego, La Jolla, California, USA, amoturi@ucsd.edu; Shelby Myrman, University of California, San Diego, La Jolla, California, USA, smyrman@ucsd.edu; Siya Rajpal, University of California, San Diego, La Jolla, California, USA, sirajpal@ucsd.edu.


ACM XXXX-XXXX/2025/6-ART
https://doi.org/10.1145/nnnnnnn.nnnnnnn

### 1.2 Overview of ScopeFace

This paper introduces ScopeFace, a modern digital stethoscope system designed to meet the real-world needs of clinicians. ScopeFace captures and preserves high-quality heart and lung sounds, stores them securely in the cloud, and enables users to review and share recordings across devices. By combining clear audio playback, intuitive user interface design, and data storage, ScopeFace offers a portable and collaborative auscultation solution.

Unlike many bulky digital stethoscopes or traditional acoustic models, ScopeFace is engineered for extreme portability. The compact hardware is small enough to fit comfortably in a scrub pocket or coat, allowing clinicians to carry it easily during rounds or between patient visits. This lightweight form factor ensures that digital auscultation becomes a frictionless part of daily clinical routines, rather than an additional burden. Portability also opens new use cases in remote, resource-constrained, or home care settings where traditional tools may be impractical.

ScopeFace was developed through interdisciplinary collaboration between hardware and software teams, emphasizing clinical usability, technical robustness, and future extensibility. The system is designed to be lightweight and affordable enough for widespread adoption while providing the core functionality needed to improve confidence in auscultatory assessments. Importantly, it supports both in-person and remote diagnostic workflows, helping to close gaps in continuity of care.

### 1.3 Key Contributions

This paper presents the design, implementation, and potential impact of ScopeFace, a portable, cloud-connected digital stethoscope. The first contribution is the development of a novel hardware prototype that amplifies and records heart and lung sounds with high fidelity. Designed with clinical usability in mind, the device preserves subtle auscultatory signals that are often missed in traditional settings due to noise, distraction, or time pressure.

Second, this paper introduces a lightweight mobile interface, developed in SwiftUI, that enables users to review, annotate, and organize audio recordings in an intuitive way. This interface streamlines clinician workflows by allowing on-the-go playback and session management. Paired with our backend infrastructure, the application supports secure cloud storage and enables better collaboration across care teams. While ScopeFace is not currently HIPAA-compliant, it is built with future compliance in mind, prioritizing secure data handling and user privacy as foundational elements.

A third major contribution is the device's compact form factor. Unlike traditional digital stethoscopes, which are often bulky or tethered to larger systems, ScopeFace is small enough to fit comfortably in a pocket. Its lightweight design allows for greater mobility during clinical rounds and makes it viable for use in telehealth, fieldwork, and home care scenarios. This emphasis on portability ensures that ScopeFace integrates seamlessly into daily practice without adding burden or complexity.

Finally, the paper provides a critical discussion of ScopeFace's broader implications in modern healthcare. It explores how the device could enhance diagnostic certainty, facilitate remote second opinions, and support medical education by allowing students to replay and study real-world auscultation sounds. Collectively, these contributions demonstrate how ScopeFace leverages simple but powerful technologies to modernize a core clinical practice and make auscultation more precise, shareable, and accessible.

## 2 Related Works

Stethoscope technology has progressed from analog acoustic devices to digital stethoscopes capable of capturing, amplifying, visualizing, and analyzing cardiopulmonary sounds. This shift is motivated by the growing demand for improved diagnostic precision, better listening ability, remote monitoring, and integration with machine learning models to help clinicians make more informed diagnoses.

Traditional acoustic stethoscopes have minimally progressed since first being introduced and are limited by ambient noise, user variability, and lack of data storage. Digital stethoscopes address these limitations by incorporating technology like microphones, piezoelectric sensors, or capacitive diaphragms for high-fidelity signal acquisition [13]. To mitigate frictional noise and background interference during auscultation, many commercial devices now incorporate active noise cancellation. Devices like the Thinklabs One Digital use capacitive membrane-based technology to isolate low-frequency signals and amplify them up to 100×, connecting directly to earbuds for improved acoustic isolation, costing around $499 [10]. Additionally, the 3M Littmann CORE Digital Stethoscope offers up to 40× sound amplification and integrated active noise cancellation, retailing at approximately $330 [1]. Although Littmann is often regarded as a premium brand in clinical practice, prior studies have shown that higher cost does not necessarily correspond to superior diaphragm sound quality at clinically relevant frequencies when compared to more affordable alternatives [2].

Recent advances in embedded AI have made it more feasible to add intelligent features such as diagnostic capabilities directly within the digital stethoscope to help clinicians detect abnormalities [5]. Digital stethoscopes have shown strengths in early diagnosis of cardiovascular disorders like murmurs and arrhythmias [11]. FDA-cleared devices like the Eko CORE and CORE500 incorporate AI-based algorithms to detect atrial fibrillation and low ejection fraction in real time [6]. These devices also allow users to visualize waveforms on the stethoscope head and access recording, annotation, and patient history features via a companion mobile app, which requires an additional subscription fee. Wearable and mobile stethoscope platforms, like StethoMe and Eko DUO, support remote auscultation and asynchronous consultations, which gained attention during the COVID-19 pandemic [12].

Digital stethoscope performance is typically assessed in terms of frequency response, sensitivity, noise tolerance, and diagnostic reliability. A study by Ang et al.[3] evaluated four commercial digital stethoscopes and found significant variability in signal quality across clinical and teleauscultation conditions. Other works have evaluated the ability of digital auscultation systems to assist in emergency or low-resource settings, emphasizing the importance of reliable hardware and seamless integration with clinical environments [8, 9].

Since the use of stethoscopes is deeply embedded in clinical practice, end-user acceptance plays a critical role in the adoption of new digital alternatives. A recent study evaluated a prototype digital stethoscope with active noise cancellation and concluded that while users liked the AI functionalities, they faced challenges with respect to reliability and training [7]. Additional works have emphasized the need for ergonomic design, integration into electronic medical records, and compatibility with mobile health infrastructure to ensure practical deployment within established clinical workflows[4].

## 3 Technical Material

### 3.1 Project Requirements

This section outlines the requirements for ScopeFace hardware needed to meet the minimum viable product goals set by the team. These requirements can be divided into functional and nonfunctional requirements.

*3.1.1 Functional Requirements.* The functional requirements of the product relate to its core behaviors and usage. These requirements are likely customer-facing and are heavily driven by customer research.

- **Acoustic Signal Acquisition**: The device must be able to acquire and record audio in the frequency range associated with heart and lung sounds (20 - 2000 Hz).
- **Playback**: The device must be able to transmit the data to the frontend so that the user can playback audio.
- **Portability**: The device must be compact and lightweight so that the user can carry it consistently.
- **Durability**: The device must be durable and capable of withstanding drops, slight water exposure, and day-to-day wear.

- **Simplicity**: The device must be simple to use and require minimal setup and intervention from the user to operate.

*3.1.2 Non-Functional Requirements.* The non-functional requirements relate to its quality and performance. These requirements are critical for user adoption and usage.

- **Extended Battery-Life**: The device must be able to operate for significant periods of time without access to a power supply.
- **Robust Communication**: The device must be able to transmit audio with consistency and repeatability with minimal data loss during transmission.

## 3.2 Design Decisions

The project requirements drove many major design decisions related to hardware selection and integration.

*3.2.1 Microcontroller.* The microcontroller was the most important component of the system. It relates to all functional requirements. The team considered a series of microcontrollers and evaluated their abilities when related to cost, documentation, connectivity, and performance. The ideal microcontroller would be low-cost, well documented, have WiFi/Bluetooth connectivity, and have strong processing capabilities. The options considered were the ESP-WROOM-32, ESP32-S3, RPi Pico, and Arduino Nano ESP32. To evaluate the potential microcontroller, a decision matrix can be utilized. A scoring scale between 1 and 5 was utilized, with 1 being the best score and 5 being the worst.

| | ESP-WROOM-32 | ESP32-S3 | RPi Pico | Arduino Nano ESP32 |
|---|---|---|---|---|
| Cost | 2 | 2 | 1 | 5 |
| Documentation | 1 | 2 | 3 | 5 |
| Connectivity | 1 | 1 | 3 | 1 |
| Performance | 4 | 1 | 5 | 3 |
| Totals | 8 | **6** | 12 | 14 |

Based on the decision matrix, the ESP32-S3 was chosen. The ESP-WROOM-32 was chosen as a secondary fallback since the ESP32-S3 has less documentation and is a newer device. The ESP32-S3 is low-cost, has ample documentation, has WiFi and Bluetooth capabilities, and is capable of onboard AI inference. Once the ESP32-S3 was evaluated, a smaller form factor device called the XIAO ESP32-S3 was used in its place to fit the consolidated footprint of the stethoscope. This device uses the same SoC as the ESP32-S3, with the only difference being the form factor of the microcontroller.

*3.2.2 Audio Hardware.* The next major component of the device is the audio hardware. A major decision was to decide between an analog and a digital microphone. An analog device provides the benefit of custom circuitry and filtering, however, suffers from issues like power supply interference and the time to build out the electronic system to acquire data. A digital device is simpler to implement and often has built-in filtering and an analog-to-digital converter (ADC).

To make this decision, common analog and digital microphones were chosen for testing. The MAX4466 analog microphone and the INMP441 digital microphone were chosen for comparison. During testing, the MAX4466 audio suffered significantly from power supply interference and required more components to acquire data from since the ESP32-S3 does not have a built-in ADC. On the other hand, the INMP441 performed well and was quick to set up with minimal additional components needed. For this reason, a digital microphone was chosen for the hardware setup.

Next, a microphone needed to be identified that would be able to acquire frequencies related to heart and lung sounds. This requirement led to the DMM-4026-I2S-R microphone which has a frequency range of 20 Hz - 20 kHz. The benefit of this microphone is that the company makes an evaluation board consisting of hardware filtering and amplification, as well as a simple interface identical to the INMP441, leading to a simple swap of the microphones.

The Inter-IC Sound ($I^2S$) protocol is a standard communication protocol for digital audio. This is well documented with the ESP32, making it simple to implement and allowing for simple integration with the microcontroller.

*3.2.3 Communication.* Based on the chosen microcontroller, there were two options for communication: WiFi and Bluetooth. These options were evaluated on their throughput, data transfer rates, range, and power consumption. WiFi excelled in all areas except for power consumption, where it would be more costly than Bluetooth. Apple places limitations on Bluetooth devices through iOS, which would significantly impact the user experience. This includes only allowing Bluetooth Classic servers to appear in the traditional Bluetooth page in Settings, as well as throttling throughput and data transfer over Bluetooth to the device. Given these drawbacks and the complexity to implement Bluetooth, WiFi was chosen. The device connects to WiFi and transmits data directly to the storage database.

*3.2.4 Human-Hardware Interfaces.* User interaction with the device was an important design decision to ensure a seamless user experience. The two aspects of this interface were the onboard RGB LED and a user-triggered button. The LED was chosen over an onboard display for simplicity and a cleaner design. The LED illuminates in three different states: initialization, recording, and uploading. The states of the RGB LED can be seen below, corresponding to the illuminating color.

| **State** | **Color** |
|---|---|
| Initialization | Yellow |
| Recording | Red |
| Uploading | Green |

The button was required to be low-profile, require a deliberate press to trigger, and have a sufficient tactile response for users to discern that the button has been pressed. A panel-mount momentary push button was chosen that would integrate into the physical design. This allowed for the button to be triggered and return to its original state automatically, and fit in a compact design.

*3.2.5 Storage.* The device needed to store audio data so that the frontend could retrieve the data to serve to the user. Major cloud providers such as AWS, GCP, and Firebase were considered. These services were fairly evenly matched given the extensive support for each on the ESP32. The deciding factor was the connectivity with authentication, as well as credits provided for use of Firebase services.

## 3.3 System Design

The overall system was laid out using the components of the system mentioned above.

*3.3.1 Design Schematic.*

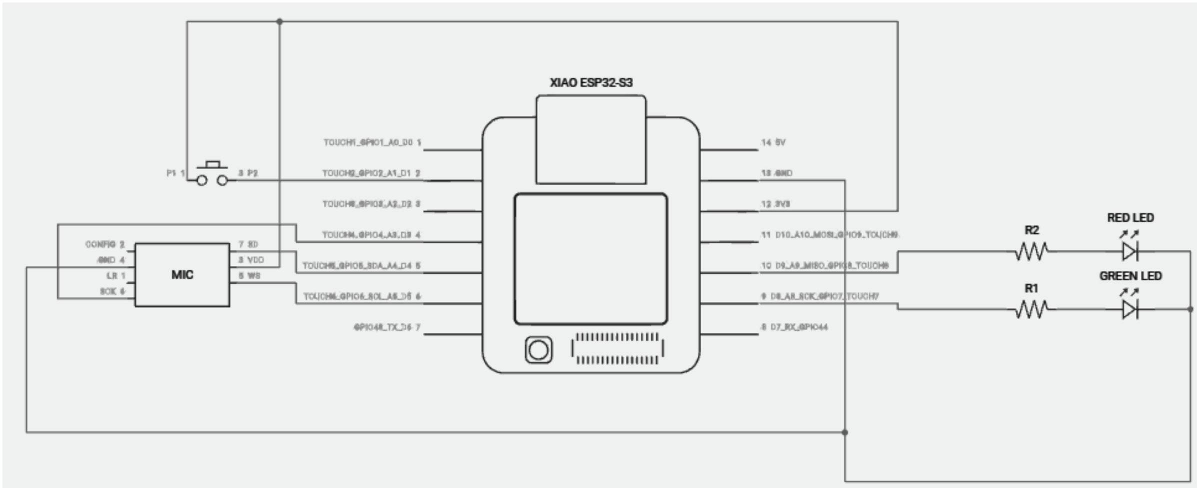


Fig. 1. ScopeFace Electronics Schematic

*3.3.2 System Workflow.* The system workflow can be seen in Figure 2. There are two main workflows triggered by a button press. The first is the initialization of the device as it is brought out of sleep mode. The second is recording and uploading of the heart sounds from the device to the database.

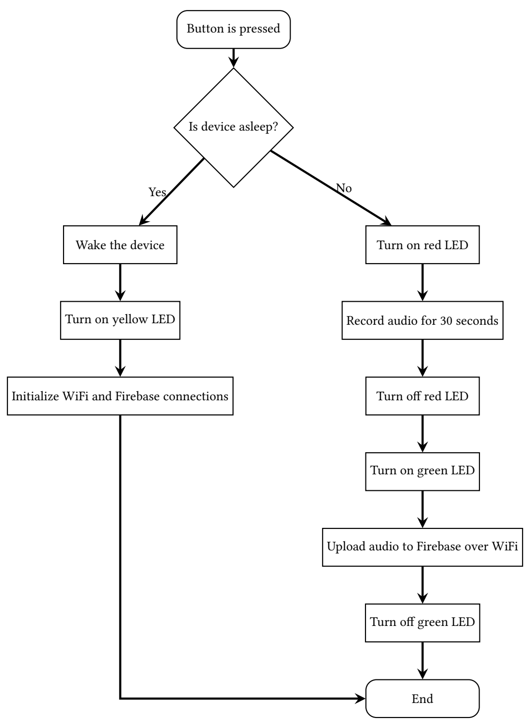


Fig. 2. System Workflow Diagram

### 3.4 Device-App Data Transfer

To enable transfer of audio recordings from the hardware to the mobile application, we adopted a decoupled architecture using Firebase Cloud Storage as an intermediary layer. This design allowed the hardware and software teams to work independently and reduced the complexity of direct device-to-app integration.

*3.4.1 Data Transfer Process.* When a user presses the button and completes a recording, the device uploads the resulting .wav file to Firebase Storage using a uniquely generated session ID and timestamp. The file is stored in a predefined bucket path accessible to authenticated users through Firebase's SDK.

On the iOS side, when the user loads new sessions, the app checks for new session metadata associated with the signed-in user's device ID. When a new recording appears, the app retrieves its metadata and generates a download URL, allowing the user to stream the audio file directly without needing to download it locally.


Fig. 3. ScopeFace Data Transfer Visualization

*3.4.2 Key Benefits.* This architecture offers several key benefits:

- It enables asynchronous data transfer between hardware and software.
- It improves the potential marketability of ScopeFace by allowing for easier integration with existing electronic health record systems
- It decouples the embedded firmware from the mobile UI, reducing risk during development.
- It allows users to access recordings from any authenticated device with minimal latency.

## 3.5 Embedded Program

The program onboard the device is written in C++. The decision between MicroPython and C++ was made based on documentation as well as execution speed. The design is broken up into modules that each serve different functions. The modules can be broken up as follows:

- Initializations
- I2S protocol helpers
- Firebase helpers
- Main acquisition program
- Timeout check

*3.5.1 Initialization.* The initialization functions serve to initialize the core components of the program. This includes WiFi, Firebase connection, peripherals (i.e. LED, button, PSRAM, buffer allocation), $I^2S$ initialization, and pin the main audio program to a core.

*3.5.2 $I^2S$ Protocol Helpers.* The $I^2S$ protocol helpers allow for the initialization of the audio interface with the digital microphone. This initializes the pins for the serial clock, data communication, and word select. Additionally, it sets key parameters for audio acquisition including the sampling rate, channels, and bits per sample.

We know that to mitigate issues like aliasing, missing low frequencies, and audio compression a few parameters had to be set. The sampling rate can be determined using the following equation based on the Nyquist-Shannon Sampling Theorem:

$$F_s = 2 \times \text{Highest Frequency}$$

The highest frequency of interest are the lung sounds which are approximately 2000 Hz. The default sampling rate is 44100 Hz, which is well above the calculated minimum sampling rate needed of 4000 Hz. This mitigates aliasing and missing of lower frequencies such as heart and lung sounds. The WAV file format was chosen to prevent audio compression issues as it is a lossless, uncompressed audio format.

*3.5.3 Firebase Helpers.* The Firebase helper functions allow the ESP32 to communicate with Firebase. To perform this communication, the Firebase-ESP-Client library is used. This library allows the device to connect to Firebase Storage and utilize a callback function to determine the data upload and communication status.

*3.5.4 Main Acquisition Program.* The main acquisition program is how heart and lung sounds are recorded and uploaded from the device. To ensure a simple recording is taken at a time, a simple synchronization mechanism is utilized. A flag is set when the button is pressed, which is checked to begin a recording. If the flag is already on, a new recording will not start. Once the flag is set, the WAV header is written to the buffer. Then, audio is acquired from the microphone for a pre-set amount of time (in this case, 30 seconds). Once recording is complete, a filename is generated, and the buffer is uploaded to Firebase Storage using the filename as the path.

*3.5.5 Timeout Check.* A key feature of the device is a timeout sequence. To ensure the device can operate for extended periods of time without charging, the device must timeout and enter sleep mode. ESP natively supports light and deep sleep modes. Given that no state needs to be saved between power cycles, deep sleep is the best choice. Additionally, deep sleep consumes much less power, making it ideal for an edge device like ScopeFace.

To bring the device out of sleep mode, the button was used as the pin it was connected to was an RTC GPIO, which remains active during deep sleep. When the button is pressed, the device runs its setup and waits on the next button press to begin recording.

The timeout check is done through an activity flag. This flag is set every time the recording sequence is activated and at the end of upload. This ensures that the device does not timeout and go into sleep mode while recording or uploading. The timeout occurs after two minutes of inactivity based on the activity flag. After this timeout occurs, the WiFi is disconnected, and the device is put into deep sleep.

## 3.6 Audio Signal Processing

The audio uploaded from the device needed to be processed to remove noise and to better isolate the lower frequencies of the heart and lung sounds. To do this, a Firebase Function was deployed to process the newly received files in place. This function would also append to the metadata of the file, allowing for the data to only be retrieved if it has been processed.

To process the data, the first step was to apply a low pass filter to the WAV file. A cutoff frequency of 150 Hz was chosen to isolate the heart sounds and a cutoff frequency of 2000 Hz for lung sounds. Once this filter was applied, the audio was amplified with a gain of 20.0. After this, the file was written back to the storage bucket. Below is a diagram of an audio sample before and after processing. The spikes in the processed data show isolation of heart sounds as opposed to the consistent waveform of the unprocessed audio.

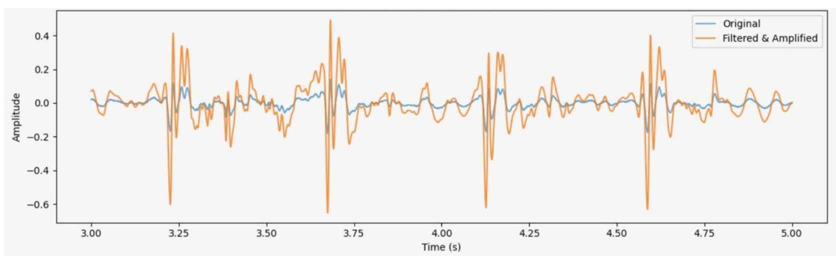


Fig. 4. Signal Processing: Original and Processed

## 3.7 Device Modeling

Once the electronics system was completely designed for ScopeFace, the next step was to design the enclosure to house these components.

*3.7.1 Physical Requirements.* The physical requirements of the device were based on the project requirements, as well as limitations imposed by the nature of the electronics system. These requirements are laid out below:

- The device must be compact and lightweight
- The device must be durable and able to withstand drops, slight water exposure, and day-to-day wear.
- The device interface must be simple and intuitive to use

Given the time constraints, limited resources, and cost constraints, the best option for manufacturing was 3D printing. To select a material, a few options were considered including PLA, PETG, Nylon, and Resin. Based on the properties of these materials, the best option was PETG, which was more suited for acoustic applications. This was based on previous research done to design a simple, open source digital stethoscope [2].

*3.7.2 Physical Design.* The design was created using an iterative process, incorporating user feedback in each iteration. Additionally, changes were made due to tolerancing issues and iterations in electronics hardware.

Fig. 5. ScopeFace Full Design

*Microphone and Diaphragm Holder.* The microphone and diaphragm holder is the base of the physical design. It contains standoffs in which heat-set threaded inserts are placed to secure the device together. Two standoffs are used to secure the internal structures (i.e. Battery Holder, Microcontroller Mounting Plate), and the other two are to secure the top cover to the device.

The microphone is press fitted into place, and the housing contains a relief slot in the event that the microphone needs to be removed. The diaphragm is attached the same way it would be on a traditional stethoscope, with the rubber ring being stretched onto the bottom surface of the part. The space between the microphone receiver and the diaphragm is conically shaped to better promote sound transfer between the surfaces.

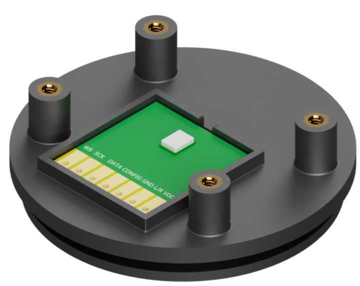

Fig. 6. ScopeFace Microphone and Diaphragm Holder

*Battery Holder.* The battery holder is designed to hold a 350 mAh lipo battery. The battery slides into place, and has an opening on its back side to route the power cable out to connect to the microcontroller power supply. The battery holder is also set lower than the standoffs to maximize space usage and sits just above the microphone.

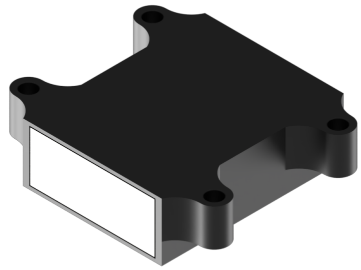

Fig. 7. ScopeFace Battery Holder

*Microcontroller Mounting Plate.* The microcontroller mounting plate houses the XIAO ESP32-S3 using a press fit design as well. The device is slid into place, elevated so that the power supply cables below the device can route outwards toward the battery holder. The plate also lines up with the top cover opening to allow for charging to occur via the USB-C port.

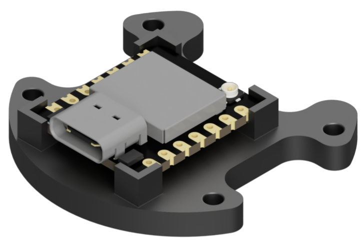

Fig. 8. ScopeFace Microcontroller Mounting Plate

*Top Cover.* The top cover houses the momentary push button, RGB LED, and provides an opening for charging of the device. This part fits over the internal components of the device and secures to two standoffs on the microphone and diaphragm holder. The device contains two locating features used to ensure a consistent assembly. The momentary push button is secured using a threading design and the LED is secured using a press fit design.

Fig. 9. ScopeFace Top Cover

### 3.8 Device Testing

Testing was done to ensure that the device met the non-functional requirements laid out in the Project Requirements section. This included repeatability testing to ensure no dropped data or connections over a significant number of attempts as well as calculations to determine the battery life of the device.

For the repeatability test, the device was set up to upload a sequence of 100 recordings, each 30 seconds. The goal was to record the upload time and any failures. The device had no recorded failures and an average upload time of 17.829 seconds. Below is a graph which displays the upload time distribution over the 100 recordings.

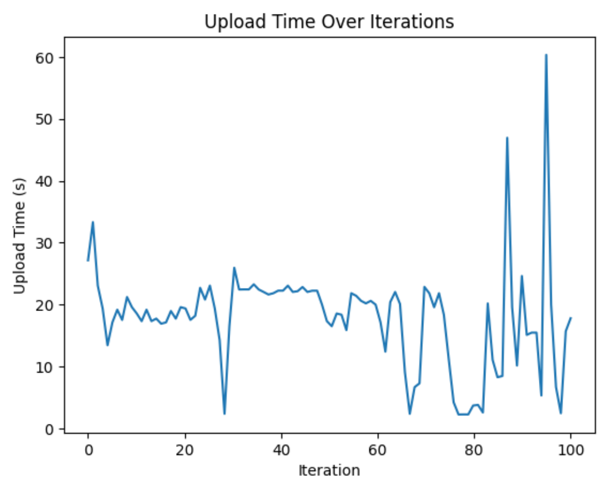


Fig. 10. Repeatability Test

The battery life was calculated using the battery capacity and documentation provided by Seeed Studio regarding the max power consumption in normal state and deep sleep states. The lipo battery has a capacity of 350 mAh and according to the documentation the max power consumption of the microcontroller is approximately 100 mA and 14 $\mu A$ for normal and deep sleep states respectively. We assume a situation where the user activates the device out of deep sleep each time they see a patient, and the device remains out of deep sleep for approximately two minutes per patient. This assumes that the user records the four main areas associated with auscultation. We assume on average the user sees twenty patients a day. From this, we can perform some basic calculations to determine the battery life of the device.

$$\text{Battery Life} = \frac{\text{Capacity}}{\text{Usage}} = \frac{350 \text{ mAh}}{0.014 \text{ mA} \times (24 - 2.667) \text{ hrs} + 100 \text{ mA} \times 2.667 \text{ hrs}} = 1.311 \text{ days}$$

This shows that the device would last 1.311 days under high load. This scenario is an overestimate of usage, and the battery life would likely last longer than the calculated value.

## 4 Milestones - Hardware

At the start of the quarter, we set out to develop a fully functional and portable digital stethoscope as our MVP that was capable of capturing, amplifying, recording, and visualizing heart and lung sounds. We successfully met this goal, and below we detail the key hardware milestones achieved and the pivots made along the way to realize this outcome.

- **Deliverable 1 (Week 3): Hardware Schematic Diagram (Amrita)**
  *Proof:* Document with all major hardware components mapped out
  *Milestones:*
  - Determine and list major components needed based on final product requirements
  - Determine connections and additional interfacing components
- **Deliverable 2 (Week 3): Hardware Selection Decision (Shayan)**
  *Proof:* Bill of materials based on hardware schematic diagram
  *Milestones:*
  - Research and downselect microcontrollers
  - Test microphones for heart and lung sound frequencies
  - Decide on filtering approach (analog and/or digital) to isolate target frequency bands for heart (typically 20–150 Hz) and lung (100–1000 Hz) sounds
- **Deliverable 3 (Week 5): Initial Product Prototype (Shayan)**
  *Proof:* Prototype with fully integrated hardware components
  *Milestones:*
  - Order and assemble hardware
  - Wire components, create documentation and diagram
- **Deliverable 4 (Week 5): Initial Prototype Code (Amrita)**
  *Proof:* Run and test hardware code, graphs showing frequency response before and after filtering, video demonstration
  *Milestones:*
  - Microphone data acquisition and processing
  - Recording functionality
  - Implement digital filtering to isolate relevant sound frequencies
  - Transmit data from microcontroller to device

  *Pivots:* We initially explored analog filtering for heart and lung sound isolation, we ultimately transitioned to post-acquisition digital filtering. This pivot was driven by hardware limitations and our need for tunable filtering and amplification parameters during development.
- **Deliverable 5 (Week 7): Design and Build Casing (Shayan)**
  *Proof:* Pictures of assembled product, tested with casing
  *Milestones:*
  - CAD design, dimensioning and tolerancing
  - 3D print and assemble casing
  - Test full product functionality
- **Deliverable 6 (Week 7): Transmit Data to Database (Amrita)**
  *Proof:* Documented process with test data samples in database
  *Milestones:*
  - Communication from device to external gateway
  - Route data from gateway to database
  - Repeatability testing to ensure consistency and quality

  *Pivots:* Our original plan to use Bluetooth transmission encountered roadblocks due to iOS compatibility restrictions. We pivoted to using Wi-Fi, which introduced a new challenge: signal interference caused by internal adhesive used to mount the Wi-Fi range extender. This required hardware debugging and layout adjustments to restore stable data transmission.

- **Deliverable 7 (Week 9): Refine Prototype Design (Shayan)**
  *Proof:* New CAD design focused on ruggedness and ergonomics
  *Milestones:*
  – Feedback from medical professionals
  – Improve usability, add drop and water resistance
  *Pivots:* Based on initial user feedback, we intended to enhance the device's ruggedness. However, due to time and prototyping constraints, we were unable to implement full drop and water resistance in the final MVP.
- **Deliverable 8 (Week 9): Create MVP (Amrita)**
  *Proof:* MVP in hand
  *Milestones:*
  – Reprint casing, assemble product
  – Full functionality testing
- **Deliverable 9 (Week 10): Product Documentation (Shayan and Amrita)**
  *Proof:* Full documentation, specifications, and use cases
  *Milestones:*
  – Compile and write product documentation
  – Specifications sheet with safety and use case info

## 5 Milestones - Software

### 5.1 Overview

At the beginning of the project, we defined a series of software milestones to guide the development of ScopeFace, with a focus on core functionality, modular architecture, and user-centered design. Over the course of the quarter, we successfully achieved most of these goals, adapting to shifting constraints while maintaining progress toward a functional, end-to-end system.

### 5.2 Foundational Decisions and Setup

We began by finalizing our technology stack and documenting key decisions through Architectural Decision Records. Our frontend was developed using Swift and SwiftUI to ensure a smooth native iOS experience. Firebase was selected for user authentication and cloud data storage, and a Python backend using FastAPI was implemented to support future expansion.

To support maintainability and team collaboration, we established a structured GitHub repository, clear setup guide, and automated CI/CD pipelines for both frontend and backend code. These workflows provided automatic linting and test execution on every pull request, reinforcing code quality throughout the development process.

We created a set of high-fidelity Figma wireframes that guided the development of each user-facing component for the MVP, particularly the login flow and recording interface. We also created wireframes extending beyond the MVP, including session history views, FAQs, and account pages. These designs ensured visual consistency and helped us align on navigation, interaction patterns, and content hierarchy.

In parallel, we produced detailed software architecture diagrams and database schematics using Miro for both our MVP and extended system. These diagrams informed team task division, aided component prioritization, and supported communication across our team.

### 5.3 Core Feature Development

Core application features, including user authentication, audio recording management, session review, note-taking, and tagging, were implemented and integrated across the frontend and backend. These features formed the backbone of the app's interactive experience and established a solid base for future iterations. The fully functional MVP app interface supports reviewing sessions, renaming them, and attaching clinician notes.

Notably, we also implemented scrollable waveform visualizations for recorded audio. This allows users to visually navigate through a session and jump to specific time points, offering improved control over playback and enabling more efficient and targeted listening. The waveforms enhance the user experience while supporting clearer communication and annotation of auscultation data.

*5.3.1 Integration Strategy.* To simplify integration between the software and hardware components, we adopted Firebase Storage as a shared interface. Rather than directly coupling frontend and hardware logic, the hardware uploads .wav recordings to Firebase, where they can be accessed and streamed by the iOS app. This decoupled architecture helped us streamline integration by reducing tight interdependencies and allowing each subteam to work asynchronously. However, this decision limited our ability to have realtime audio playback.

### 5.4 Deferred Milestones and Challenges

Some goals, such as full end-to-end test coverage and deployment via TestFlight, were intentionally deferred. TestFlight was deprioritized due to the cost of an Apple Developer account, and instead, we conducted internal usability testing using the Xcode simulator. While unit tests were implemented for core functionality, broader integration and end-to-end tests were pushed due to time constraints and remain a target for future development.

## 6 User Feedback

To evaluate the usability and clinical relevance of ScopeFace, we conducted remote demonstration sessions with three healthcare providers. During these live Zoom meetings, we showcased the functionality of both the hardware device and the mobile application, providing a walkthrough of the recording process as well as core app features such as login, playback, session review, tagging, and waveform navigation.

Overall, users responded positively to ScopeFace's simplicity, portability, and collaborative potential. Participants highlighted ScopeFace's potential to support clinical education, cross-team collaboration, and second-opinion workflows.

> *"ScopeFace would boost my confidence in my assessments. It's easy to use, and I love that I can playback recordings for my care team. As a new nurse, being able to confirm recordings with other clinicians would help my learning and development."*
> — Claudia, RN, Scripps Health

> *"ScopeFace would make it easier to communicate about my patients to my colleagues. In the future, real-time listening would be a game changer!"*
> — Sriram, MD, UCI & Anschutz Medical

> *"ScopeFace would help me confirm findings and get second opinions on subtle murmurs. In the future, I'd love the option to send recordings to doctors remotely."*
> — Kristen, PA, Brigham and Women's Hospital

These interviews reinforced key design choices and surfaced valuable insights for future iterations, including real-time listening capabilities and secure sharing for remote consultation. Notably, our first round of feedback came from close friends and family members, which likely adds a more positive bias to our results. Going forward, user feedback will continue to play a central role in refining ScopeFace, and we will prioritize gathering input from a broader, unbiased group of clinicians.

# 7 Future Work

## 7.1 Overview

If ScopeFace were to be developed into a market-ready product, several areas of further work would be necessary to support real-world clinical deployment and adoption. These directions are informed by our current prototypes, system design, and user feedback.

## 7.2 Deployment and TestFlight

Currently, the app is only functional on the iOS simulator due to the cost of registering for an Apple Developer account. A necessary next step would be to obtain a developer license, distribute builds through TestFlight, and begin testing the app on physical devices. This would enable broader usability testing and facilitate eventual submission to the App Store for pilot use.

## 7.3 HIPAA Compliance

Before any clinical rollout, ScopeFace would need to be brought into full HIPAA compliance. We would switch to a HIPAA compliant database and adjust patient visibility on the app interface if needed.

## 7.4 Mobile App vs. EHR System Integration

A key architectural decision moving forward will be whether to continue developing ScopeFace as a standalone mobile app or to refactor it for integration with existing electronic health record (EHR) systems such as Epic. Integration could improve data centralization, streamline clinician workflows, and support institutional security policies. However, it may also constrain ScopeFace's usability in home care, educational, or resource-limited settings.

## 7.5 Patient Assignment, Record Management, and Sharing

A major extension of the current system involves associating recordings with individual patients as well as allowing clinicians to send recordings and notes to their care team to support collaboration continuity of care. Additional features would include:

As referenced above, integrating richer patient-oriented features would be dependent on the decision to develop a mobile app rather than integrate with an existing EHR system. Features may also be abandoned based on HIPAA compliance.

## 7.6 Machine Learning for Auscultation Support

In the long term, ScopeFace would incorporate machine learning models to assist with interpreting heart and lung sounds. This might include:

- On-device or server-based audio analysis
- Detection of murmurs, wheezing, or irregular rhythms
- Summaries of algorithmic findings to assist clinicians

These tools would be designed to support (not replace) clinical judgment, especially for early-career practitioners.

## 7.7 Live Playback

One important aspect of a stethoscope is the listener's ability to make adjustments in real time. To accommodate this, ScopeFace would need to have the ability to stream the audio data live as well as record it for upload. This would come in the form of connecting a Bluetooth headset or a set of wired headphones to the device. To do this, a digital to analog converter (DAC) and amplifier would need to be incorporated into the design.

## 8 Conclusion

ScopeFace contributes to the modernization of medical tools to improve the quality of patient care. By addressing common limitations of traditional stethoscopes, ScopeFace enables healthcare providers to make more confident and accurate diagnoses, all in a small and lightweight form factor. An integral contribution of the project is its integration with a companion app that stores recordings, helping bridge gaps between care teams and facilitating better coordination among healthcare providers. The incorporation of signal processing techniques improves the clarity of auscultation recordings by minimizing background noise, allowing for easier and more reliable interpretations of heart and lung sounds. Overall, ScopeFace enables healthcare providers to operate with confidence, ensuring they don't miss a beat.

Throughout the development process, significant progress was made on both the hardware and software components of ScopeFace. The successful implementation of signal acquisition, signal processing, and a compact, portable design highlights the device's strong potential for practical use in clinical settings. Additionally, the full system integration demonstrates the feasibility of creating an end-to-end system that supports diagnosis, documentation and communication within healthcare ecosystems.

Looking ahead, there are several promising directions for continued development of ScopeFace, including enhanced user interfaces and enabling live playback through Bluetooth connection or wired headphones, and the incorporation of machine learning classification algorithms to assist with the interpretation of heart and lung sounds. ScopeFace's commitment to delivering smarter, affordable, and more accessible diagnostic tools lays a strong foundation for future innovation and impact in patient care.

## References


[1] 3M Littmann. 2023. 3M Littmann CORE Digital Stethoscope. https://www.littmann.com/en-us/home/f/b5005222000. Accessed: 2025-06-12.

[2] Spencer Chambers Mahmoud Al-Alawi Kliment Yanev Tarek Loubani Alexander Pavlosky, Jennifer Glauche. 2018. Validation of an effective, low cost, Free/open access 3D-printed stethoscope. *PLoS ONE* (2018).

[3] Y. Y. Ang, L. R. Aw, V. Koh, and R. X. Tan. 2023. Characterization and Cross-Comparison of Digital Stethoscopes for Telehealth Remote-Patient Auscultation. *Medicine in Novel Technology and Devices* (April 2023). doi:10.1016/j.medntd.2023.100256

[4] A. Deniz-Garcia, H. Fabelo, A. J. Rodriguez-Almeida, G. Zamora-Zamorano, M. Castro-Fernandez, M. D. P. Alberiche Ruano, T. Solvoll, C. Granja, T. R. Schopf, G. M. Callico, C. Soguero-Ruiz, and A. M. Wägner. 2023. Quality, Usability, and Effectiveness of mHealth Apps and the Role of Artificial Intelligence: Current Scenario and Challenges. *Journal of Medical Internet Research* 25 (May 2023), e44030. doi:10.2196/44030

[5] Motaz Faroq A. Ben Hamza and Nilam Nur Amir Sjarif. 2024. A Comprehensive Overview of Heart Sound Analysis Using Machine Learning Methods. *IEEE Access* 12 (2024), 117203–117217. doi:10.1109/ACCESS.2024.3432309 Preprint available via ResearchGate.

[6] Eko Health. 2023. CORE500 Stethoscope FDA Clearance and AI Capabilities. https://ekohealth.com.

[7] Tanvir Joarder, S. N. B. K. Tune, A. A. Islam, A. M. Shahunja, M. I. Haq, S. Bhowmik, K. M. Shahunja, K. Rahman, A. Chisti, E. L. Winch, and S. El Arifeen. 2023. End-user acceptability of a prototype digital stethoscope to diagnose childhood pneumonia: A qualitative exploration from Sylhet, Bangladesh. *BMC Digital Health* 1 (2023), 26. doi:10.1186/s44247-023-00027-y

[8] Clare Koning and Adinet Lock. 2021. A systematic review and utilization study of digital stethoscopes for cardiopulmonary assessments. *Journal of Medical Research and Innovation* 5, 2 (July–December 2021), 1–11. doi:10.25259/JMRI_2_2021

[9] Rohini Kotha, Caleb Streitmatter, Andrew Serdiuk, Nasrin N. Aldawoodi, and Robert S. Ackerman. 2024. Cardiac Remote Monitoring Devices and Technologies: A Review for the Perioperative Physician and Telemedicine Providers. *Cureus* 16, 2 (February 2024), e53914. doi:10.7759/cureus.53914

[10] Thinklabs Medical LLC. 2002. Electromagnetic diaphragm design for digital stethoscope. US Patent US6406412B1.

[11] Jia Jie Seah, Jie Zhao, Yujie Wang, and Heow Pueh Lee. 2023. Review on the Advancements of Stethoscope Types in Chest Auscultation. *Diagnostics (Basel)* 13, 9 (April 2023), 1545. doi:10.3390/diagnostics13091545
[12] StethoMe. 2021. Smart wireless stethoscope for home use. https://stethome.com.
[13] Swathi Swarup and Amgad N. Makaryus. 2018. Digital stethoscope: technology update. *Medical Devices: Evidence and Research* (2018).

## A Appendix A: Bill of Materials

| **Part** | **Quantity** | **Unit Cost** | **Amount** |
|---|---|---|---|
| XIAO ESP32-S3 | 1 | $16.99 | $16.99 |
| 3.7V Lipo Battery 350mAh | 1 | $7.99 | $7.99 |
| JST-PH 2.0 Male and Female Connector Cable | 1 | $0.69 | $0.69 |
| Momentary Push Button | 1 | $4.50 | $4.50 |
| M2 x 12 mm | 2 | $0.08 | $0.16 |
| Heated Threaded Inserts | 4 | $0.22 | $0.88 |
| M2 x 25 mm | 2 | $0.18 | $0.36 |
| Microphone | 1 | $2.66 | $2.66 |
| RGB LED | 1 | $0.09 | $0.09 |
| Stethoscope Diaphragm Cover | 1 | $8.99 | $8.99 |
| | | **Total** | $43.31 |

Table 1. ScopeFace Bill of Materials