Thesis Portfolio

Creating a Temperature-Tracking Door Locking System for COVID-19
(Technical Report)

Face2Gene: Using Facial Recognition to Aid in Diagnosing Rare Genetic Disorders
(STS Research Paper)

An Undergraduate Thesis

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Amanda Rein

Spring, 2021

Department of Electrical and Computer Engineering

Table of Contents

Sociotechnical Thesis

The application of automation technologies into the healthcare industry has revolutionized medical systems and processes. It is no longer necessary for patients and their loved ones to rely upon their doctor(s) to memorize all symptoms and cues of every disease, nor do they have to wait for their doctor(s) to comb through literature in order to reach a proper diagnosis. A major example of a technology eliminating this massive inefficiency is Face2Gene, a mobile application that utilizes facial recognition to scan a person's face for characteristic facial cues and match them to a massive catalogue of data on genetic disorders. The freely-available application returns a list of most likely diagnoses and their probabilities within seconds, alleviating the workload for patients and their physicians. Patients can also have far more faith in the accuracy of their diagnoses with this application of machine learning, too, because of the massive amount of data supporting the finding(s).

Another major example of how automation technologies can maximize efficiency and reduces peoples' workload is automatic temperature screening devices. During the COVID-19 pandemic, many workers were tasked with spending their days standing at the entryway to an office/indoor space, scanning entrants' temperatures and asking them a few questions about recent exposures to the virus before admitting them. The Capstone project outlined in this prospectus proposes a design for a system that automatically senses someone approaching a doorway, scans their temperature from afar, then unlocks the door if the person's temperature is in a healthy range. Nobody is allowed to enter a space if they have a fever and thus could be contagious. The system also sends these temperatures to a Web dashboard that allows stakeholders to monitor the traffic into/out of an indoor space, through the number of people admitted and the number denied.

This STS research paper and Capstone project are therefore tied together by their ability to address inefficiencies in healthcare. Both of them involve automating processes that would otherwise waste the time of health care professionals/officers of public health. While the STS research paper provides a persuasive argument that the application Face2Gene has been molded by society to fit users' needs, the Capstone project involved the creation and testing of a physical prototype for the temperature-sensing door locking mechanism. As such, their deliverables are quite different. The Capstone project taught me how to design and create a tangible product beyond a written report. The STS research paper, on the other hand, taught me how to employ an STS framework to make a case for a particular viewpoint.

Creating a Temperature-Tracking Door Locking System for COVID-19

A Technical Report submitted to the Department of Electrical and Computer Engineering

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Amanda Rein
Fall, 2020

Technical Project Team Members
Andy Hui
Greg Vavoso
Jack Shefer
Matthew Bain

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Signature ____ Date ___11/24/2020__
Amanda Rein

Approved ____ Date ___11/24/2020__
Harry Powell, Department of Electrical and Computer Engineering

Abstract

This project's final deliverable is an automated health screening door connecting temperature measurements to the door lock. When a prospective entrant approaches a door outfitted with this system, they will place their wrist under a mounted box, triggering a motion detector and awakening a non-contact temperature sensor. The temperature sensor will take a temperature reading from the person's wrist, and that reading will be interpreted within a MSP432 microcontroller to determine whether to unlock the door. If the temperature is above a healthy range, the MSP432 will communicate to the door that it should lock and prevent that person's entry to the space until they can return with a normal temperature reading. At the same time that the MSP432 makes the locking decision, the temperature reading is sent to a Raspberry Pi server, which stores all temperature readings in a database. This database is utilized to render a web dashboard for stakeholders to this door locking system. Anyone interested in seeing the number of people allowed to (or prevented from) entering the space, or other statistics on door usage, may find this information on the website.

Background

The COVID-19 pandemic has wreaked havoc on the global economy. Commercial businesses, such as retail, restaurant establishments, and entertainment venues, have been forced to reopen to avoid bankruptcy, despite the serious health risks. Consumers also itch to return to a sense of normalcy, leave their homes in which they have been confined for months and support local businesses. Educational institutions are also reopening their doors amidst the pandemic, as students desire to remain on-track in their education. While reopening the economy, supporting local businesses, and keeping children in classrooms are all noble pursuits, they have created a dire need for ways to discern individual's health before they potentially infect the others in these spaces.

Ideally, individuals would be tested for COVID-19 before they are permitted to enter a shopping mall, restaurant, gym, movie theater, or classroom, but this is not possible for multiple reasons. For starters, there are no existing testing options that can return results in a short enough time to make this feasible. Secondly, there has been a shortage of COVID tests worldwide, making it incredibly difficult and expensive for corporations or institutions to acquire even enough tests to give students before they return to a residential school, for example. The University of Virginia, a prominent and powerful institution in the United States, is not even able to test students more than a few times per semester unless the student requests a test with a credible case for being exposed to the virus.

For these reasons, many businesses and schools have resorted to more generic screening metrics to gauge a person's overall health. A multitude of businesses, such as doctors offices and gyms, have defaulted to measuring the temperature of individuals seeking entry. Body temperature is thought to be correlated with contraction of COVID-19 because fevers are one of the major symptoms of the virus. This requires the business to staff at least one employee at the entryway to manually screen temperatures. Not only is this a financial concern; this poses a health risk for the person who must interact with potential entrants, getting close enough to take their temperatures. Even with non-contact devices, they would need to be relatively close to the entrant for an accurate scan. The Absolute Chicanery Capstone team saw the inefficiency of this system that had been adopted by many and wondered why an automated option hadn't filled this gap in the market.

After some research on prior art, it is clear that the systems currently automating temperature scans before entry are very expensive. In fact, Advanced Imaging estimates that the average temperature scanning kiosk costs between $2,000 to $7,000 [1]. A primary example of one of these expensive kiosks is Meridian's Personal Management Kiosk, which starts at $2,895. This kiosk, like most others of its kind, includes facial recognition technology so that the temperature readings can be associated with a person's identity [2]. While this level of identifiability may have its proper use cases, it is a very controversial technology. Many people are uncomfortable with the idea of this information being collected about them without consent.

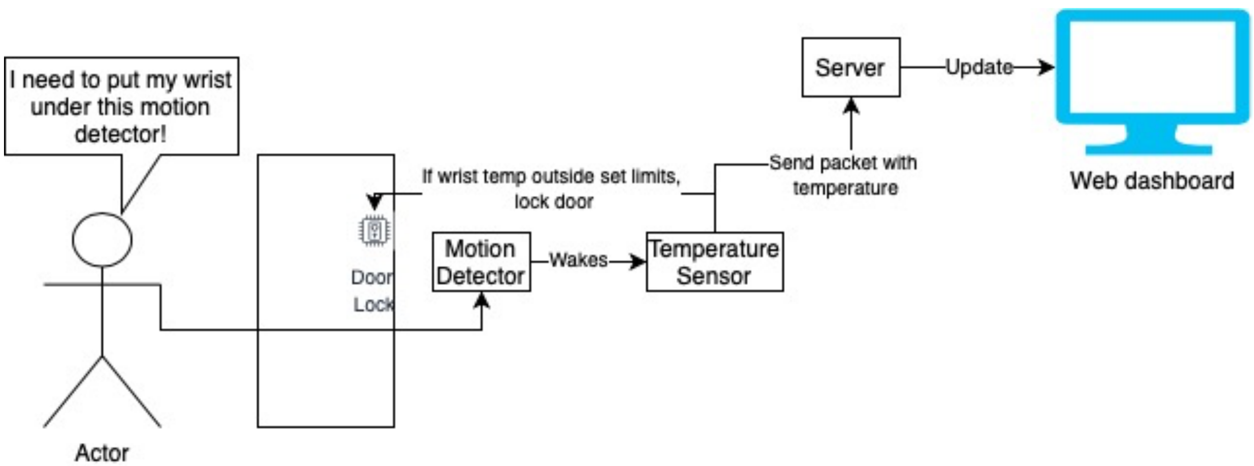

Figure 1. System Flow Diagram from User Perspective

To address this gap in the market, the team presents a system to automate the temperature screening process, preventing individuals with higher-than-average temperatures from entering a space without privacy concerns involving storing personal data. This system also promises to be cost-effective when produced at scale, in comparison with competitors on the market today.

Constraints

Design Constraints

This project was under a considerable time constraint, as it had to be completed in one semester. The team had under four months to complete the design, from start to finish, while also juggling other classes and responsibilities as full-time students.

The Capstone projects were required to include either a myRIO or MSP43X microcontroller, limiting the CPU capabilities for the system. Additionally, each project had to include a custom PCB design, created via Multisim and Ultiboard. The manufacturing constraints on these boards included that they could be no more than 30 square inches in size, with a maximum of 50 drill holes per square inch and 10 millimeter drill hole size. The manufacturers were also unable to drill any slats or non-circular holes into the PCB, which affected which components were ultimately selected. Additionally, any lines or traces had to be 5 millimeters in thickness.

The team was also constrained by a lack of a crimper; WWW Technologies (a company with which the team worked to solder the smaller components onto the PCB) did not possess one, and the costs of purchasing one were far out of the budget.

There were constraints involved with utilizing a Raspberry Pi for the server, as well. The Raspberry Pi has 1 GB of RAM and a 32 GB disk, constraining how much data the server can store at any given time.

There were additional constraints associated with creating the housing for this system, as well. 3D print designing required precision down to the millimeter for component through-holes;

our team was forced to opt out of printing screw mounts into the design since printing errors were too large for our precise dimensions. Additionally, closing the unit atop the shell required additional hours of printing time as well as support material so we had to add acrylic sheets on top to finish the assembly.

Economic and Cost Constraints

This Capstone project was capped to a maximum budget of $500. This was the primary cost constraint for the design. While this seems like an ample budget, it did not permit for purchasing a highly accurate, medical-grade non-contact temperature sensor. This caused issues with inaccurate temperature readings during testing.

If this product were taken into high volume production, many costs would be eliminated. Around $90 was used on testing and revisions. This included multiple board revisions and part revisions.

Another significant cost came from using development boards. For example, the MSP432 LaunchPad board was approximately $24, while the MCU chip itself would be less than $5 in volume. A cheaper chip could also have been used to further reduce costs. The Raspberry Pi 3B+ ($35) could similarly be removed from the system and replaced by a computer or server already present in the building.

External Standards

There are a few key external standards that this project takes into account. Firstly, safety standards for this project include FCC regulation of radio frequency devices. This project contains both intentional and unintentional radiators. The intentional radiators consist of the selected Bluetooth module and the Raspberry Pi. Both these have been approved by the FCC for use in short range communications contexts [3]. The rest of the circuitry and digital devices are categorized as unintentional radiators. We have aimed to keep these devices outside of the 30 MHz - 960 MHz frequency range so they are exempt under 47 CFR 15.101 (b) [4]. The MSP and various circuitry will stay below this range while the Raspberry Pi is clocked above it.

Fire safety is something that should be considered whenever implementing a locking mechanism on an outward facing door to a building. Safety is an important consideration when designing doors that are in route of egress. Every three years, the Virginia Uniform Statewide Building Code (VUSBC) adapts the International Building Code (IBC) updated by the International Building Code Council. Section 1010.0.9.3 of the IBC states that:

> "Where electrical systems that monitor or record egress activity are incorporated, the locking system shall comply with Section 1010.1.9.7, 1010.1.9.8, 1010.1.9.9, 1010.1.9.10 or 1010.1.9.11 or shall be readily openable from the egress side without the use of a key or special knowledge or effort." [5]

The VUSBC alters this requirement by stating the following [6]:

"Change Item 2 of Section 1010.1.9.3 of the IBC to read: 2. In buildings in occupancy Groups B, F, M and S, the main exterior door or doors are permitted to be equipped with key-operated locking devices from the egress side provided:

1. The locking device is readily distinguishable as locked.

2. A readily visible durable sign is posted on the egress side on or adjacent to the door stating: THIS DOOR TO REMAIN UNLOCKED WHEN THIS SPACE IS OCCUPIED. The sign shall be in letters one inch(25 mm) high on a contrasting background.
3. The use of the key-operated locking device is revokable by the building official for due cause."

To satisfy these conditions for our project, we procured a fail safe electric strike. The strike allows the door latch to be released if pulled from inside the building, but it controls the door latch if attempted to be pulled from the outside of the building. Since this is "readily openable" from the inside, or "egress" side, the door meets inspection [7]. To increase the fire safety of the system, the fail safe strike unlocks when there is a loss of power - allowing emergency personnel into the building. Also, the high side switch is configured to only lock the door when the microcontroller is supplying a signal- otherwise if the microcontroller is faulty, the system remains unlocked even if the building's power is operational [8].

Aside from safety standards, extra care must be taken when dealing with health data and medical records. While this device does not fall under the jurisdiction of the Health Insurance Portability and Accountability Act (HIPAA), the US Department of Health and Human Services (HHS) guidelines for HIPAA compliance were useful when evaluating the ethical consequences of device installation [9].

Additionally, standards exist for mechanical assembly of devices like this system. Particularly, under IPC standards [10], we classify our project as a Class 1 General Electronic Product where function is the primary goal. Further, we achieved a Level C design producibility where significant manual assembly is tolerated. This is an area that would require improvement if the device were to be sold commercially.

Finally, this project will rely on standardized patterns for both internal and external communication. Internally, communication between the MSP and the Raspberry Pi will use Bluetooth [11] technology due to its compatibility with the Raspberry Pi and it's low energy requirements. The Bluetooth specification is maintained by the Bluetooth Special Interest Group (SIG). Each door-mounted MSP device will use a Bluetooth component (controlled via UART [12]) to send data to the Raspberry Pi hub. Then, interaction with the Pi will be done over IEEE 802.11 WiFi [13] because of its wireless nature and use in the HTTP and SSH protocols that follow. Management and development for the Pi will be accomplished using SSH [14] for security and wireless access. Management and access to the temperature sensor will also rely on the I2C communication standard [12]. Externally, users will consume temperature data using the HTTP 1.1 protocol [15] for access from any available device and browser on the network. Use of these communication standards will ensure interoperability of hardware components (Bluetooth emitters, receivers) as well as software tools (SSH clients, browsers).

Tools Employed

The primary computer aided design for this project was done with the National Instruments suite of Multisim 14 [16] and Ultiboard [17]. These tools allowed rapid schematic creation and PCB footprint design. Additionally, the plastic housing to contain our device was designed with AutoDesk AutoCAD 2021 [18].

Assembly and testing took place in UVA's NI Lab which provided various additional resources. The Virtual Benches and accompanying digital multimeters and logic analyzers were critical for powering and testing our prototypes. The soldering equipment was important for final assembly of our PCB. Lastly, UVA's mechanical engineering department provided access to 3D printing capabilities that ultimately built the plastic housing for our device.

Software creation came with it's own set of tools. Throughout the project, software was tracked in a Git [19] repository for version control and easy code sharing among team members. The embedded programming made use of Texas Instruments' Code Composer Studio [20] development environment and accompanying compiler [21]. During debugging, PuTTY [22] was used to read from the microcontroller's serial port. To access the Raspberry Pi during development, OpenSSH [14] provided remote terminal access.

The final software also relied on some important libraries. The embedded code made extensive use of the Texas Instruments DriverLib [23] suite and the FreeRTOS [24] real-time operating system. The Raspberry Pi relied heavily on the Raspberry Pi OS [25] operating system as well as the PySerial library [26], the SQLite database system [27], the Flask web framework [28], and the Apache2 web server [29].

Ethical, Social, and Economic Concerns

Environmental Impact

The system manages environmental impact by reducing overall energy usage with the incorporation of a motion detector. While the motion detector was not a necessary component to the system, it makes the system far more energy-efficient because it avoids the need for the temperature sensor to be constantly scanning for temperature readings. The motion detector will only wake the temperature sensor once motion is detected, reducing the temperature sensor's energy consumption. That being said, the system does consume some level of energy constantly, as the motion detector must be able to detect motion at all times.

Significant portions of the project can be disassembled and re-used in other contexts later on. Components like the Raspberry Pi, the MSP432 microcontroller, and the digital sensors can all be unplugged with relative ease for use in future Capstone projects or in other academic pursuits. Similarly, if one of these parts breaks, it can be removed from the system and replaced by an identical but functional part; that way, the entire system does not need to be thrown away. For that broken part, electronics recycling is freely available in the Charlottesville area through the McIntire Recycling Center, for example [30]. This is more complicated for components soldered to our PCB, but these can still be disassembled and re-used to a certain extent.

Sustainability

The automated health screening door lock device is a highly sustainable system as its individual components can easily be replaceable or updated while still maintaining a consistent and effective result. The motion detector allows the other components of the system to stay dormant until usage is required meaning overall longevity of the system is high. All the power in our project comes from electrical wall outlets, both on the microcontroller and the server side, so there is no reliance on disposable batteries.

In terms of the electrical components, the main concern is the constant power consumption of the motion detector. However, the passive infrared sensor is very inexpensive and requires low power so it won’t wear out easily. The Bluetooth module also poses a sustainability question because of the range of communication. The HC-05 chip used for our Bluetooth communication, though, exceeds expectations and allows for 3 Mbps of modulation with a complete 2.4GHz radio transceiver and baseband. Additionally, the module auto-reconnects within 30 minutes when disconnected from out of range connections.

Finally, the mechanical casing of the system is highly durable and can maintain drops due to the lightweight of the device. The main shell is formed from 3D printed thermo-plastic filament, which is sustainable but causes a waste issue due to its un-recyclability [31]. In addition, the clear 5 mm acrylic sheet enclosing the system is also durable but also imposes the same issue on un-recyclability similar to the thermo-plastic, which often goes straight to landfills [32]. Finally, initial prototypes failed to protect the internal components from rain and other weather conditions, which presents problems for system durability but could be further developed before production.

Health and Safety

The design for this system prioritized health and safety, largely by ensuring that the door locking mechanism follows fire and safety regulations. The lock we chose is specially shaped to allow the strike to leave if the change in the latch comes from the inside [7]. Therefore, regardless of power availability, a user can exit the building. The strike needs power supplied continuously to remain locked, and the high side switch adds redundancy in case the system fails but the building's power is still working. The high side switch controlling the lock also has a fault detection functionality that allows different errors to be detected, reported, and fixed in a timely manner [8]. While this default unlocked state would mean that people with fevers may be allowed to enter the space if the system is not functioning properly, this would just be similar to the system not having been present rather than posing a risk to those who could be trapped inside.

Another relevant safety consideration is the spread of contagious disease, which the system hopes to decrease. The system was designed to be hands-free and even make it unnecessary for users with fevers to touch the doorknob. When a fever is detected, the LEDs facing out of the housing indicate that the temperature was invalid and that the door remained locked. As such, potentially contagious users never need to make physical contact with the device or the door to prevent the spread of germs.

There is a health risk to collect an artificially high temperature reading, as it can cause anxiety for the user to see that their temperatures may be out of range when they are actually healthy. This is a risk that is worth the reward of a system that maintains the health of those inside, though, especially since the risk is mitigated when the user can wait a moment until they have cooled off (i.e. if they approached the system after jogging or partaking in an activity that spiked their body temperature) and then try again to enter. There is no limit to how many times a user may request entry to the space via the system.

Manufacturability

Manufacturability is an important consideration for any potential consumer product. During prototyping however, our emphasis was focused on completing the functionality of the device under the given time constraints which is why we aimed for an IPC Level C design producibility where manual assembly is permitted. That being said, many of the PCB components could be adjusted for more manufacturing-friendly surface mounts and configurations. Additionally the MSP432 Launchpad could be swapped out for the surface-mount package, or even a less powerful microcontroller. The external casing would be quite amenable to manufacturing at high volume by switching from 3D printed plastic to injection molding or other techniques. The limiting factor for manufacturing efficiency would likely be the manual placement of the sensors and the installation of the wiring between them.

Ethical Issues

Unlike some of its competitors mentioned previously, this system does not utilize facial recognition technology, nor does it collect any personal information on the individuals associated with each temperature scan. As such, the primary ethical concerns in storing user data are negligible. There do exist corner cases for identifiability, however. For example, if timestamps shown on the system's dashboard could be correlated with entrance times observed by a bystander or video camera. These considerations may be mitigated in the future by restricting dashboard access or implementing an encryption scheme.

Additionally, it is possible that the placement of the wrist scanner disadvantages certain people in the population. For example, a person without arms or a toddler that cannot reach the temperature scanner may have trouble gaining entrance to the space.

Another potential area for concern is how skin pigmentation affects temperature measurement reliability in our device. Producing a system which disproportionately allows members of certain ethnicities to enter a building over other ethnicities would present serious ethical barriers to the success of the project.

Intellectual Property Issues

Patent US10692596B2 introduces a health kiosk that includes a "physiological measurement apparatus connected to the computing unit". This kiosk includes a weight scale, blood pressure monitor, and a heart rate monitor. That being said, some of them allow for the addition of blood pressure cuffs, EKGs, ECGs, blood glucose measurers, and thermometers. A web server stores the measurements taken by the kiosk, then can transmit them on the network as a URL or HTTP cookie of an HTTP request. The kiosk includes a screen for more user interaction/communication, such as an LCD that is a touch screen. It also may include keyboard, mouse, microphone, or other external tools that the user may utilize to communicate with the kiosk. Some versions may include technology for automatic user identification, such as a camera for facial recognition or a fingerprint scanner [33].

Patent US20180158555A1 is a patent for another medical kiosk, though this one focuses on tele-medicine and check-in for medical appointments or prescription pick-up. The kiosk includes video conferencing capabilities for tele-medicine appointments, and the kiosk includes user input that "can include one or more components selected from the group consisting of a key pad for identification and/or data entry, … motion sensor, sound sensor temperature sensor, …" [34].

Patent US9256719B2 is a patent for a smart mirror that includes an infrared camera to record body temperature, among other sensors/devices collecting biometric data. This system also includes facial recognition software [35].

Since all of these patents do not include a door locking mechanism controlled by the temperature readings, Absolute Chicanery's Capstone deliverable is patentable. Additionally, the lack of any cameras or other ways of gathering user identities makes the system stand out from existing patents.

Detailed Technical Description of Project

Overall System

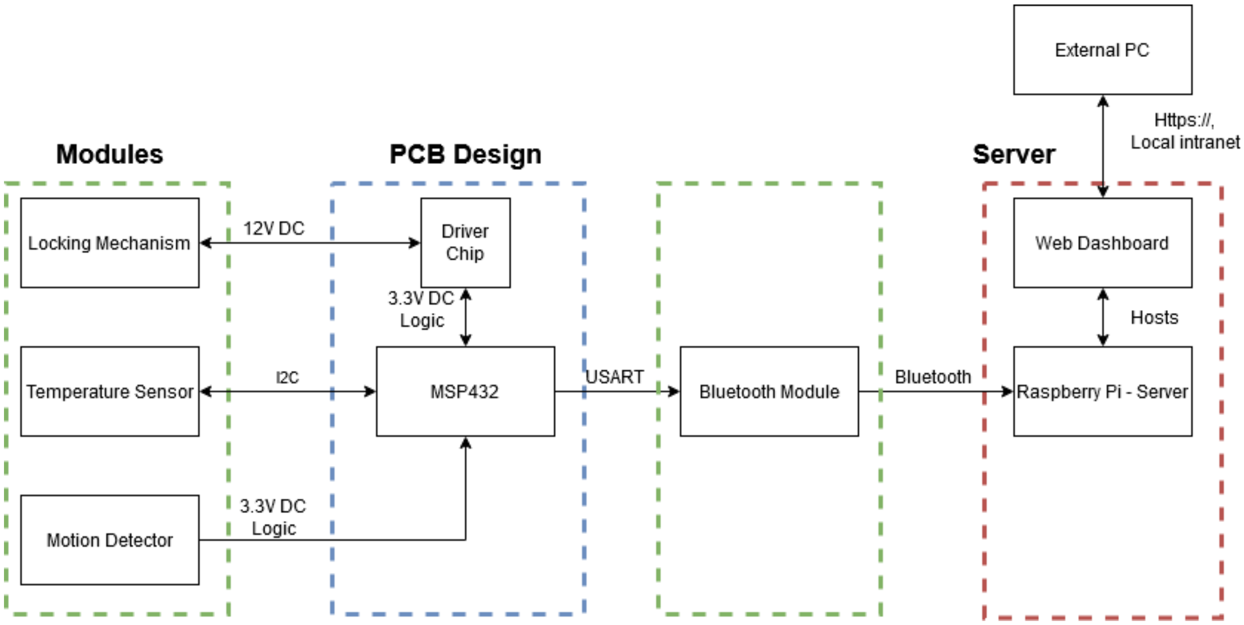

Figure 2. High level system architecture.

There are 7 main areas of the project: the power supply, the door lock, motion detector, temperature sensor, Bluetooth module, embedded system, and server/web dashboard. Figure 2 shows how the different systems interact as well as the expected communication signals used on each connection. The different modules in green interact with the MSP432 microcontroller in blue and the server architecture in red. These colors provide conceptual distinction among subsystems. The MSP432 acts as a central hub that controls the logic and communication of all modules. The MSP432 sends data to the Bluetooth module in order to communicate with the raspberry pi server and display the dashboard.

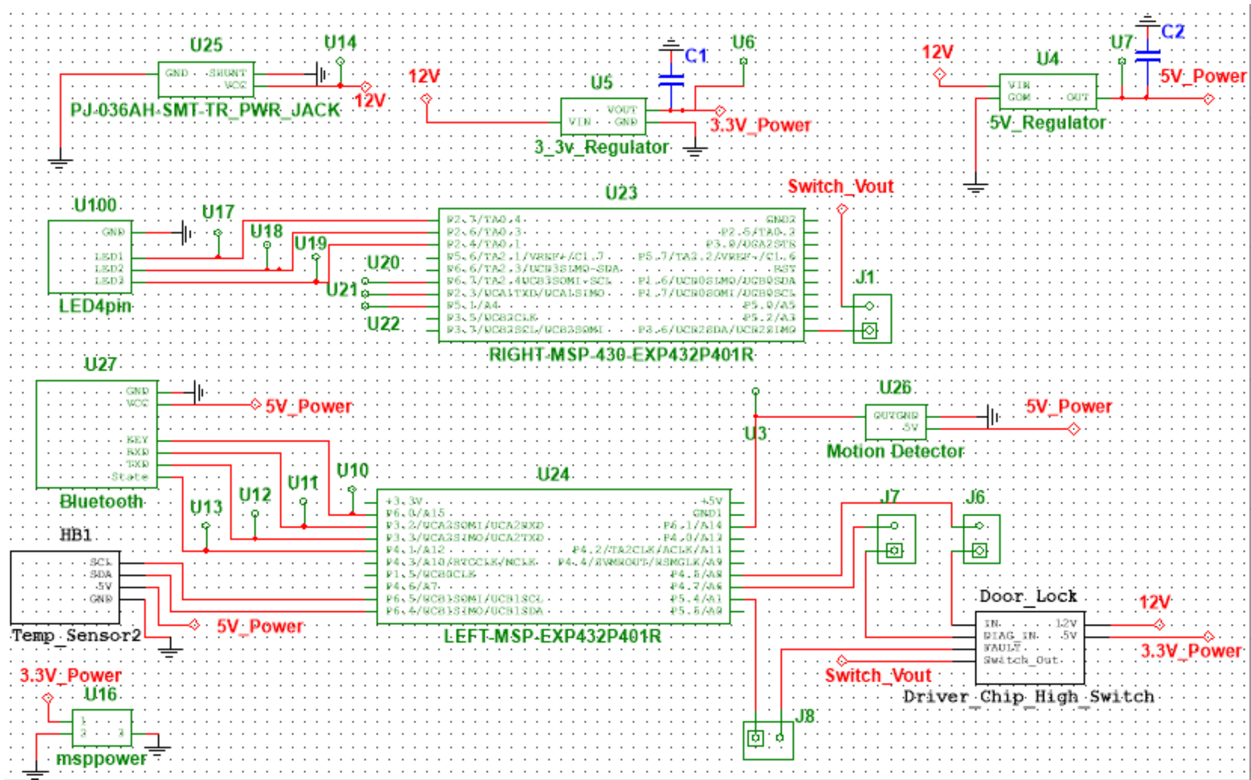

Figure 3. Top level schematic diagram.

Figure 3, provides a more detailed schematic description of how the subsystems managed by the microcontroller interact with one another, including test points and MSP432 pin assignments. The power supply subsystem runs along the top of the schematic while the connections to the MSP432 run around the outside. The LED board is connected to the right hand side of the MSP432, while the other modules connect to the left hand side, for organization. The Bluetooth module and motion detector connections are represented by MOLEX on-board connectors, as well as the lock and temperature sensors once inside their respective hierarchical blocks. The two hierarchical blocks are examined more closely in their sections, but the main connections are shown in Figure 3.

Power Supply

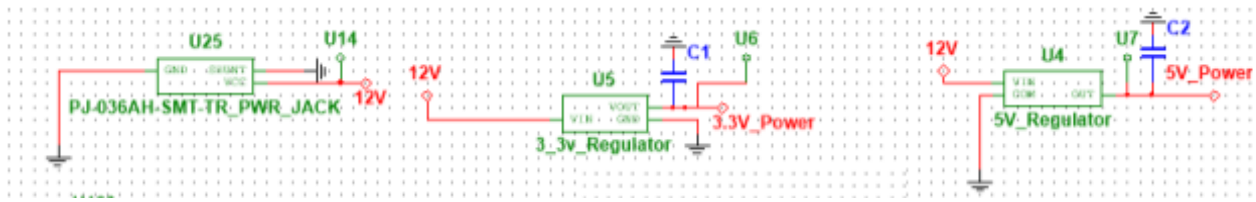

Figure 4. Schematics for power supply subsystem.

Different subsystems in this project have different power needs, adding complexity to how we must power the device. The lock strike itself requires 12 V when locked, which led us down the path of powering the device from the wall outlet rather than using a battery. A 12 V power adapter for a laptop computer was used to plug into the wall and convert to DC. This adapter was plugged into our board via a surface-mount power jack. Two linear voltage regulators were used to step the 12 V supply down to both 5 V and 3.3 V for different components in the system. The Bluetooth module, temperature sensor, and motion detector were all powered from the 5 V supply whereas the microcontroller itself used the 3.3 V supply.

Current requirements also play a role in the power supply design. The Bluetooth module has a maximum current requirement of 30 mA and the temperature sensor has a maximum current draw of 16 mA [36]. The motion detector has the least amount of current draw with a 3 mA maximum [37]. The microcontroller was allocated 75 mA based on the recommendations form its data sheet, and finally the door lock and driver chip use a combined 285 mA of current at maximum [38]. Therefore, the regulators need to be able to handle these maximum current draws. The 5 V regulator is rated for a typical load of 150 mA whereas our system only needs 68 mA of current, total [39]. The 3.3 V regulator can handle 250 mA where the system only needs 91 mA [40]. The off- board LEDs were powered directly by the microcontroller's digital pins since their current requirements of 2 mA were low enough [41]. In total, the system needs about

410 mA of current to operate effectively. The DC power connector we used can handle 5 A, so it is more than sufficient [42].

Door Lock

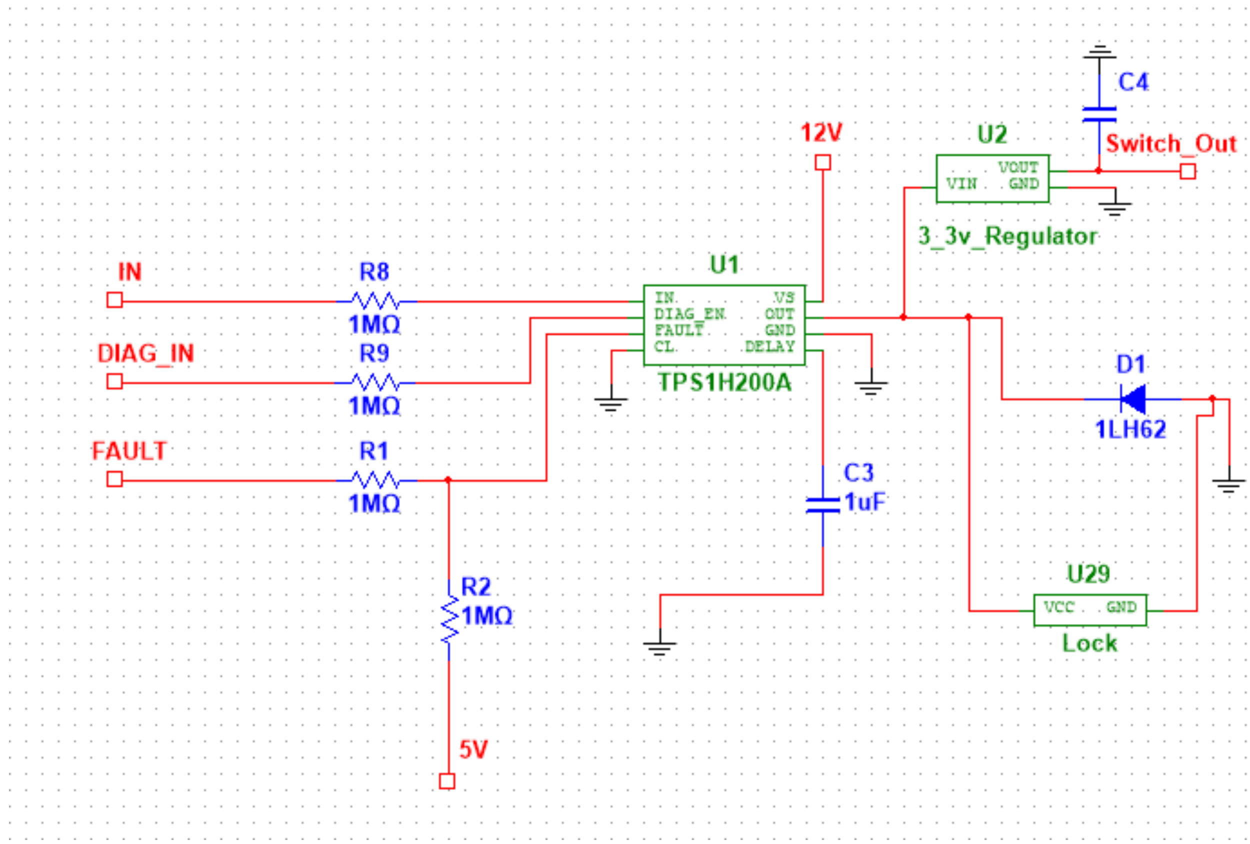

Figure 5. Switch and door locking subsystem schematic.

The locking mechanism for the system is a fail safe lock designed for general use. It is controlled by a chip on the custom manufactured PCB that intercepts the 12 V power supply and connects through the 12 V VCC connection to the strike. This can be seen in the schematic in Figure 5 and the data flow diagram in Figure 6**.** The chip is controlled by the MSP432 through the IN and DIAG EN 3.3 V digital logic controls. When IN is high, current from the 12 V supply is connected to the strike, locking the door. When IN is low, the 12 V power supply is disconnected from the strike, and the strike unlocks. The DIAG EN control enables fault protection on the chip and diagnostic information to be relayed to the microcontroller through the FAULT pin, which can pinpoint a fault according to a truth table.

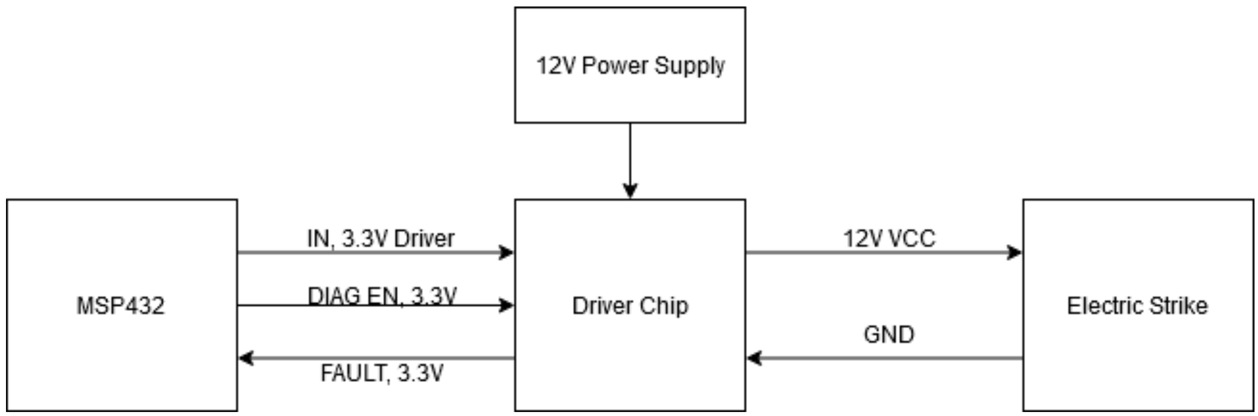

Figure 6. Component interaction in the door lock subsystem

Motion Detector

The motion sensor removes the need for our temperature sensor to continually collect temperature readings, which would require an unnecessary amount of power. The motion sensor also ensures that the system only collects temperatures to send to the web dashboard and utilizes it for deciding whether or not to lock the door *when someone has approached*. It would be problematic for both power consumption and system usability if the device constantly sent ambient temperature readings over Bluetooth to the server.

We decided to utilize a Passive Infrared (PIR) motion sensor, as it will allow us to constantly sense motion in a low-power manner. PIRs are composed of pyroelectric sensors that detect movement via elevated levels of infrared radiation. The sensor is split into two halves (technically, it is composed of two individual sensors). If both sides detect the same amount of infrared, it is considered idle. Then, when anything warm (i.e. a human wrist) passes the sensor, it will detect the movement as an imbalance in infrared readings between the two halves.

As for the components that make up this device, there is a JFET transistor at the core, which has low noise and built-in buffering for high impedance. Atop this JFET is a lens, a very low-cost plastic dome that drastically increases the detection area. Using optics concepts, this lens condenses a large field of view into a small one, which is ideal for the tiny sensor that takes the infrared readings. The lens is split up into sections, each of which is a Fresnel lens (a type of lens that is made up of concentric rings to concentrate into a thin beam).

The PIR has 3 digital pins: power, ground, signal. The PIR has digital output (high or low voltage readings; high when motion is detected). The power is 3-5 V DC, though 5 V is recommended and what we used in the system. When the PIR detects motion, the signal pin goes to “high”, outputting 3.3 V. The PIR used in this project is from Adafruit and has a trimpot that allows users to easily adjust the sensitivity. The sensitivity range for this device is up to 6 meters (20 feet), with a 110-70 degree range of detection. This will adequately cover anyone who is coming up to the door to scan their wrist temperature. There is also a trimpot to adjust how long the output stays high after a motion detection occurs (10-20 second range).

In our project, the signal pin is repeatedly polled when the device is idle. When the signal pin goes high and motion is detected, processing flow is yielded to the temperature sensor to take and compute temperature measurements. This represents the rising edge of the signal pin but we also take advantage of the falling edge as it transitions back to ground. This is the point where the system is ready for another scan so the door is re-locked and the readiness is indicated on the LEDs.

Temperature Sensor

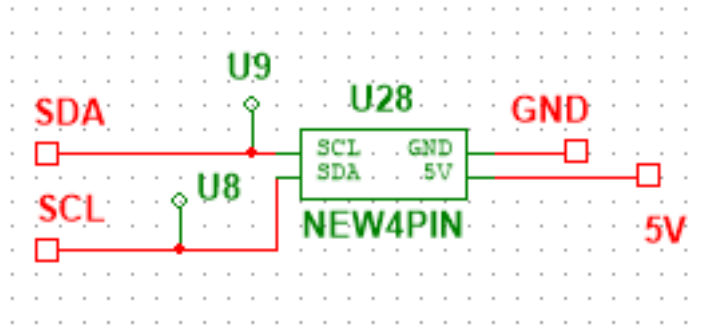

Figure 7. Temperature Sensor Circuit Schematic Pins

Our device incorporated the MLX90614 infrared temperature sensor manufactured by Melexis. Specifically, the temperature sensor was purchased through a third-party vendor which came as a GY-906 module chip. This contactless temperature sensor module uses infrared rays to measure the human body temperature by converting the signal from the sensor to a digital value and communicating to the microcontroller using an I2C protocol. Specifically, the contactless temperature measures the intensity of the emitted infrared energy radiated from a living being, which is directly proportional to the object's temperature as stated in the Stefan-Boltzmann law. The overall device consists of two chips integrated into a single sensor – one sensing the IR temperature and the processing unit converts that signal to a digital value.

There are three main functions that allow the temperature sensor to operate. First is signal processing. The device is controlled by an internal state machine, which controls the measurements and calculations of the object and ambient temperatures. Second is amplification. A low noise low offset amplifier with programmable gain is implemented for amplification of the IR sensor voltage. Third is the SMBus 2 Wire Protocol. The module comes with digital SCL and SDA pins for use with SMBus. These pins, shown in Figure 7 above, are decoded by an integrated circuit on the chip and allow the microcontroller to alter the state of the sensor and request temperature readings.

This device can output a temperature using either PWM output or SMBus (a protocol that extends I2C). Our system uses the temperature sensor as the slave in a master-slave I2C configuration, while the MSP432 microcontroller acts as the controlling device. When a read command is sent over the shared bus, the temperature sensor responds with a 15 bit temperature value that can be converted to a Kelvin temperature using division in the microcontroller logic. Figure 8 illustrates the I2C communication that takes place for a read command as observed with a VirtualBench Logic Analyzer. The red labels indicate information written to the bus by the microcontroller whereas the yellow labels indicate data written by the temperature sensor. Table 1. Example calculation of Fahrenheit temperature based on sensor response. below illustrates the steps taken to decode this sensor response to a Fahrenheit temperature.

Figure 8. Screen capture of I2C communication with temperature sensor.

High Data Byte	39h = 57
Low Data Byte	F5h = 245

Full Temperature Payload	39F5h = 14,837
Decoded Temperature, Kelvin	14837 / 50 K = 296.74 K
Decoded Temperature, Fahrenheit	75.46 F

Table 1. Example calculation of Fahrenheit temperature based on sensor response.

There are many key specifications that make this temperature sensor unique. The ambient temperature range for detection is between -40°C to 125°C whereas the object temperature range is -70°C to 382.2°C. The field of view of the sensor lens sits at 90 degrees and the measurement resolution to determine accuracy is at 0.02°C. In terms of power, the operating voltage is at 53.3 V after using a voltage regulator from the 12 V AC adapter input, and the supply current is at 1.5 mA. Finally, it is important to note that the Melexis datasheet indicates that the ideal distance between sensor and object is between two and five centimeters.

As will be discussed later on, we had issues with the accuracy of the sensor likely relating to the large field of view. While ambient room temperatures were always measured accurately, the temperature varied drastically based on how far away the wrist was from the sensor. A measurement taken with a wrist 1 cm from the sensor could be significantly different from a measurement taken with a wrist 5 cm away.

Bluetooth Module

Our selected Bluetooth module exposes a UART interface which can send bytes one-by-one to paired devices. Once a 15-bit Fahrenheit temperature has been decoded from the temperature sensor’s encoding, a two-byte Bluetooth packet is constructed with a re-encoded lower resolution temperature. The first byte contains metadata and error correction information and the second byte contains the temperature payload. Temperatures are encoded as a 7-bit unsigned integer representing the number of tenths of a degree over 95 °F. Thus, in 7 bits we can encode temperatures from 95.0 - 107.6 degrees Fahrenheit in degradations of 0.1. Table 2, below, specifies the formal packet structure.

Byte 1: Header								
Bit #	7	6	5	4	3	2	1	0
Use	0	Lock Fault		Allow	-	Checksum		
Byte 2: Data								
Bit #	7	6	5	4	3	2	1	0
Use	1	Encoded Temperature						

Table 2. Bluetooth Packet Structure.

The most significant bit of each byte is used as a sequence number, differentiating header bytes (0) from data bytes (1). The “Lock Fault” bits store an encoding of the state of the lock. There are three potential failure cases that the driver chip can output, or these could be zero to indicate no errors. The “Allow” bit is set to 1 if the door was unlocked as a result of this reading and set to 0 if it remained locked. Bit 3 of the header byte is unused for now and always set to 0.

The “Checksum” bits store the number of bits set to 1 in the entire packet, modulo 8. This can be used to detect if there were errors in transmission, acting as an additional layer of reliability on top of the L2CAP Bluetooth protocol. The “Encoded Temperature” bits store the unsigned integer described above. To reconstruct the temperature, the following formula can be applied:

$$Temperature\ (F) = 95 + 0.1 * Encoded$$

These Bluetooth packets are then sent to the server’s Bluetooth receiver, where they are persisted for future use.

Embedded Programming

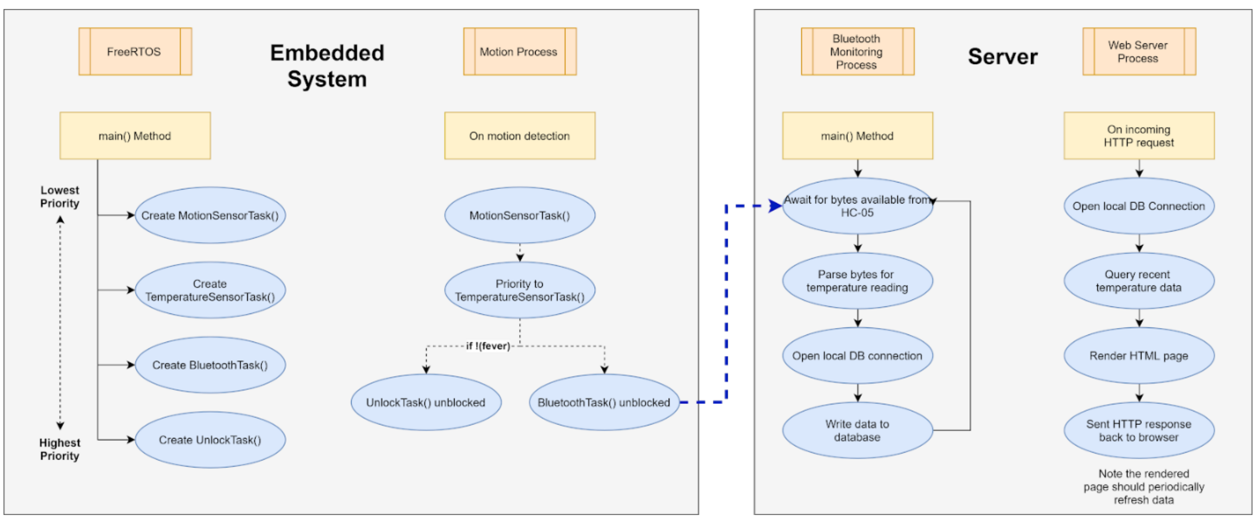

Figure 9. Embedded System diagram

The MSP432 uses a real-time operating system to manage the timing between different subsystems. The FreeRTOS kernel was chosen for its ease of use and reliability, as well as being distributed under the MIT open source license [24]. The FreeRTOS kernel was chosen over the TI-RTOS for its portability and level of documentation, allowing for other microcontrollers to be used should this project continue.

The FreeRTOS kernel uses the term “task” to define any code to be scheduled. The architecture of this project was to assign each subsystem a task. The following tasks were created: Bluetooth task, Lock task, Temperature task, and Motion task. FreeRTOS uses a priority based preemptive scheduling technique. In this project, semaphores were used to control flow from task to task.

Generally, the flow from task to task is as follows: motion, temperature, lock, then Bluetooth. The motion sensor task runs periodically until motion is detected. Once motion is detected, a semaphore is released to allow the temperature task to run. The temperature task then collects 128 temperature samples. The MSP432 uses its I2C module to send a request to access the temperature sensor’s RAM, which contains an encoded Kelvin object temperature. After successfully reading a temperature, the temperature task releases semaphores for both the lock task in and the Bluetooth task, in that order. The lock task determines whether or not to toggle the locking mechanism. The Bluetooth task assembles a 2-byte packet and sends it wirelessly to the Raspberry Pi. Each of these tasks takes their semaphore to indicate that they have handled the current door-entrance attempt. Once all tasks have completed, the RTOS returns to execution of the lowest priority task, motion detection. From here, the system is ready for a new temperature scan.

Another strategy employed in the design of the embedded system was to begin with “driver” code. Each subsystem would have a driver source file associated with it for use as a small test program for the subsystem’s functionality. For example, “TemperatureSensorDriver.c” was the driver for the temperature sensor subsystem and repeatedly took temperature measurements. This file could be compiled without the use of the RTOS, which was useful for isolated testing. The strategy was to design driver files and test each subsystem before bringing each of the driver functions together in the RTOS code. This also allowed isolated testing of the RTOS as well.

Server & Web Dashboard

The Raspberry Pi server had two primary functions: to receive temperature data from the microcontroller and then to evidence that data in a usable dashboard. Figure 10 summarizes the high-level architecture of the server. The Bluetooth monitoring process and the web server process were both implemented by the team whereas a database library (SQLite) was used as-is in the middle. This was an important design decision because it allowed the Bluetooth code and the web dashboard code to be developed independently and tested separately with the use of a dummy database file.

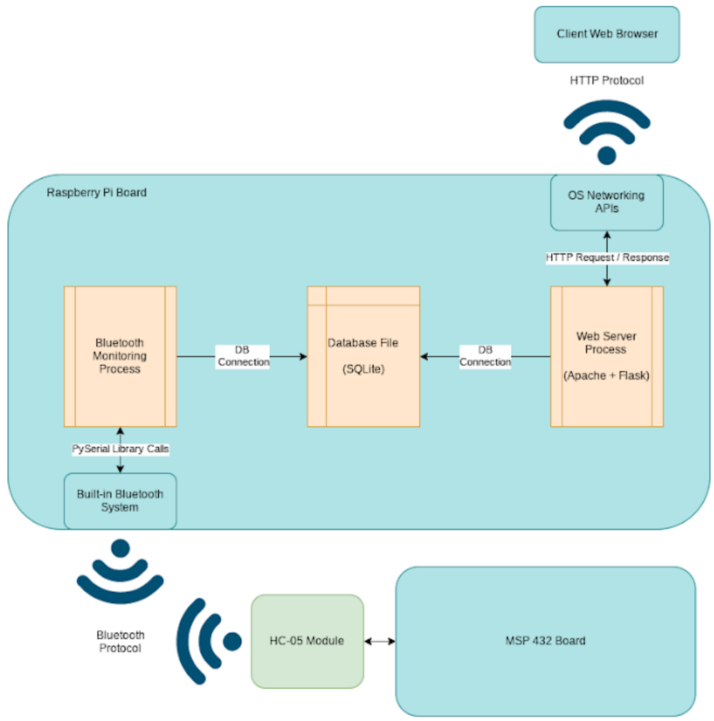

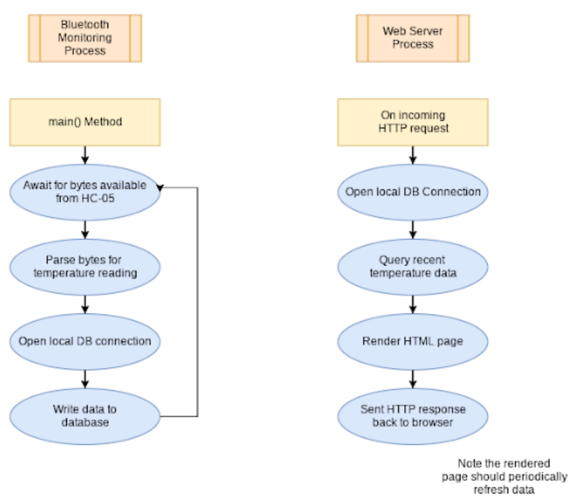

Database Schema and Example Data				
id	door_id	timestamp	temperature	decision
1	1	2020-09-01 09:21:12.00	98.4	ALLOW_THROUGH
2	1	2020-09-02 10:13:52.00	97.9	ALLOW_THROUGH
3	2	2020-09-11 16:00:00.00	100.2	LOCK_OUT

Figure 10. Raspberry Pi Architecture

To get the Bluetooth process working, it first had to be manually paired with the HC-05 on the microcontroller. Using the MAC address of the Bluetooth module, the Raspberry Pi paired and connected with the device. Then, once paired the device was available in the */dev/rfcomm0* file on the server's filesystem. Once paired and available, our code could automatically listen for packets and decode them using a finite state machine. When a complete transmission occurred as defined in our packet structure, the resulting information was written to the database.

The database had the table schema shown in Table 3. For each temperature reading received, a row was written to this database table with all fields populated.

Attribute	**Type**	**Description**
door_id	string	A human-readable indicator of which door the message came from. This is included for easy extension of the project to multiple doors reporting to the same server, but was only ever set to "Door 0" in our initial implementation.
fault	integer	A representation of the fault signal that the locking mechanism's driver chip reports. By looking at this integer, one can discern

		between potential errors like “no load connected” or “insufficient power”.
temperature	float	The actual temperature scan reported by the microcontroller. A fahrenheit measurement between 95 and 107.6.
decision	string	Restricted to the strings “ALLOW_THROUGH” and “LOCK_OUT”, an indication of what the microcontroller decided to do for this temperature scan. This is useful in case the temperature cutoff becomes configurable in the future.
timestamp	timestamp	The time that this scan occurred.

Table 3. Server database schema.

Separate from this data reception step was the web server and dashboard rendering process. A small website was written using the Flask framework to show the dashboard on incoming HTTP requests. Figure 11 shows a portion of the final dashboard configured to show the last 30 days worth of data. Here, the temperature scans are plotted for the user to see and then some sample statistics are calculated. Following these statistics is a table enumerating all of the scans (though the image is cut off to just show the first one). This dashboard refreshes its data automatically every 2 seconds to provide immediate feedback when new scans are made. The timespan to view data for is configurable on the dashboard, but restricted to the choices “Last Hour”, “Last Six Hours”, “Last Day”, “Last Week”, “Last two Weeks”, and “Last 30 Days” for security purposes.

This dashboard is available via the Raspberry Pi's IP address to all devices on the local area network. This was accomplished using the Apache2 web server library. Apache2 was configured to point all incoming network traffic on port 80 (the HTTP port) to forward to the Flask application. Using Apache2 also provided benefits like reliability and constant availability whenever the Pi was powered on.

Show data for: Last 30 Days

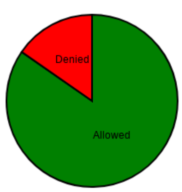

Statistic	Allowed	Denied	All
Count	11	2	13
Average	98.63	101.80	99.12
Max	99.6	102.3	102.3
Stdev	0.69	0.71	1.36

Decision	Fault	Temperature	Timestamp
ALLOW_THROUGH	0	99.6	2020-11-05 14:34:52

Figure 11. Final Web Dashboard.

System Assembly

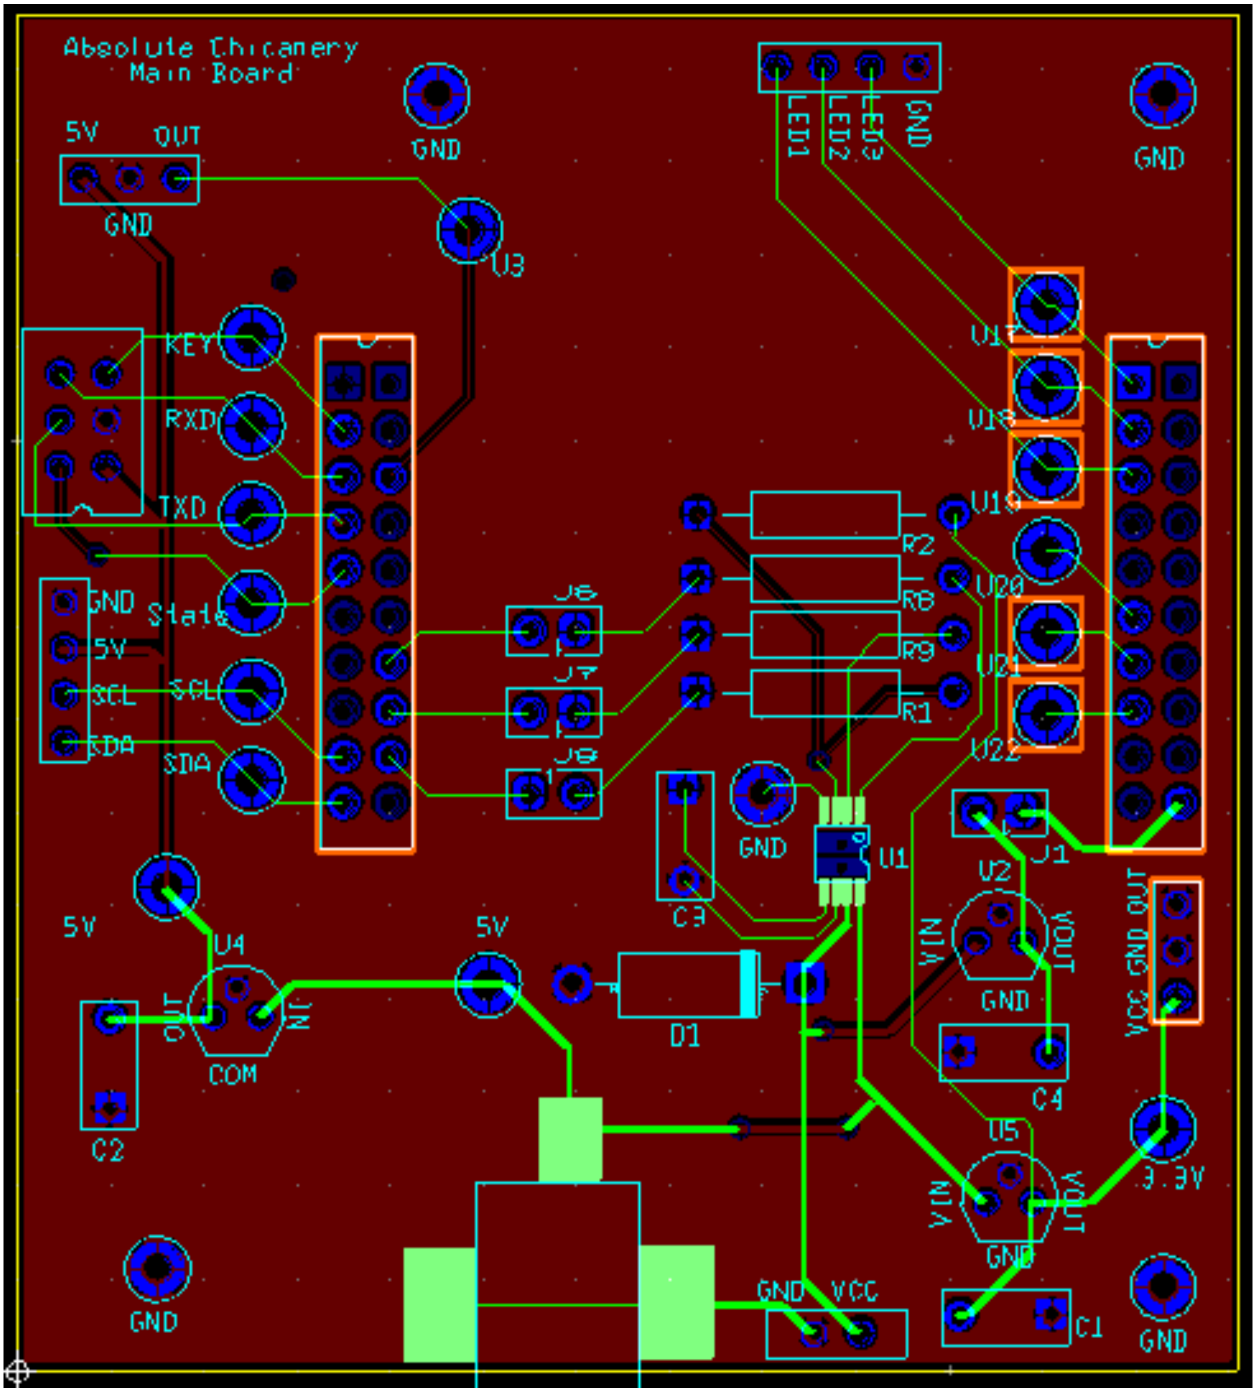

Figure 12. Final PCB layout for MSP432 header board.

Pictured in Figure 12 is the final main PCB revision used for the project. The board went through 3 major revisions, two of which were fabricated and tested. Each component on the board, barring standard components, was created in Ultiboard using the footprint creation wizard. The board is organized by power supply requirements. The left side of the board is for components that require the 5 V power supply. The components consist of the MOLEX connections to the temperature sensor, bluetooth module, and motion detector. The right hand side of the board is for the components that require 3.3 V and 12 V. The components that requires 3.3 V is the microcontroller, while the components that require 1 2V are the driver chip high side switch, MOLEX electric strike connection, and the two 3.3 V regulators used for the microcontroller and to step down the VOUT of the switch to implement fault logic.

The top of the board has a MOLEX connector that leads to the LED sideboard, shown in Figure 13. This board is separate from the main board in order to allow for flexibility in the location of the LEDs when placed in the housing. These LEDs convey to the user if they should wait, enter, or were denied. It consists of low power LEDs and drain resistors to set the current flowing through them.

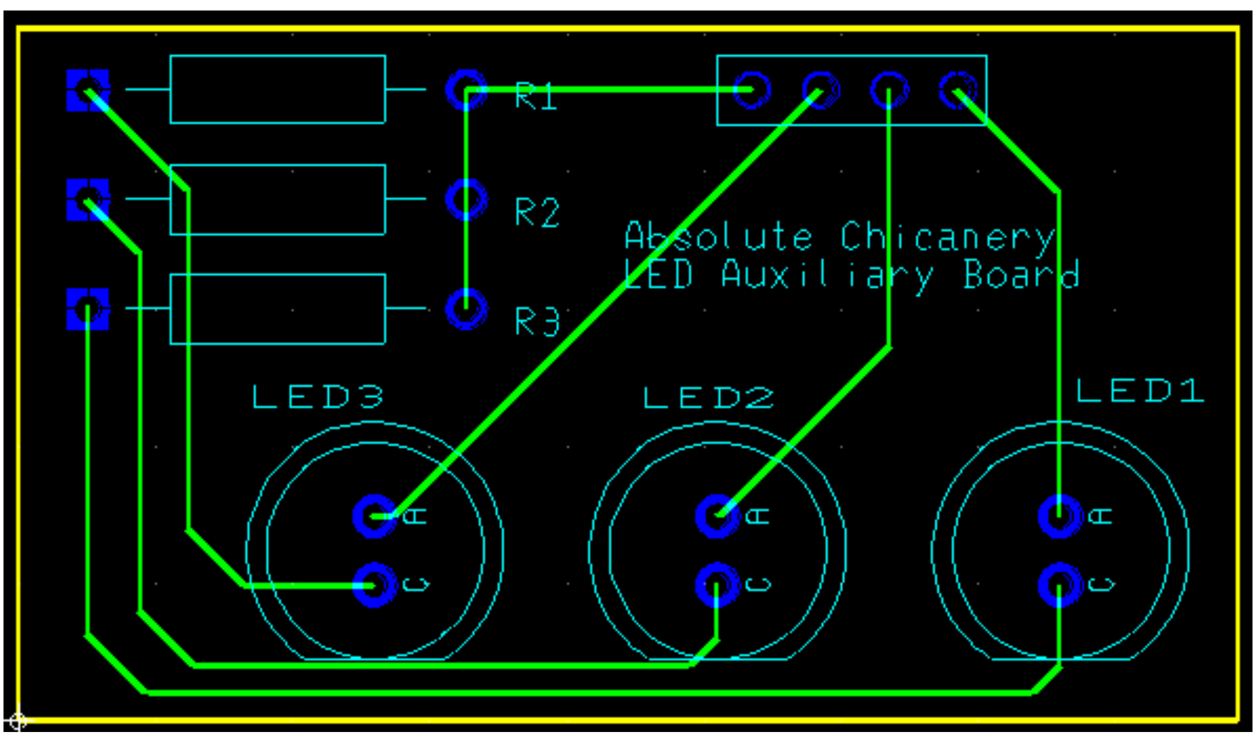

Figure 13. PCB layout for LED sideboard.

The design for the housing unit is shown from an aerial view in Figure 14. From left to right, the design features 3 holes for the power supply, temperature sensor, and motion detector. The LED panel faces outward through three holes cut post fabrication on the right hand panel. The unit is covered with a piece of clear acrylic to aid in debugging the system after it has been assembled. The microcontroller is mounted on the back panel of the housing and the Bluetooth module is to be attached on either side of the vertical piece of the system. We chose an "L" shaped design to clarify wrist placement and keep ambient temperature constant across the temperature sensor by blocking sunlight. This housing is intended to mount on a door frame at slightly above waist level, though this was never fully achieved during testing due to concerns about needing to further test the device outside of the lab. The fully assembled housing used during testing is shown in Figure 15.

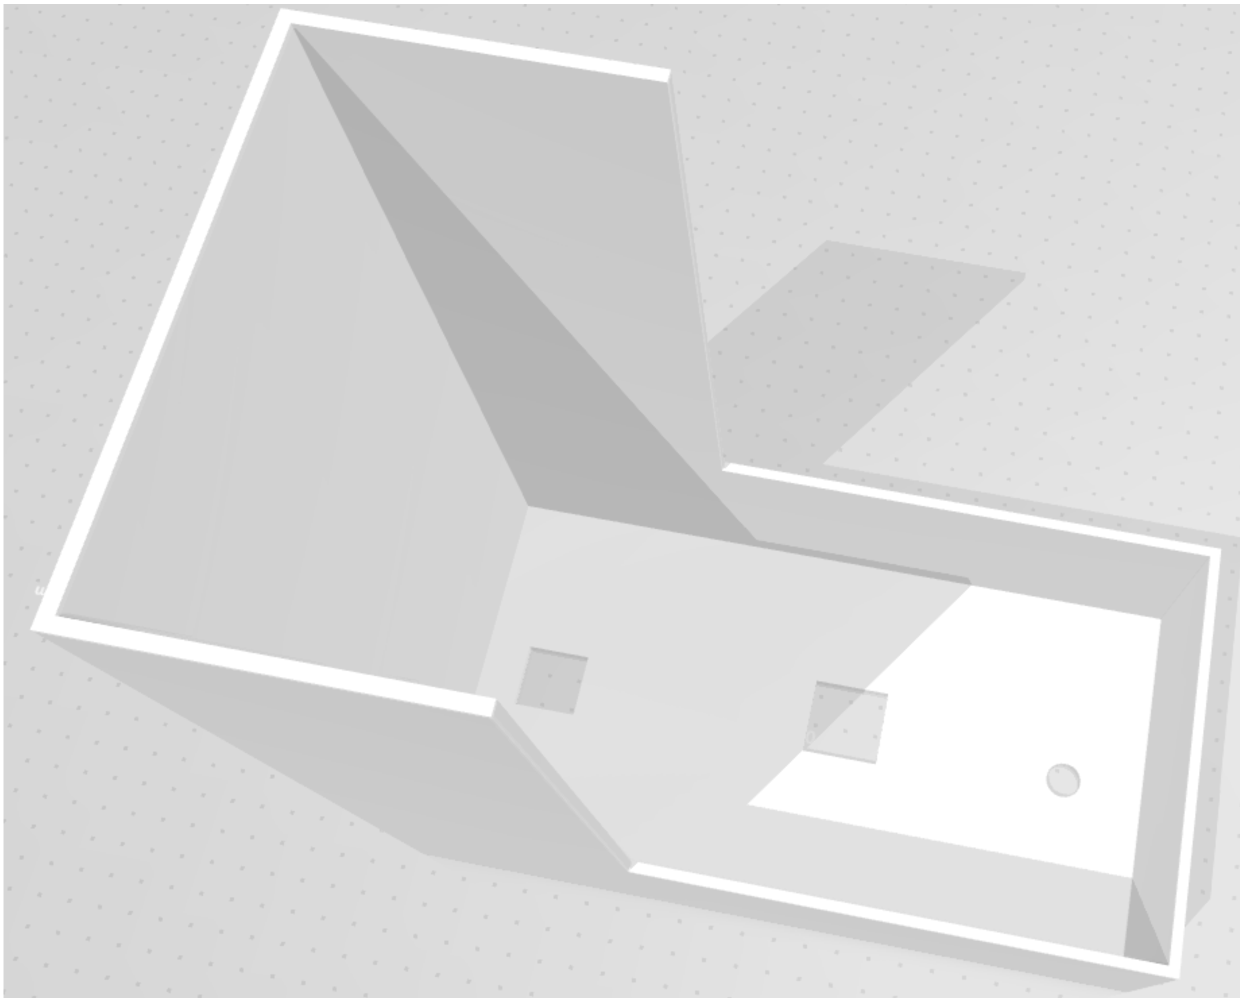

Figure 14. CAD model of mechanical housing.


MSP432
Indicator LEDs

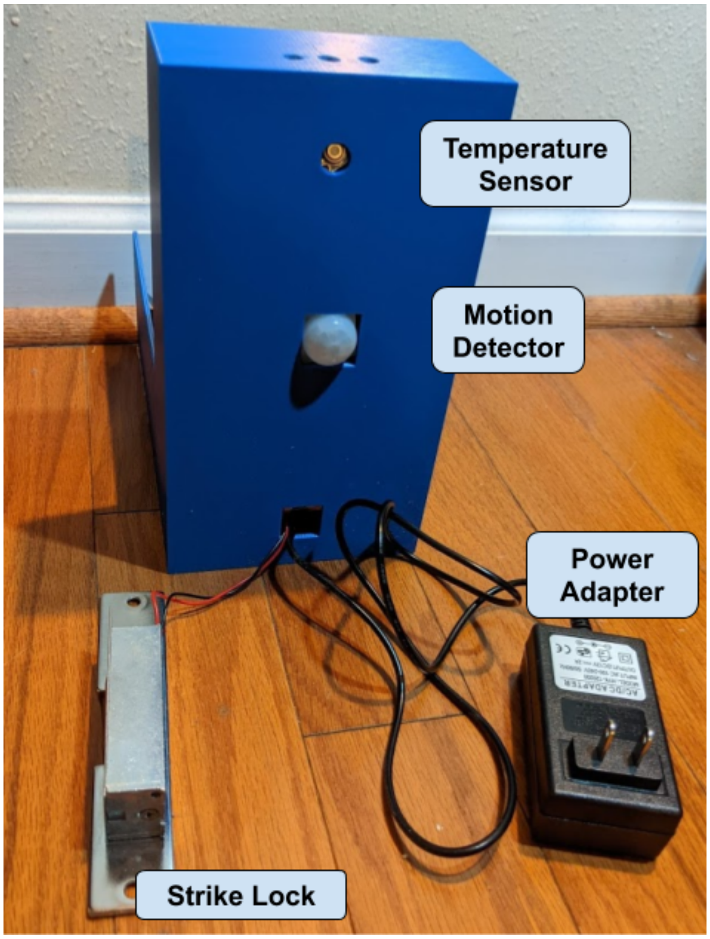

Figure 15. Assembled system prototype, from above (left) and below (right).

Project Time Line

Figure 16 and Figure 17 show the team's proposed timeline for execution and what ended up actually happening, respectively. At the beginning of the semester, greater parallelism was achieved as part determination, schematic and footprint design, and embedded coding could all happen independently of one another for each subsystem. As the semester went on, the more complex components like the temperature sensor and Bluetooth module took more time to get working independently when compared to the motion detector and lock. While we initially aimed to have the entire system prototyped on the breadboard by October 15, this did not end up happening until the first week of November. Final PCB design and mechanical assembly thus were pushed back to the last week before Thanksgiving as we had to wait for some last-minute parts to arrive. Thankfully, due to thorough prototyping on the breadboard the system worked as anticipated when fully installed with little need for adjustment.

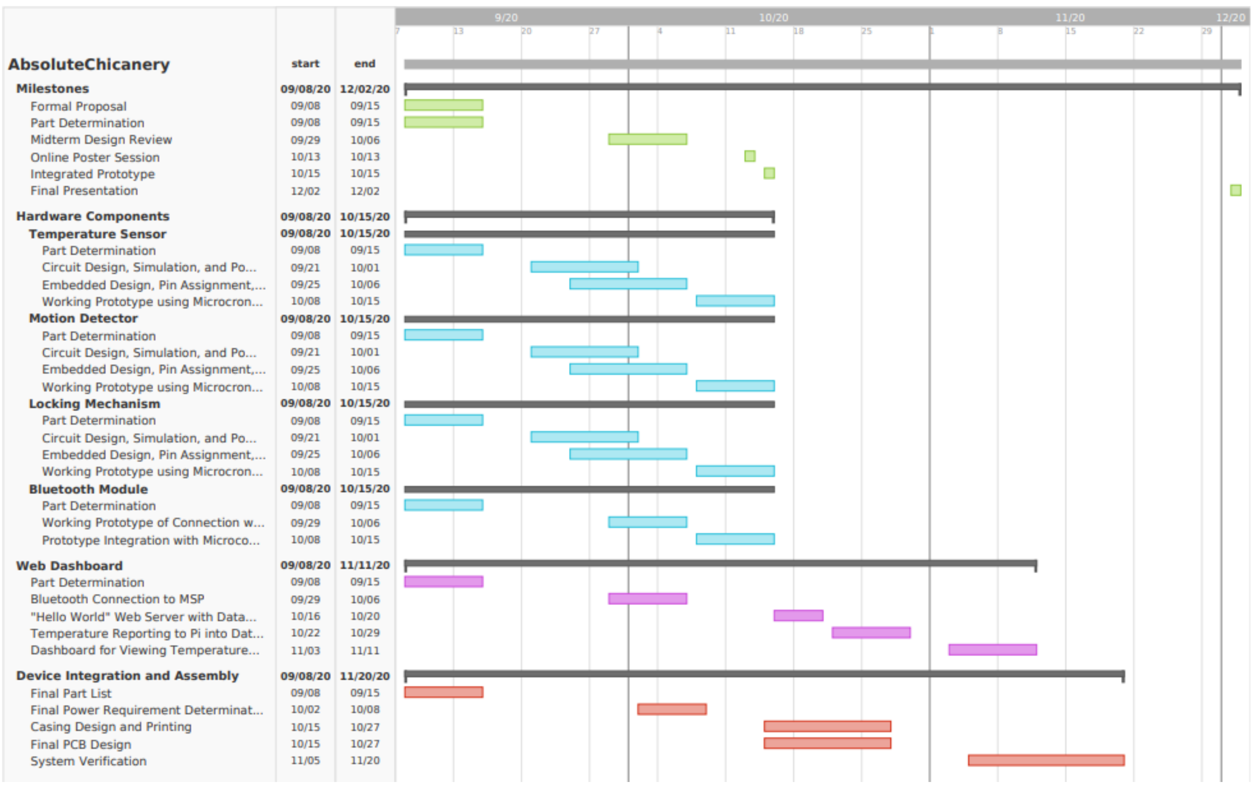

Figure 16. Proposed Project GANTT Chart

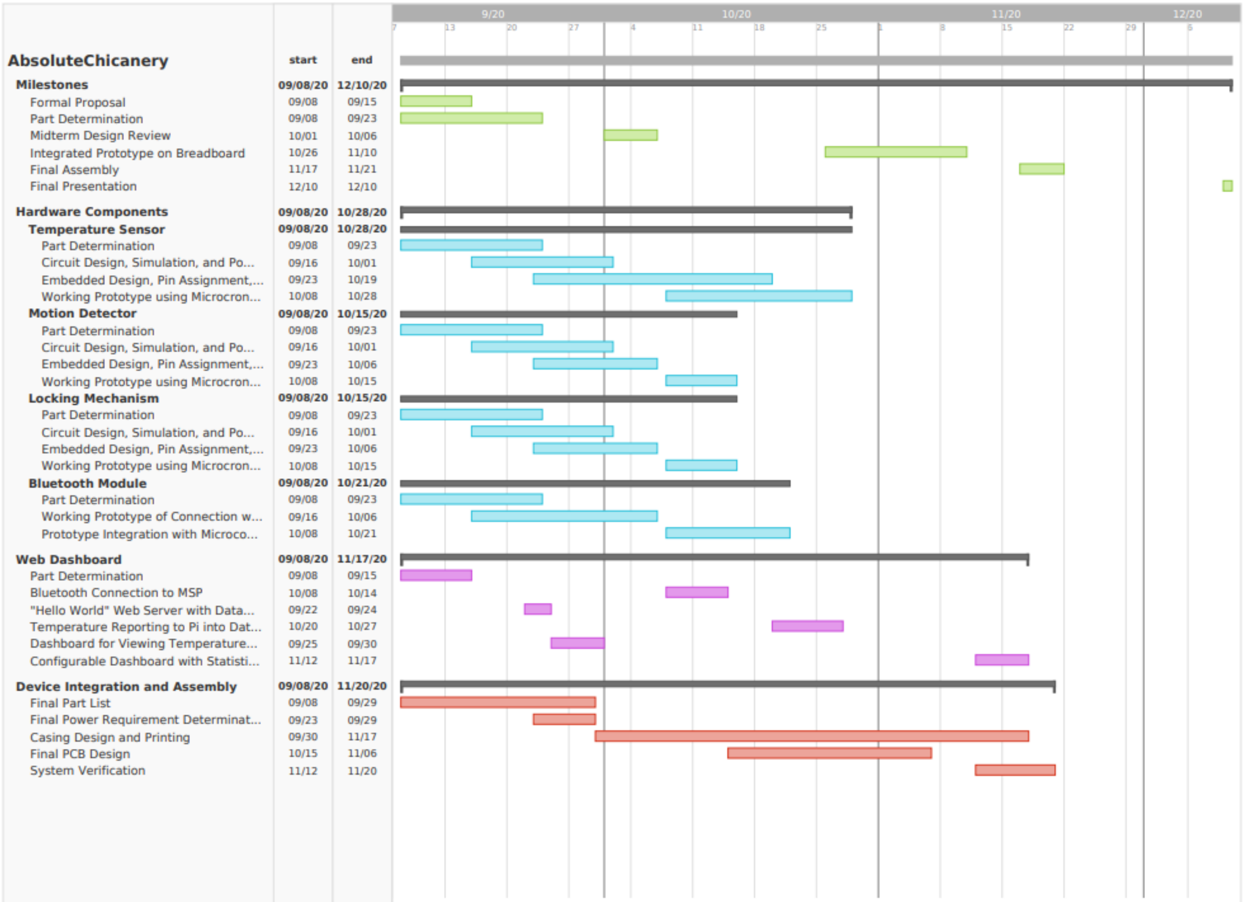

Figure 17. Final Project GANTT Chart

Test Plan

To minimize the need for adjustment when integrating the components together, a separate test plan was created for each subsystem in the project. This way, subsystems could be verified independently and any issues occurring after assembly must be due to the connections and interactions with one another.

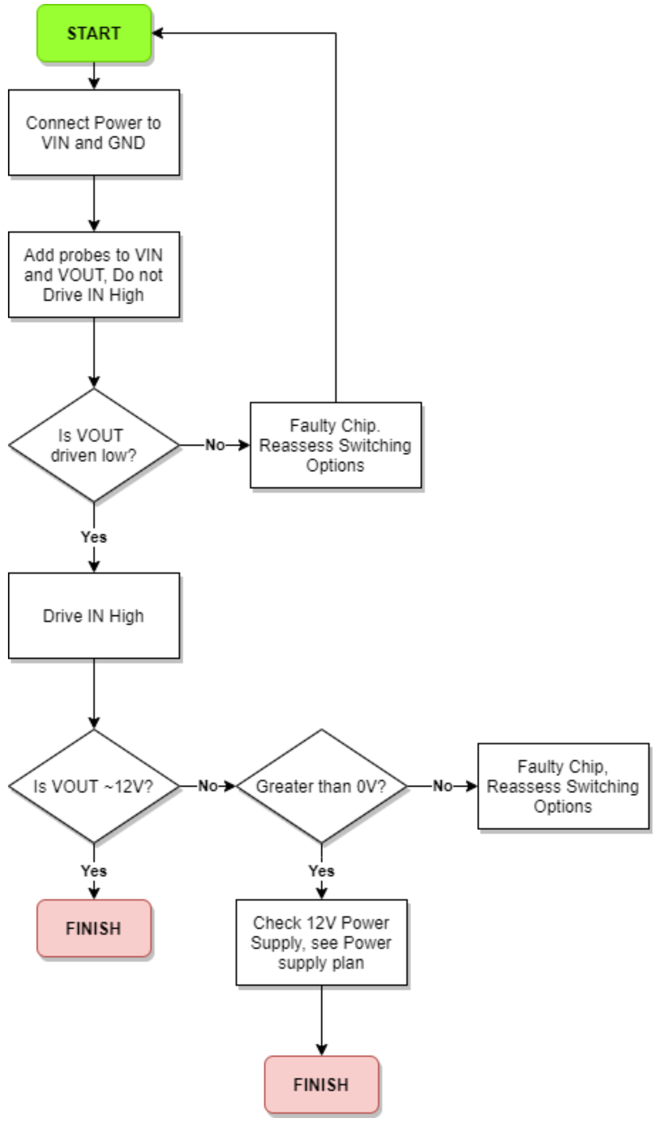

Figure 18. Driver chip test plan.

The test plan for the driver chip is shown in Figure 18. The chip has 3 inputs: the VCC to be switched on/off, the in control, and a diagnostic enable that allows fault protection to be enabled and diagnosed. The outputs to check functionality are OUT and FAULT. First, the chip is wired with 12 V in VIN and all other pins driven low or grounded. VOUT should be low at this stage of testing. If VOUT is low, then IN should be driven high. When IN is driven high, VOUT should switch to 12 V. If the above conditions are met, then the chip is operating adequately, and testing can continue.

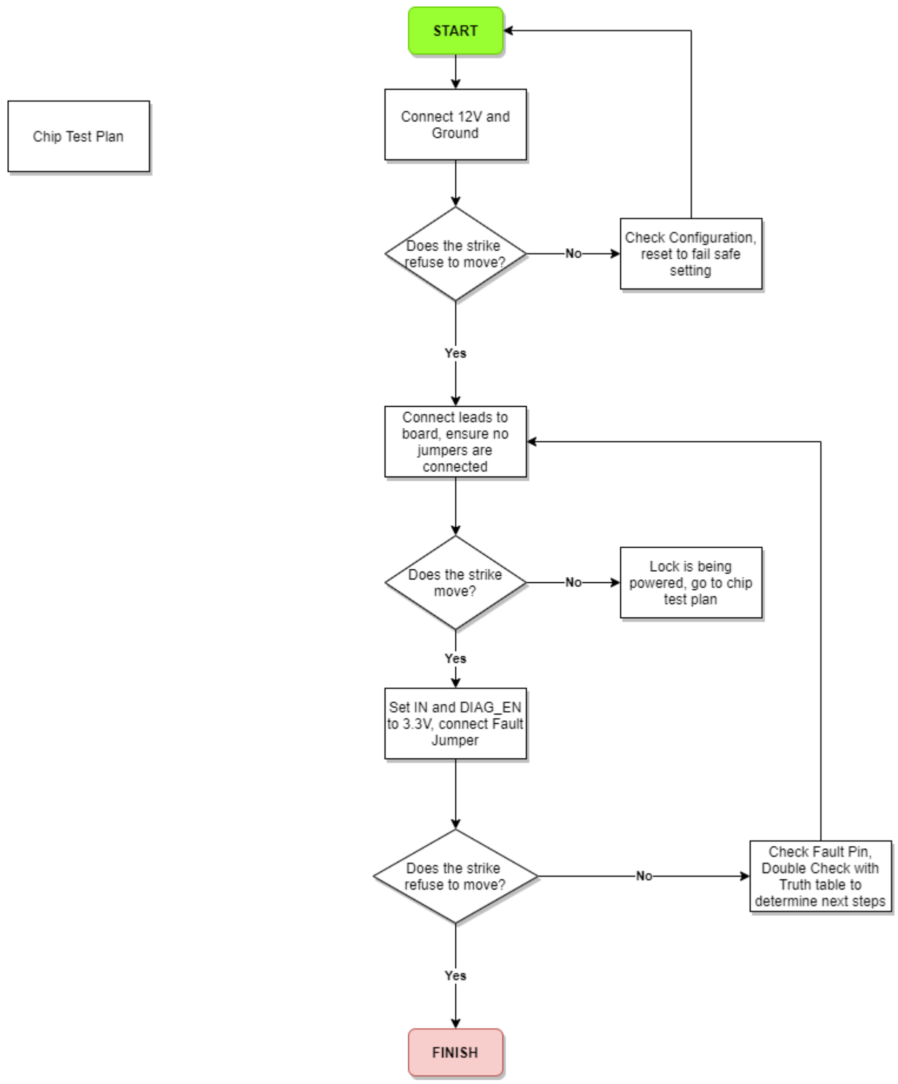

Figure 19. Locking mechanism test plan

The electronic strike only has two inputs, VIN and GND. The latch on the strike can be considered an output. Figure 19 begins by connecting VIN to a power supply and GND to the

GND of the power supply. Since the strike is classified as fail safe, it should be locked when voltage is applied. Therefore, the lock should not move when 12 V is applied to VIN. Next, power should be removed from VIN and the strike should be wired to the board to ensure the traces from the driver chip are behaving properly. No power should be applied, and the strike should move. Lastly, the driver chip should be triggered to activate the strike, and it should be locked.

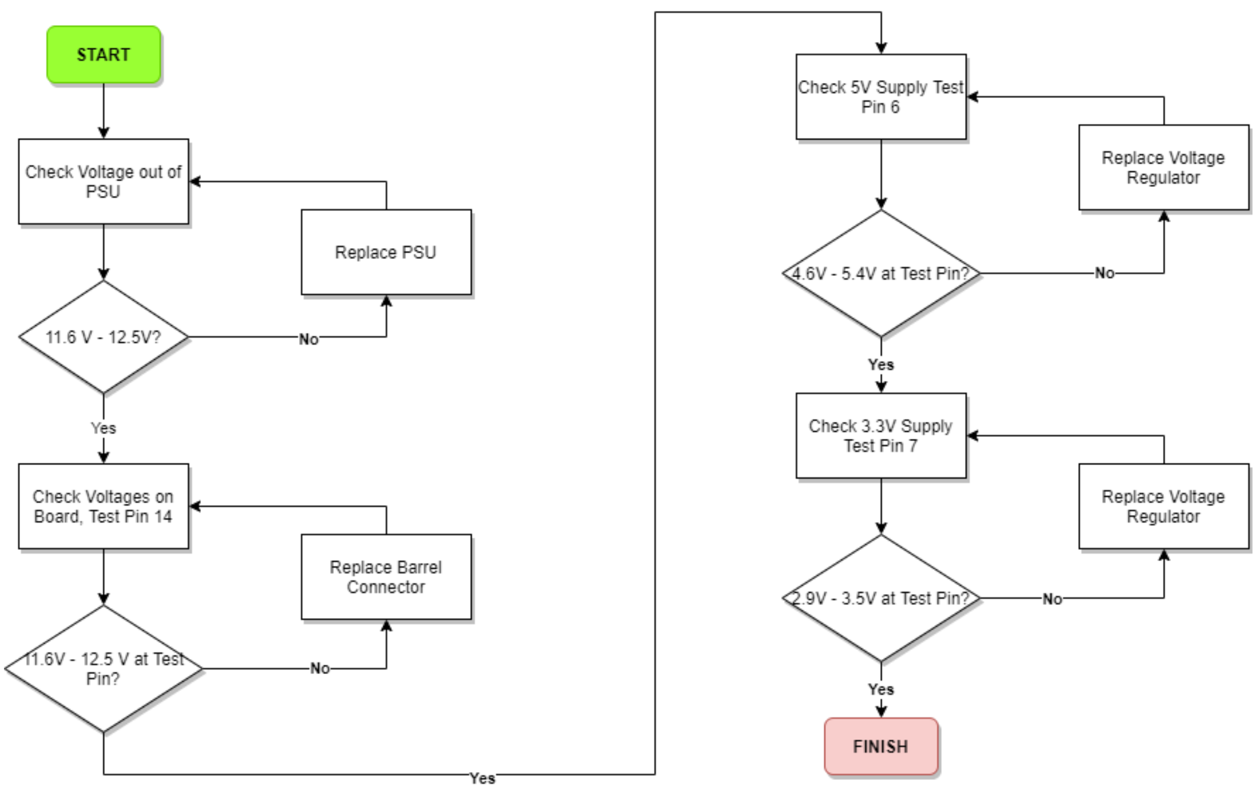

Figure 20. Power test plan

The power supply for the entire system starts with a 12 V power supply unit (PSU) from a wall outlet. This voltage is used to control the door lock, and is stepped down with linear voltage regulators to 3.3 V and 5 V for the microcontroller and the peripherals, respectively. The

PSU voltage should first be confirmed to be within 0.5 V of 12 V. Then, the PSU should be inserted in the barrel connector on the board and the voltage verified at a test point. Then, the linear voltage regulators should be verified on a breadboard. Once the regulators are confirmed to be working, they should be added one at time to the board and verified on a test point.

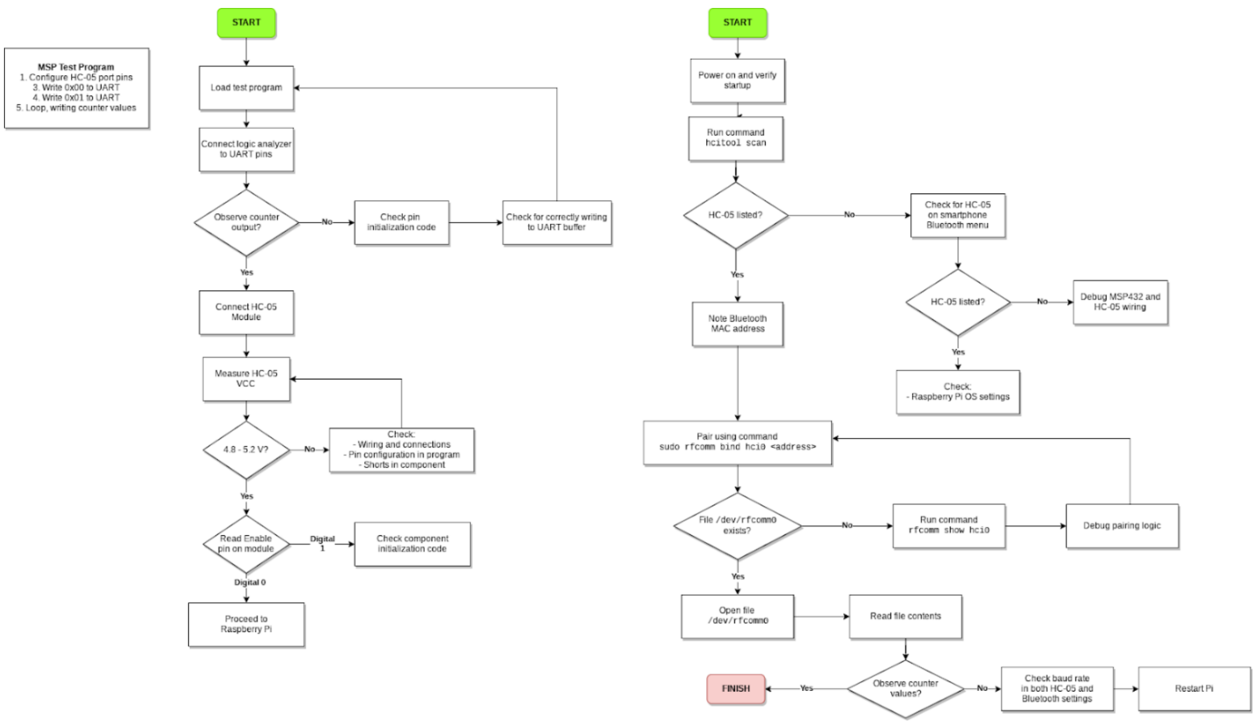

Figure 21. Bluetooth test plan, HC-05 side

Figure 21 and Figure 22, below, comprise the testing of the Bluetooth connection on both the microcontroller and the server, respectively. The HC-05 is initially loaded with a test program that repeatedly writes an incrementing counter to the UART interface. This can be read from the logic analyzer on the VirtualBench to make sure the embedded code uses the interface correctly. After plugging in the HC-05, we can make sure that VCC power is not distorted and that the code properly configures each pin on the module. On the server side, the tester must get a remote terminal on the Pi to run shell commands. First, the tester must scan for the HC-05 module, pair with it, and bind to it (the first two of these are only required on first setup). Once connected, the tester can examine the corresponding file on the device to see the Bluetooth counter values. The exact commands to run and files to examine are listed in the boxes of the test plan. Packet encoding and decoding were tested separately with unit tests in code, so they were not included in this test plan for the sake of simplicity.

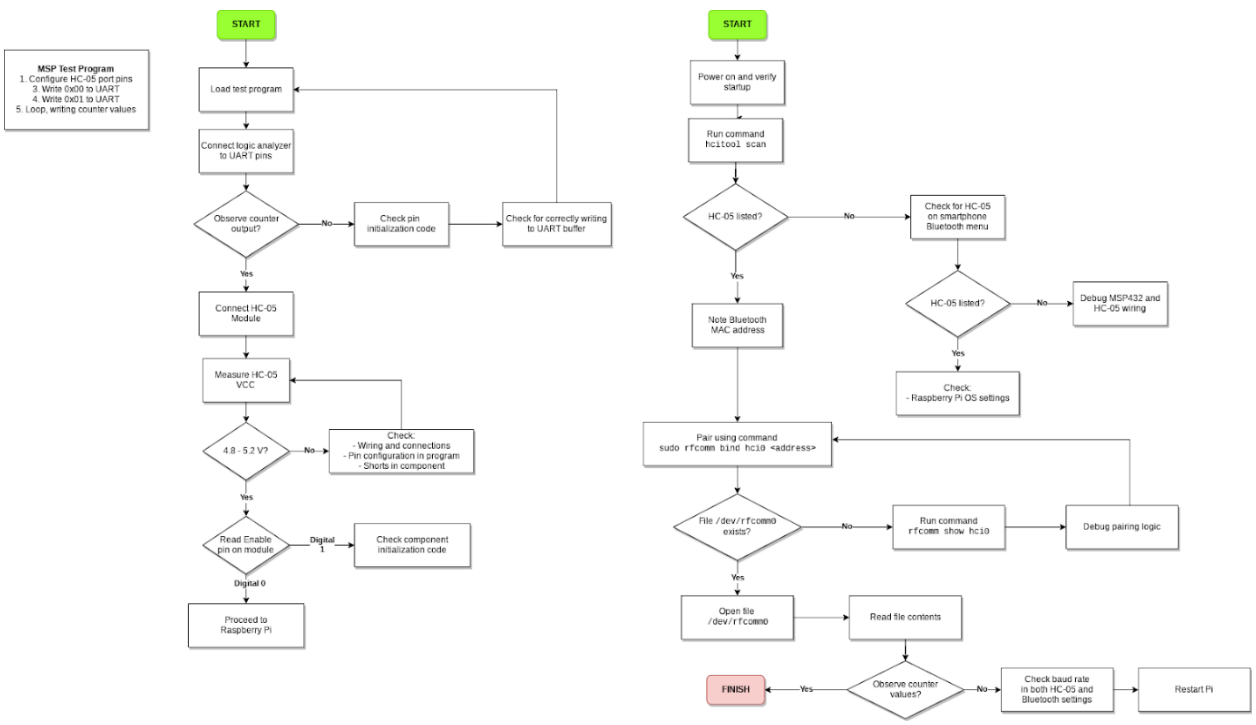

Figure 22. Bluetooth test plan, Raspberry Pi side

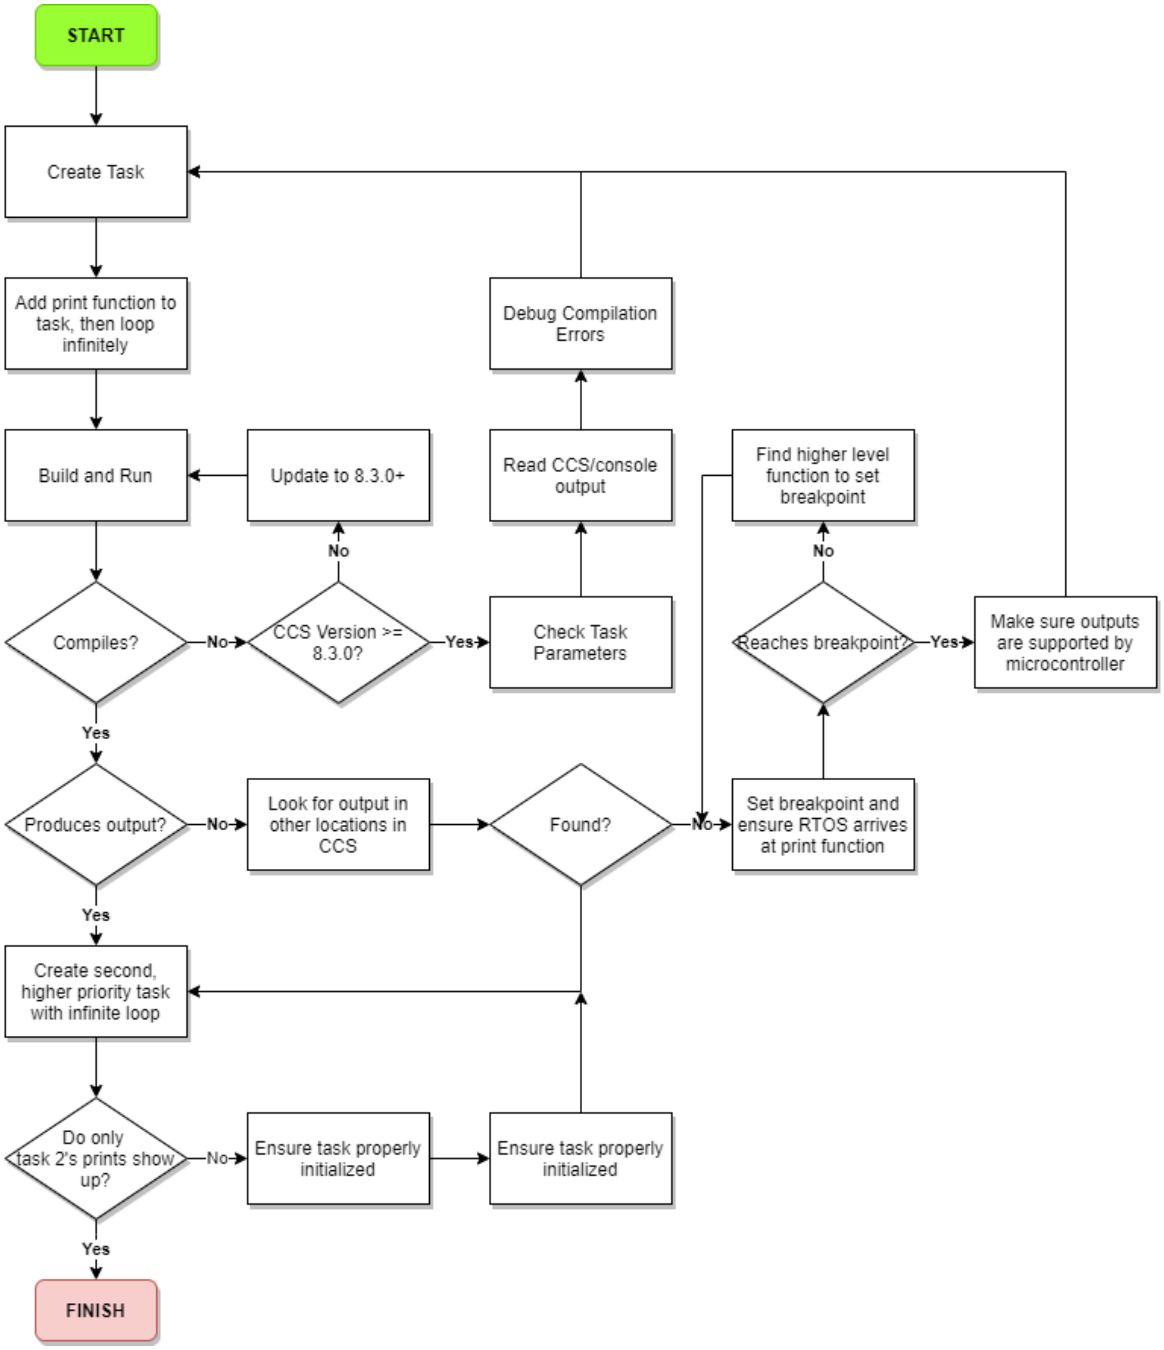

Figure 23. RTOS test plan

The RTOS testing was fairly seamless. The example projects provided by FreeRTOS greatly reduced the amount of debugging involved with compiling and designing for the

MSP432 system. The largest struggle with the RTOS was managing different clock speeds for each subsystem. For example, the UART communication for the Bluetooth module operated at a different clock speed than the I2C bus.

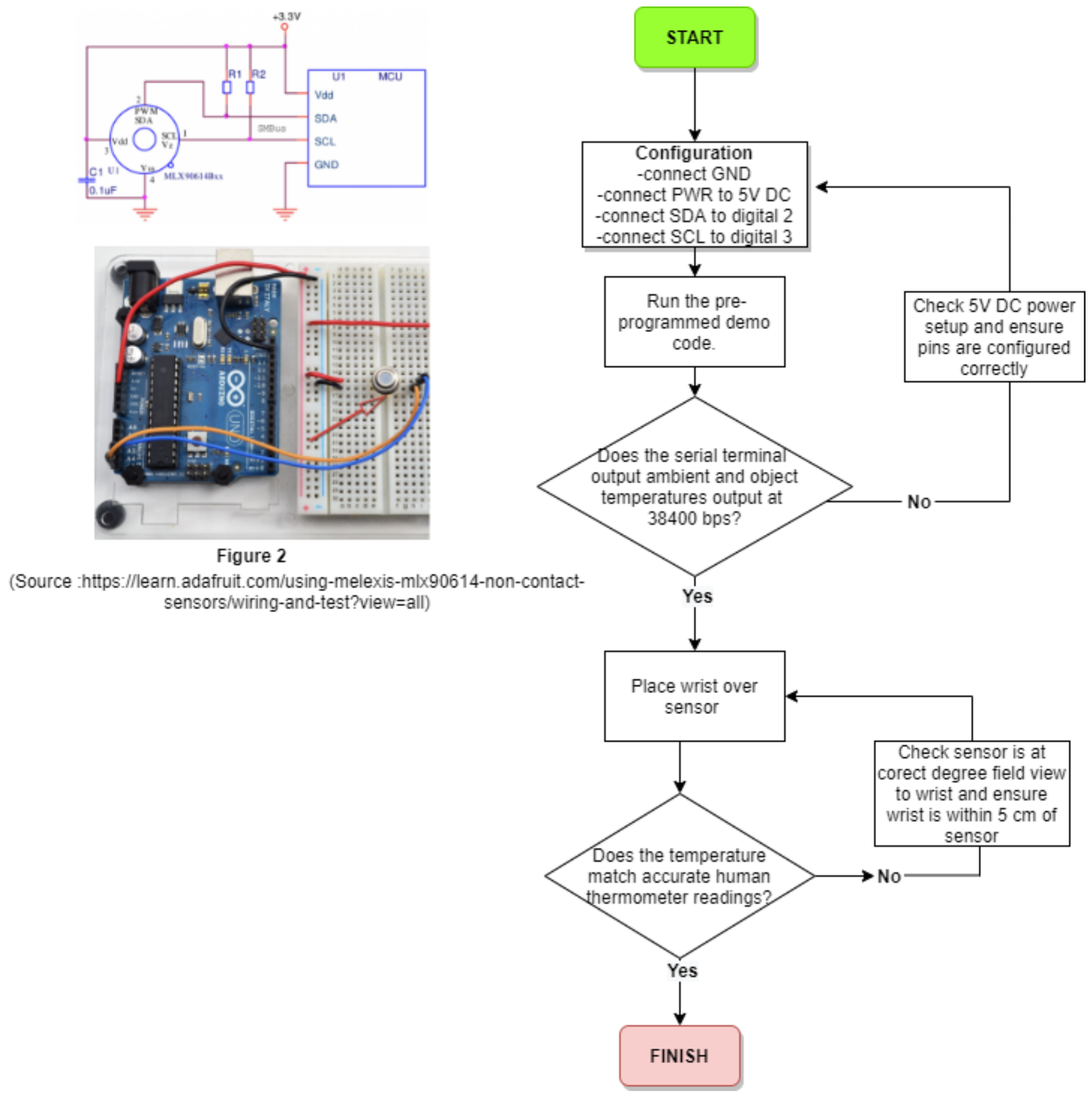

Figure 24. Temperature sensor test plan

Testing the temperature sensor module chip went as planned. Despite using the MSP432 microcontroller for the overall system, testing for this part was done using an Arduino microcontroller allowing our team to easily view serial terminal results with test code provided from Melexis Technologies. After wiring the temperature sensor to the Arduino and running the test code, we were able to get ambient and object temperature readings from the sensor shown on the serial terminal every few seconds, as expected, ensuring functionality.

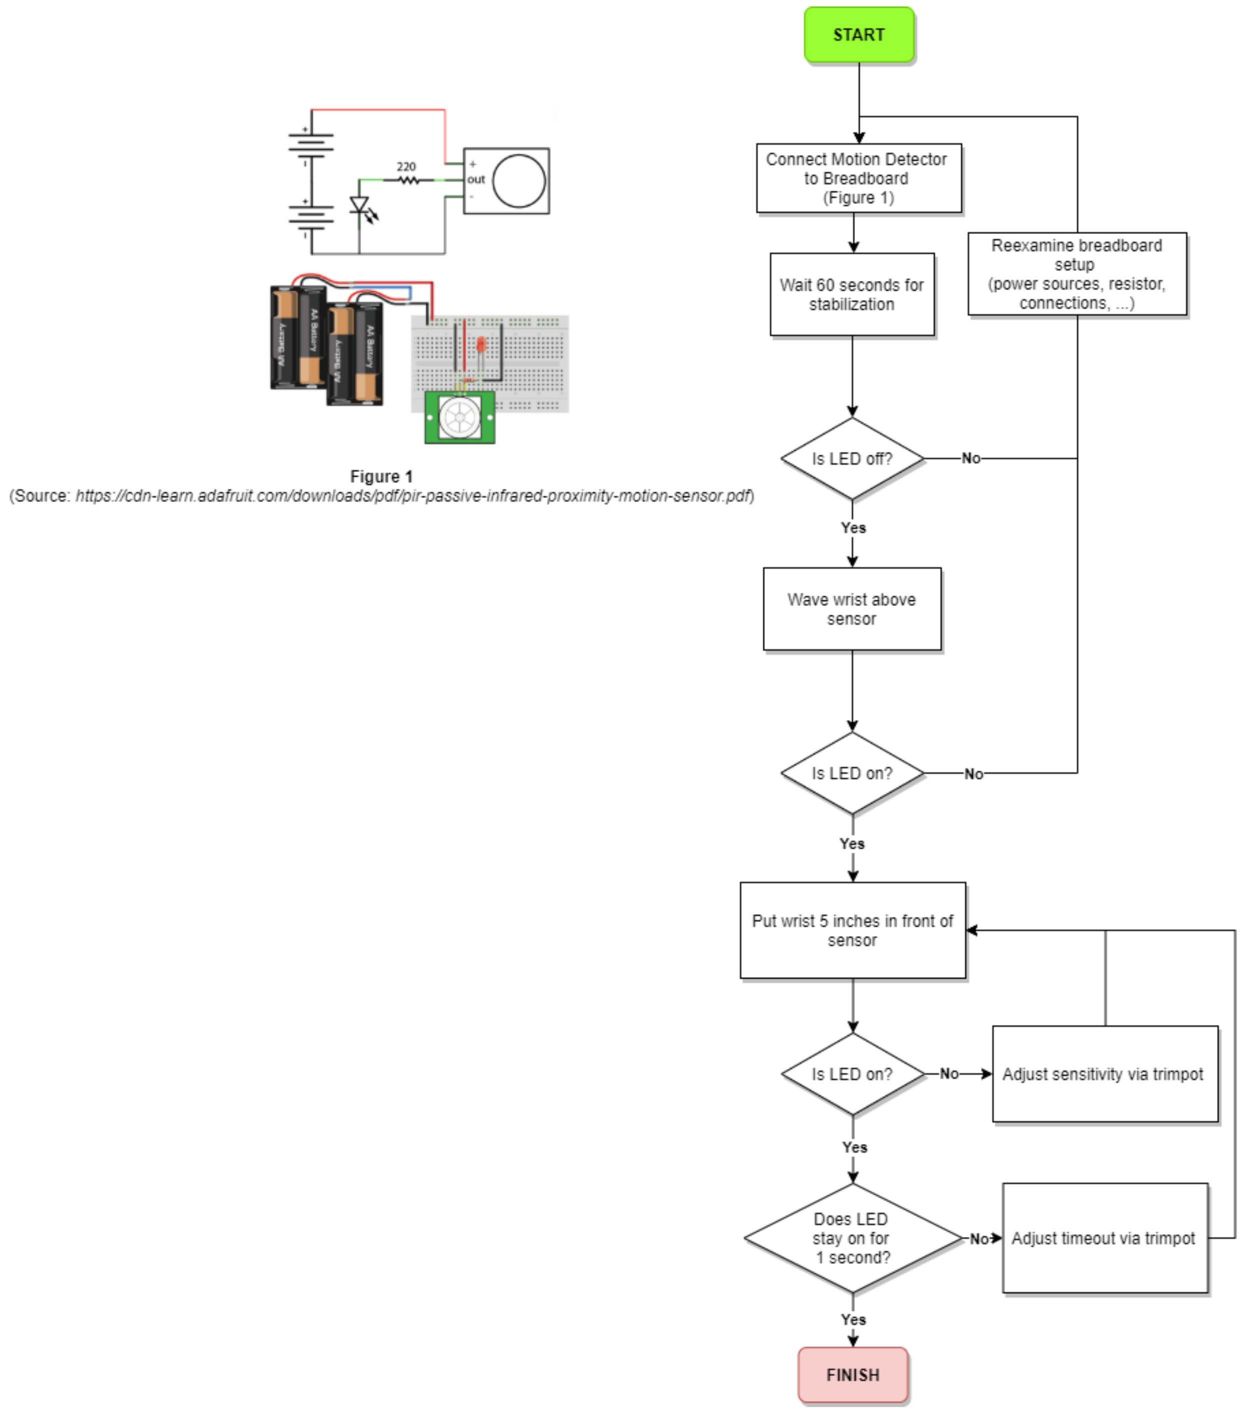

Figure 25. Motion sensor test plan

The test plan above highlights how the motion detector was tested. The test circuit shown in Figure 1 of the test plan itself comes from the datasheet, as well as the flow of steps taken to ensure that the motion detector is working properly. Since the manufacturers strongly recommend this version of a test plan, it was determined to be a trusted and comprehensive way to test this component [43].

Final Results

The final prototype of our device successfully takes non-contact temperature measurements, conditionally locks the door based on these measurements, and sends the data over Bluetooth to the Raspberry Pi to be shown on a web page. Table 4 represents the requirements for success outlined in the initial project proposal, though some features have been added and some changed slightly over the course of the semester.

For the locking mechanism section, we altered expected behavior so that the lock defaults to being locked until a valid temperature is scanned rather than defaulting to unlocked as specified in the table. This was a conscious change based on further thoughts on user interaction and fire safety. When a user scans their temperature, the lock reacts quickly, imperceptible to our group members during final testing. In the case of non-fever temperature the door unlocks, and after a reasonable amount of time afterwards (around 10-20 seconds) it re-locks. If power is lost, the door unlocks but when powered, the default state is locked. This represents full completion of the proposed functionality.

In the temperature sensor category, we have fulfilled almost all of the requirements for full credit except for that of displaying temperatures to users. Our temperature scans are non-contact using infrared, and are taken immediately once motion is detected, and are then sent to the server over Bluetooth. When motion is detected, 128 temperature scans are rapidly taken and averaged. Instead of displaying the raw temperature, door users are only shown an LED indicating the lock decision. Temperature scans can be taken back-to-back in a reasonable amount of time (the same 10-20 seconds), and one of the LED indicators is used to show when the system is ready for a new scan. As discussed with advisors and application engineers from the manufacturer throughout the semester, our temperature accuracy during testing has been noticeably erratic. We believe this is mostly due to the specific sensor we were able to purchase, since we had limited options available as a result of the pandemic. Even after averaging 128 samples, temperatures are sometimes erratic and heavily dependent on the wrist's distance to the sensor. In summary, this sensor would be unsuitable for use if going to market but we don't believe it should affect our outcome in this category. In the end, we achieved all requirements for full points except that we use LED indicators rather than displaying the raw temperature on an LCD screen.

For the dashboard, we believe we achieved full functionality. When accessed by any web browser on the local network, the page polls for data every two seconds to provide the experience of immediate updates. Temperatures over a user-configurable timespan are plotted using color to differentiate fevers from non-fevers. Statistics are calculated and shown to users about these temperatures as well as a log of all data. If the server goes down for any reason, no data would be lost since it is persisted on the server's filesystem, though no new data could be added and the website would not be available during downtime.

Points	Locking Mechanism	Temperature Sensor	Dashboard
2	• Locks immediately upon receiving temperature reading above the limit • Lock remains in locked position until receiving temperature in accepted range or within reasonable amount of time • Lock default (even with no power) is unlocked for safety	• Temperature readings are done via infrared (non-contact) • Temperatures taken as soon as new person steps up to device • Temperature readings are triggered by motion detector for energy-saving purposes • Temperatures readings are sent immediately to the server • Temperature readings are also displayed for users to see • Temperature readings can be taken back-to-	• Updates automatically as soon as new temperatures are taken • User-friendly UI including helpful graphs showing trends over time • Available as a website (not device-dependent) • Not fully dependent on server being connected to power + will not lose data if server loses power

		back in reasonable amount of time	
1	• Locks a reasonable time after the person steps up to the door, the majority of the time • Lock default is still unlocked	• Temperature readings are not displayed to user specifically but displayed via red/green LED (red = temperature outside of normal range, green = temperature OK) • Temperature readings triggered by motion most times, but may take several seconds to recalibrate + redetect for new readings	• Only available on 1 computer • Updates less frequently (i.e. once per day) with new readings • Contains 1 graph of temperatures taken over time
0	• Lock does not trigger even with temperature	• Temperature readings not triggered with any motion	• Not created

	reading outside of the normal range	• Temperatures taken are very inaccurate	

Table 4. Initially proposed project requirements and points.

While the presented project fulfills the requirements initially outlined, it is not ready for market adoption. As mentioned previously, the temperature sensor is not accurate enough for a commercial product. Another shortcoming is that the team eventually wanted to integrate multiple doors to the system. That is, multiple locks and multiple MSPs reporting data to the same server. While portions of the system were designed with this in mind (most of the server code is equipped to handle it), it quickly fell out of scope with time and budget constraints imposed this semester. We also fell short in our mechanical assembly and installation. Due to the team’s limited experience in this area and time constraints in the last week of lab access, the final assembly was not as sturdy as initially expected. To fulfill manufacturing standards and generally keep the system functioning reliably, this would need to be significantly improved if commercialized.

Costs

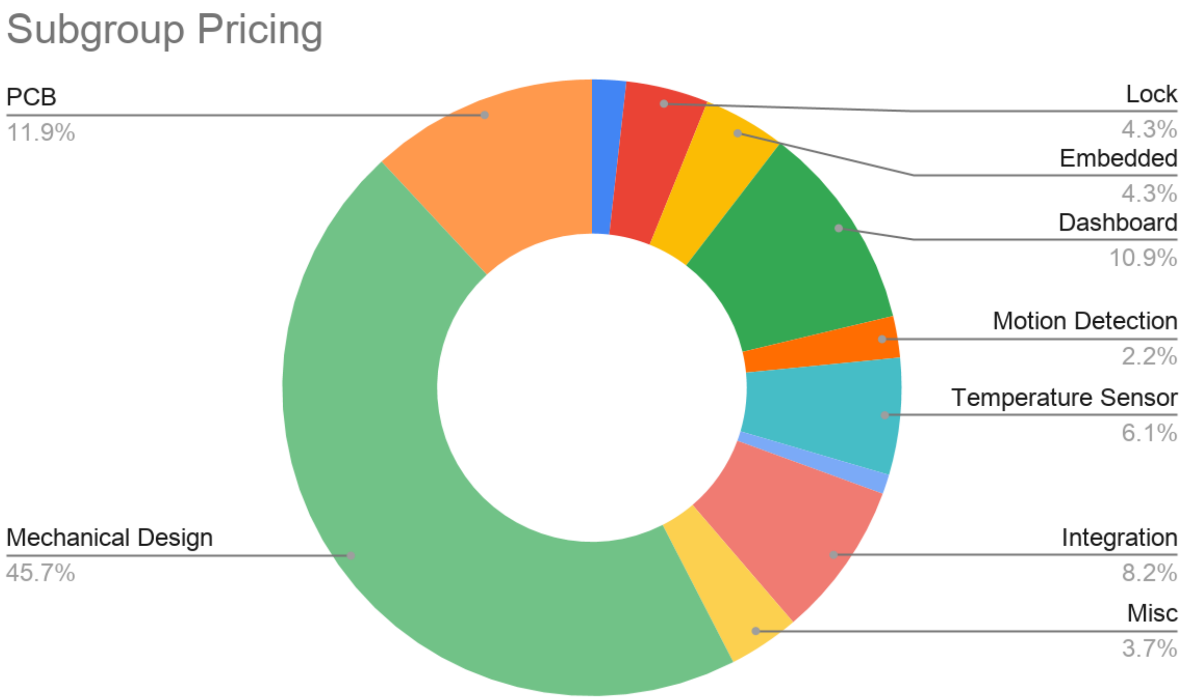

Figure 26. Budget subgroup breakdown

Each budget item was broken down into a subsystem category. The mechanical design was the most expensive, which included costs associated with 3D printing, acrylic parts, and other fasteners. The PCB was the next most expensive subgroup. This included the cost of our PCB as well as a revised version. The dashboard included the Raspberry Pi 3B+ as well as SD cards and other necessary peripherals. Most of the expense from the Embedded system came from the MSP432 LaunchPad. The integration subgroup included parts that were necessary to physically connect different subsystems, such as connectors, wires, and crimps. The overall cost of the project was $556.20.

If this project were taken into high volume manufacturing with 10,000 units, we would expect part prices to change from approximately $316 to $189. The 3D printing, which costed around $240 could be replaced with injection molding or some cheaper technique, which would be in the range of $15-20 for the price of plastic. Therefore, $200 would be a reasonable device manufacturing cost to aim for in high volume production.

Future Work

There are several major areas in which this system could be improved given more financial resources and time. The temperature sensor that has been utilized in this system should be swapped out with a far more accurate sensor, perhaps with a lens that increases precision as well.

Future versions of this system should also enhance user experience. This could include replacing the three LED indicators with an LCD screen that instructed a first-time user on how to interact with the system, provided the user with their own temperature reading rather than a binary good-or-bad, and gave the owner of the system an idea as to the system's health.

A major concern with this system is how easily it may be rendered useless by user misuse, whether a user props the door open or just allows others to enter under his/her temperature scan. This was not a problem that the team could tackle in such a short timeframe, but it may be possible to incorporate a door alarm for the door being open for too long (to discourage propping the door open), a weight sensor to discern how many people had crossed the threshold per temperature scan, or perhaps a switch from the non-contact wrist temperature scanner to an infrared camera mounted on the door frame that would scan users' foreheads. These ideas, especially the lattermost, would likely have some ethical implications that would need to be explored.

A final extension that could make the project more usable would be allowing multiple door-mounted locks to report data to the same server. This would allow building owners to monitor entrant temperatures for all access points on the same dashboard. Much of the server code was designed with this extension in mind, but it would have required multiple copies of the PCB and microcontroller which we couldn't justify from a budget perspective.

References

[1] *Temperature Kiosks: What Are They and How Much Do They Cost?,* 2020.

[2] *Temperature Screening Kiosk | Personnel Management.*

[3] *Grant of Equipment Authorization,* 2018.

[4] *Equipment Authorization of Unintentional Radiators,* 2017.

[5] I. B. C. Council, International Building Code, Country Club Hills, IL: International'Code Council, Inc, 2017.

[6] D. H. C. D. Division of Building and F. Regulation, 2015 Virginia Construction Code Part 1 of the Virginia Uniform Statewide Building Code, State Building Codes Office, Richmond, VA, 2018.

[7] *Stainless Steel Electric Power Door Lock Access Control Fail Safe NO Narrow-type 12V DC.*

[8] *TPS1H200A-Q1 40-V 200-mΩ Single-Channel Smart High-Side Switch,* Texas Instruments, 2019.

[9] *Guidance Regarding Methods for De-identification of Protected Health Information in Accordance with the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule.*

[10] *IPC Checklist for Producing Rigid Printed Board Assemblies,* Association Connecting Electronics Industry, 2019.

[11] *Bluetooth Core Specification,* Core Specification Working Group, 2019.

[12] *I2C-bus specification and user manual,* NXP, 2014.

[13] *IEEE Standard for Information technology—Telecommunications and information exchange between systems Local and metropolitan area networks—Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,* IEEE, 2016.

[14] *OpenSSH,* 2020.

[15] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach and T. Berners-Lee, *Hypertext Transfer Protocol – HTTP/1.1,* IETF, 1999.

[16] *NI Multisim,* National Instruments.

[17] *Ultiboard.*

[18] *AutoCAD Overview.*

[19] *Git.*

[20] *Code Composer Studio (CCS) Integrated Development Environment (IDE).*

[21] *Code Composer Studio (CCS),* Texas Instruments.

[22] S. Tatham, *PuTTY Use Manual.*

[23] *MSP430 Optimizing C/C++ Compiler v20.8.0.STS,* Texas Instruments, 2020.

[24] *FreeRTOS.*

[25] *Raspberry Pi OS,* Raspberry Pi Foundation.

[26] *PySerial.*

[27] G. Haring, *sqlite3.*

[28] *Flask.*

[29] *Apache HTTP Server,* Apache Software Foundation.

[30] *Recycling FAWs.*

[31] M. Griffin, *Is PLA Recyclable?,* 2020.

[32] *How Do You Recycle Acrylic Resin?,* 2012.

[33] "Health kiosk". 12 2018.

[34] "Medical kiosk and method of use". 1 2018.

[35] "Multi-biometric enrollment kiosk including biometric enrollment and verification, face recognition and fingerprint matching systems". 5 2012.

[36] *HC-05 Bluetooth to Serial Port Module,* ITead Studio, 2010.

[37] *Pyroelectric Passive Infrared Sensor,* Adafruit.

[38] *MSP432P401R, MSP432P401M SimpleLink™ Mixed-Signal Microcontrollers,* Texas Instruments, 2019.

[39] *TL750L Low-Dropout Voltage Regulators,* Texas Instruments, 2009.

[40] *L4931 Very low drop voltage regulators with inhibit,* ST Microelectornics, 2017.

[41] *WP7113LGD T-1 3/4 (5mm) Solid State Lamp,* Kingbright, 2019.

[42] *PJ-036AH-SMT-TR DC Power Jack,* CUI Devices, 2019.

[43] *PIR Motion Sensor.*

[44] T. Ylonen, *The Secure Shell (SSH) Protocol Architecture,* C. Lonvick, Ed., IETF, 2006.

[45] M. Valentine, D. C. J. Bihm, L. Wolf, H. E. Hoyme, P. A. May, D. Buckley, W. Kalberg and O. A. Abdul-Rahman, "Computer-Aided Recognition of Facial Attributes for Fetal Alcohol Spectrum Disorders," *Pediatrics,* vol. 140, p. e20162028, 12 2017.

[46] D. Stauss, M. Rogers and M. Herr, *U.S. Privacy Law Implications with the Use of No-Contact Temperature Taking Devices,* 2020.

[47] D. Siegmund, S. Dev, B. Fu, D. Scheller and A. Braun, "A Look at Feet: Recognizing Tailgating via Capacitive Sensing," in *Distributed, Ambient and Pervasive Interactions: Technologies and Contexts*, vol. 10922, N. Streitz and S. Konomi, Eds., Cham, Springer International Publishing, 2018, p. 139–151.

[48] A. Sharma, *Motion Detector Using MSP430 Launchpad and PIR Sensor,* 2019.

[49] J. Morrissey, "Fighting the Coronavirus With Innovative Tech," *The New York Times,* 6 2020.

[50] E. Maras, *Temperature check kiosks ready to tackle COVID-19,* 2020.

[51] M. Kreiger and J. M. Pearce, "Environmental Life Cycle Analysis of Distributed Three-Dimensional Printing and Conventional Manufacturing of Polymer Products," *ACS Sustainable Chemistry & Engineering,* vol. 1, p. 1511–1519, 12 2013.

[52] H.-Y. Kang and J. M. Schoenung, "Electronic waste recycling: A review of U.S. infrastructure and technology options," *Elsevier,* vol. 45, p. 368–400, 12 2005.

[53] L. Hodges, *Ultrasonic and Passive Infrared Sensor Integration for Dual Technology User Detection Sensors.*

[54] E. Grodzinksy and M. Sund-Levander, "Understanding Fever and Body Temperature," p. 23–25, 8 2019.

[55] L. Greene, "Decoded Fail Safe vs Fail Secure - When and Where," p. 3.

[56] S. Graham, *Samuel Pierpont Langley,* 2000.

[57] S. Gong, C. C. Loy and T. Xiang, "Security and Surveillance," p. 18.

[58] A. Elejalde-Ruiz, "If you get sick with COVID-19, is your employer liable? As businesses prepare to reopen, worker safety is a priority.," *Chicago Tribune,* 5 2020.

[59] CDC, *Coronavirus Disease 2019 (COVID-19),* 2020.

[60] *WP7113LYD T-1 3/4 (5mm) Solid State Lamp,* Kingbright, 2019.

[61] *WP7113LID T-1 3/4 (5mm) Solid State Lamp,* Kingbright, 2019.

[62] *What Is Multisim.*

[63] *Validity of Wrist and Forehead Temperature in Temperature Screening in the General Population During the Outbreak of 2019 Novel Coronavirus: a prospective real-world study.*

[64] *Universal Asynchronous Receiver/Transmitter (UART),* Texas Instruments, 2010.

[65] *Ultiboard,* National Instruments.

[66] "Symptoms of Coronavirus (COVID-19)," p. 1.

[67] *Product Compliance and Safety.*

[68] "Pawl & solenoid locking mechanism". Patent US5887467A, 3 1999.

[69] *MLX90614 Family Datasheet Single and Dual Zone Infra Red Thermometer in TO-39,* Melexis, 2019.

[70] *Low-noise and long-range PIR sensor conditioner circuit with MSP430TM smart analog combo,* 2019.

[71] *How to use a high current solenoid with Arduino,* learnelectronics.

[72] *GitKraken,* Git.

[73] *DriverLib,* Texas Instruments.

[74] *COVIDWISE,* 2020.

[75] CDC, *COVID-19 Cases, Deaths, and Trends in the US | CDC COVID Data Tracker,* 2020.

Appendix

Table 5. Full list of project expenses.

Part	Quantity	Individual Cost	Shipping Est.	Total	Budget Rem.
HC-05 (x2)	1	$9.88	$0.00	$9.88	$490.12
Adafruit 1512	0	$15.95	$10.00	$0.00	$490.12
MSP-EXP432P401R	1	$23.99		$23.99	$466.13
Raspberry Pi 3B+	1	$35.00	$10.00	$45.00	$421.13
32GB SD Card	1	$8.49	$6.76	$15.25	$405.88
PIR Sensor	1	$12.09		$12.09	$393.79
Temp Sensor	1	$28.99	$5.00	$33.99	$359.80

Door Lock	1	$19.85		$19.85	$339.95
Power Barrel	2	$0.96		$1.92	$338.03
5V regulator	2	$0.79		$1.58	$336.45
3.3V Regulator	2	$0.65		$1.30	$335.15
Driver Chip	1	$1.56		$1.56	$333.59
Wire Connectors				$0.00	$333.59
2 Male	2	$0.26		$0.52	$333.07
2 Female	2	$0.40		$0.80	$332.27
3 Male	2	$0.28		$0.56	$331.71
3 Female	2	$0.45		$0.90	$330.81
4 Male	2	$0.29		$0.58	$330.23
4 Female	2	$0.48		$0.96	$329.27
6 Male	2	$0.32		$0.64	$328.63
6 Female	2	$0.57		$1.14	$327.49
MSP Power Header	2	$0.37		$0.74	$326.75
Test Pins BLK	10	$0.40		$4.00	$322.75
Test Pins WHT	15	$0.40		$6.00	$316.75

Test Pins Yellow	6	$0.40		$2.40	$314.35
2.2uF Caps	5	$0.50		$2.50	$311.52
NEW 3.3V Regulator just in case	2	$0.44		$0.88	$310.64
Shunts	3	$0.13		$0.39	$310.25
Shipping + Tax	1	$10.96		$10.96	$299.29
LED RED DIFFUSED T-1 3/4 T/H	2	$0.33		$0.66	$298.63
LED GREEN DIFFUSED T-1 3/4 T/H	2	$0.33		$0.66	$297.97
LED YELLOW DIFFUSED T-1 3/4 T/H	2	$0.33		$0.66	$297.31
CONN HEADER VERT 2POS	2	$1.11		$2.22	$295.09
CONN HEADER VERT 3POS 2.54MM	2	$1.40		$0.88	$294.21
CONN HEADER VERT 4POS 2.54MM	4	$1.45		$5.80	$288.41
CONN HEADER VERT 6POS 2.54MM	2	$1.37		$2.74	$285.67
CONN HOUSING 2POS .100" CRIMP	2	$0.44		$0.88	$284.79
CONN HOUSING 3POS .100" SINGLE	2	$0.42		$0.84	$283.95
CONN HOUSING 4POS .100" CRIMP	4	$0.44		$1.76	$282.19
CONN HOUSING 6POS .100" DUAL	2	$0.53		$1.06	$281.13
CONN SOCKET 22-24AWG CRIMP GOLD	25	$0.34		$8.50	$272.63
Shipping + tax	1	$9.65		$9.65	$262.98
Acrylic Sheet	1	$13.48		$13.48	$249.50

PCB Estimate	2	$33.00		$66.00	$183.50
3D Printing ESTIMATE	1	$240.00		$240.00	-$56.50

Face2Gene: Using Facial Recognition to Aid in Diagnosing Rare Genetic Disorders

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Amanda Rein
Spring, 2021

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Signature ______________________________________ Date __4/20/2021____
Amanda Rein

Approved _____ ______________________________________ Date ___5/4/2021____
Hannah Rogers, Department of Engineering and Society

Abstract

Face2Gene utilizes facial recognition technology to analyze a “selfie” of a patient, then comes up with a list of most likely diagnoses. With over 6,000 rare genetic disorders that are often accompanied with a variety of subtle but distinct facial cues, diagnosticians need assistance in swiftly and accurately providing a diagnosis for people with rare genetic disorders. As this application has been built on a machine learning model, it has been molded to suit the accuracy and diversity needed by medical professionals and patients alike. As such, Face2Gene provides a rich subject for analysis under the framework social construction of technology (SCOT) because it is an application that has been shaped by human action. This paper will provide background information on Face2Gene, then analyze how the application has grown since its launch in 2014 under the SCOT pillars of interpretative flexibility, wider context, and closure & stabilization.

Introduction

Facial recognition technology has come a long way since 1964, when scientists first trained computers to recognize the human face (NEC, 2020). The vast improvements in camera quality, computing power, and machine learning algorithms over the past few decades have made facial recognition immensely powerful. While utilization of facial recognition technology is most often associated with law enforcement, it has proved to be beneficial in a vast majority of industries. For example, facial recognition is being used in the healthcare industry to speed up patient check-in at medical facilities, track a nonverbal patient's pain via their expression (and adjust their medication accordingly), and even to aid in diagnostics (Crompton & Diamond, 2019).

Facial recognition has the potential to solve a major problem in the healthcare industry by assisting in diagnosis of rare genetic disorders. The number of rare genetic disorders has surpassed 6,000, and each new disorder comes with an array of subtle symptoms of which diagnosticians cannot always detect nor recall. As the number of people who currently live with rare genetic disorders exceeds 300 million, there is a critical need for swift diagnoses of these disorders. That being said, patients with rare genetic disorders usually wait between 5.6 and 7.6 years to receive a proper diagnosis (Vandeborne et al., 2019). This is where facial recognition technology can come in handy, as 30% to 40% of these genetic disorders involve detectable changes to the face or skull (*Medical products: Facial recognition technology to diagnose rare genetic diseases*, 2015).

Face2Gene is the one of the first applications that utilizes facial technology to diagnose these types of rare disorders. In fact, its creators were the inventors of the first facial recognition technology that actually outperformed humans. Face2Gene allows smartphone users, physicians and patients alike, to take self-photographs and generate a list of the most likely diagnoses based on detected facial features. Since Face2Gene is built atop machine learning, the model increases in accuracy and range the more that it is utilized. (Wallner, 2017).

The application Face2Gene is a prime example of a technology that has been shaped by its users, making it a prime technology for study under the framework social construction of technology (SCOT). This research paper will first provide background information on the SCOT framework, then will detail how Face2Gene has been molded to fit its users' unique needs via the SCOT pillars: interpretive flexibility, wider context, and closure/stabilization. With an analysis under the SCOT framework, it will become clear that Face2Gene is a technology shaped by human action because of how its users have enhanced the accuracy and diversity of the application since its launch.

Conceptual Framework: SCOT

SCOT is a framework that emerged in 1987 from *The Social Construction of Technological Systems*, written by Wiebe Bijker and Trevor Pinch. The framework initially consisted of four related components: interpretive flexibility, relevant social group, closure and stabilization, and wider context. Interpretive flexibility suggests an open design process, with intergroup negotiations forming a product that is meaningful for all user groups. These user groups are called relevant social groups. A relevant social group encompasses one particular interpretation of the technology; in other words, members of a group assign the same meaning to the technology. The design phase is finally complete once each relevant social group has decided that the technology will meet their needs. The third component of SCOT, closure and stabilization, pertains to conflicts in the design of a technology (in this case, the application Face2Gene). If a relevant social group sees a problem with the current design, the design phase is not over. Closure is finally achieved when the design phase has ended, which can only occur when the technology is considered noncontroversial by all relevant social groups. Pinch and Bijker categorized two kinds of closure and stabilization: closure by redefinition and rhetorical closure. In closure by redefinition, closure is achieved when the problem is redefined so that the problems posed by social groups are no longer relevant. In rhetorical closure, a decision is made that problems no longer exist; in other words, all groups agree to consider the problem solved in its original context. Lastly, wider context encompasses the cultural and political environment in which development occurs. It is essentially the background information necessary to understand the motives and interactions between relevant social groups, as well as external constraints on the design. It will be clear why SCOT is an appropriate framework for analyzing Face2Gene in the subsequent sections of this thesis, in which the application is explained within each pillar of SCOT.

Many scholars have since critiqued SCOT for its over simplicity. For example, all relevant social groups rarely have an equal say in the design process, nor are all members of a group identical in their beliefs. This is particularly true with the development of Face2Gene's machine learning model, as they began training the model with only a few dozen clinicians. As such, the training of the Face2Gene model in the three years between its beta launch in 2014 and public launch in 2017 was limited in perspective and scope (Mack, 2017). As such, the entire relevant social group of patients and their families did not get a say in the design process of the application. Additionally, the relevant social group of medical professionals was not necessarily given equal say, as the vast majority of that group did not have the chance to beta test of the application. As such, Face2Gene is not entirely adherent to this particular SCOT concept. That being said, the application's founders and initial clinician beta testers made an intentional effort to diversify the application, such that the Face2Gene database included phenotypic/genotypic information for over 7,000 rare genetic disorders at its public launch (Mack, 2017). By training Face2Gene to recognize rare disorders that they themselves may not have ever needed to diagnose/treat, the stakeholders that had a voice in the design process also incorporated the needs of patients and clinicians beyond their own circles.

Another concept was later added to the SCOT framework called technological frame, which attempts to address these criticisms by providing room for more analysis of the structure of thinking within a relevant social group (Klein and Kleinman, 2002). Without this concept, there was little room for the technical background to be incorporated into a SCOT analysis. While the technological frame adds more realistic complexity to analysis, it will not be mentioned in this particular thesis. Instead, the subsequent analysis of Face2Gene will address the four original pillars of SCOT. After defining Face2Gene under each of these pillars and drawing the paper to a close with a conclusion section, it will be clear how the application has been shaped by human action rather than itself making an impact on society.

Interpretative Flexibility

The relevant social groups must first be defined, in order to demonstrate how Face2Gene applies to other pillars of SCOT. The relevant social groups for the application Face2Gene include: the developers of Face2Gene (a startup company called FDNA), medical professionals tasked with diagnosing rare genetic disorders, and their patients (as well as patients' loved ones) who seek diagnosis. Note that these groups are inferred based on the major users of the application and may not be entirely comprehensive (but will be suffice for this analysis). These groups were inferred based on the major users of the application and those benefitting from its success.

These relevant social groups see Face2Gene as a means to different ends. For the developers of Face2Gene, they see their application as a means to aid medical professionals in diagnostics and patients in getting answers to their symptoms. The “About” section of the FDNA site includes the statement “Since its founding in 2011, FDNA continues to aid clinicians, researchers and genetic testing labs in finding answers and treatments for hundreds-of-millions of patients globally living with a genetic disease” (*About FDNA*, 2021). That being said, the founders of FDNA also saw a gap in the market that would allow them to achieve personal career success and worldwide recognition. On their website, FDNA states “Our MISSION is to disrupt clinical genomics by integrating NGP into the genetic testing workflow” (*About FDNA*, 2021).

For medical professionals, Face2Gene is an invaluable aid to diagnosis. The meaning of the application hinges upon its ability to make suggestions on the subtle facial cues that can be hard for physicians to detect or connect to one of the thousands of rare genetic disorders with such features. In an article that was published in the British Journal of General Practice, general practitioner and author Willian Evans writes “Knowledge of all 6000-8000 rare diseases, increasing each week, is impossible for any one individual. The response ‘How can we know about all rare disease?’ is not only true but also feeds into the cognitive barrier that prevents clinicians contemplating a rare disease at all” (Evans, 2018). Karen Gripp, MD and medical geneticist as Nemours/Alfred I. duPont Hospital for Children and professor of pediatrics at Thomas Jefferson University, stated in a 2020 interview “Sometimes the typical facial features [of a genetic disorder] are present, but they are so mild that a human expert has trouble identifying them. An algorithm is at times superior at identifying these” (Grifantini, 2020).

For patients with rare genetic disorders, as well as their loved ones, they see Face2Gene as an unbiased, reliable tool for compiling a list of likely diagnoses with just a selfie as input. As stated by FDNA founder Dekel Gelbman, “Patients living with a genetic disease often times go years without receiving a diagnosis, which means they are also not receiving the appropriate treatment or therapies that are available for some of these syndromes. By aiding in reaching a faster, more accurate diagnosis, patients can seek appropriate treatments and care that can improve their quality of life and extend their lives” (AI Time Journal Editorial Staff, 2019). As the various relevant social groups attached to Face2Gene are assigning different meanings to the application, it becomes necessary for all of these separate use cases and interests to be equally accounted for in the design process. Since it is rarely the case that all stakeholders have equal role in the design process and because Face2Gene unlike most technologies is actually actively shaped as it is utilized by stakeholders (the model strengthening to fit all users’ needs as it is utilized), Face2Gene is still applicable under SCOT’s interpretative flexibility.

Wider Context

The SCOT pillar referred to as “wider context” encompasses all sociocultural and political backgrounds taking place at the time of the development and adoption of the technology. This is where it is especially clear that Face2Gene was not a technology with a purpose that was fixed during the design phase. The application was launched in 2014 in “stealth mode” to gather enough data to train the system, with a few dozen beta users sharing patient photos. By 2017 when the application was officially launched to the public, the application’s database already had information for over 7,000 rare diseases. Face2Gene expressed a desire to create the world’s largest repository of rare disease big data through users inputting selfies to

Face2Gene. Evidently, the founders of the application saw the value in allowing users to morph the application's model to fit their own needs (Mack, 2017).

A multitude of researchers and physicians have since published results of utilizing Face2Gene to diagnose various disorders. In January of 2016, the first study was published based on initial findings in the accuracy of Face2Gene in recognizing Cornelia de Lange syndrome. The average detection rate of medical experts was measured to be around 77%, while the new Face2Gene technology had a detection rate of 94%. (Basel-Vanagaite et al., 2016). Another study utilizing Face2Gene was released several months later, in December of 2017, by a group of researchers that ran tests of its ability to detect Down Syndrome. The application showed recognition of Down Syndrome in Caucasians with 80% accuracy, compared to Africans at 36.8% accuracy. The researchers then added more training data to the application that included cases of Africans with and without Down Syndrome and were able to increase the recognition to 94.7% accuracy. While this study revealed bias in the original data set supplied alongside Face2Gene, it also revealed how further training for particular ethnic groups makes the application incredibly powerful for any group. (Lumaka et al., 2017). In November of 2017, another group of researchers released results of a study applying Face2Gene to detecting two disorders: Emanuel (ES) and Pallister-Killian Syndrome (PKS) which further proved the breadth of knowledge that Face2Gene is capable of holding (Liehr et al., 2018). In December of 2017, just months after the Face2Gene application was launched, a group of researchers published a study in the Official Journal of the American Academy of Pediatrics detailing the use of Face2Gene to diagnose fetal alcohol spectrum disorders (Valentine et al., 2017). FDNA also published their own study in January 2019, utilizing Face2Gene's image analysis framework, DeepGestalt, with proof of the platform outperforming clinicians in three experiments. This particular study highlights improvements in the DeepGestalt model, including the dataset's growth to 17000 images representing 200 syndromes. The increased efforts to grow the training set resulted in 91% top-10 accuracy in identifying a person's genetic disorder in the same study, which utilized 502 different images to come up with this result. (Gurovich et al., 2019). Through studies such as these, it is evident how the underlying algorithm of Face2Gene is capable of being morphed to fit the needs of its users both in accuracy of diagnosis and breadth of knowledge.

Closure and Stabilization

SCOT also emphasizes closure and stabilization as signs of a technology shaped by society. As mentioned in the Conceptual Framework section, are two types of closure under the SCOT framework: rhetorical and redefinition of the problem. Face2Gene has experienced closure under redefinition of the problem, since it is merely an *aid* to diagnosis rather than a tool to provide a certain, singular disorder as the cause of a patient's symptoms. The context has been switched to aiding diagnosticians, rather than replacing them. Donald Basel, MD and medical director of the Genetics and Genomics Program, said about Face2Gene "It never gives you a diagnosis. It gives you an answer to consider". With the list of the ten most likely possibilities for diagnosis provided by the Face2Gene analysis, as well as the patients' medical history, Basel can make more confident and speedy diagnoses. According to Basel, "Machines can never replace humans. The power of the machine is only as good as the information going in and the people who interpret it. The machine can only take all of the data and process it, but it can't prioritize it". As of 2019, Basel and his colleagues had utilized Face2Gene to diagnose over 1,000 patients (Etter, 2019). Clearly more than Basel and his colleagues have accepted

Face2Gene; it has been adopted by over 70% of the world's geneticists across 2,000 clinical sites and in over 130 countries (*About FDNA*, 2021).

It should be noted that relevant social groups such as patients, their loved ones, and medical professionals have questioned the privacy of Face2Gene. Face2Gene has a comprehensive Data Sharing & Protection Policy that addresses any ethical concerns once raised by relevant social groups. To begin, patients' personal identifiers such as names, addresses, and social security numbers are not collected by the application. There are still concerns around the facial images and associated test results or any other information inputted by the physician (including, but not limited to, clinical observations and family history). As such, some may argue that Face2Gene does not make sense as an object for SCOT analysis as it has not been accepted entirely by the relevant social group encompassing patients and their loved ones, calling into question whether Face2Gene is an application shaped by its users. There is certainly a valid counterargument that Face2Gene founders did not incorporate users' opinions into the initial design of the application, and if they had then perhaps the application would not have utilized facial recognition at all. That counter argument certainly implies that users did *not* shape Face2Gene, nor is it an example of SCOT. That being said, Face2Gene's founders have since met the concerns of users to achieve closure for their application. In order to address privacy issues, Face2Gene's developers installed a "de-identification" process for facial photos in their data set so that it is not possible for someone who gains access to the application database to discern someone's identity. The application is applicable to HIPAA, EU data protection regulations, and meets the concerns of the relevant social groups (*Data Sharing & Protection Policy,* 2017). As such, a consensus has been established, achieving closure amongst the relevant social groups and ending the design phase of Face2Gene.

Conclusion

As technology becomes more powerful and pervasive, people view technological artifacts themselves as holding the power to shape the structure of society. That being said, this concept of technological determinism is often not accurate, especially when it comes to technologies that employ machine learning. As machine learning algorithms improve upon themselves by user input, they specifically grow in accuracy and diversity based on users' needs. As such, users hold a great deal of power in molding these applications, even if they were not directly consulted during their initial design. Face2Gene is an example of a technology that has been directly shaped by human action. Its founders incorporated the other major relevant social groups, physicians and their patients, in the design process by launching in "stealth mode" to give the application's algorithm years to be shaped by its users to best fit their needs. Now that the application is public, it still is being shaped by the multitude of studies undertaken on it, growing more accurate where it is most needed by society. Through this analysis of Face2Gene under the SCOT framework, it is clear that Face2Gene is an application that has been shaped by human action because of how its users have molded the application to suit their needs.

In the future, it is likely that more applications of machine learning are created and subsequently popularized to meet society's needs. Face2Gene itself will continue to increase in accuracy the more data that it is provided (a.k.a. the more that its users utilize it), and it will also grow in scope as it is provided with data for new kinds of rare disorders. As CEO of FDNA Dekel Gelbman said in a 2017 interview with StartUs magazine, "The majority of clinical geneticists around the world are already using our technology. It is also used by a large body of researchers, drug developers and genetic testing labs working to find answers and treatments for hundreds-of-millions of patients globally. In the future, every person's genome will serve as their medical record. Our technology will help patients are their caregivers understand their health better and come up with a personalized care plan – this is our role in the future of precision medicine" (Wallner, 2017). Face2Gene will continue to be a prime example of SCOT as it is being directly shaped by society's needs.

References

A brief history of facial recognition - NEC New Zealand. (2020, May 26). NEC. https://www.nec.co.nz/market-leadership/publications-media/a-brief-history-of-facial-recognition/

About FDNA. (n.d.). *FDNA*. Retrieved March 5, 2021, from http://fdna.com/about-us/

AI Time Journal Editorial Staff. (2019, September 9). AI in Healthcare: Interview with Dekel Gelbman, CEO of FDNA. *AI Time Journal*. https://www.aitimejournal.com/ai-in-healthcare-interview-with-dekel-gelbman-ceo-of-fdna

Basel-Vanagaite, L., Wolf, L., Orin, M., Larizza, L., Gervasini, C., Krantz, I. D., & Deardoff, M. A. (2016). Recognition of the Cornelia de Lange syndrome phenotype with facial dysmorphology novel analysis. *Clinical Genetics*, *89*(5), 557–563. https://doi.org/10.1111/cge.12716

Crompton, P., & Diamond, A. (2019, August 15). *What are the Key Use Cases for Facial Recognition in Healthcare?* Vantage Technology Consulting Group. https://www.vantagetcg.com/what-are-the-key-use-cases-for-facial-recognition-in-healthcare/

Data Sharing & Protection Policy. (2017, March 28). *Face2Gene*. https://www.face2gene.com/data-sharing-protection-policy/

Etter, N. S. (2019, April 9). *Facial recognition and the future of diagnosis*. Children's Wisconsin. https://childrenswi.org/newshub/stories/face2gene-facial-recognition-diagnosis

Grifantini, K. (2020, May 13). Detecting Faces, Saving Lives: How facial recognition software is changing health care. *IEEE Pulse*, *March/April 2020*. https://www.embs.org/pulse/articles/detecting-faces-saving-lives/

Gurovich, Y., Hanani, Y., Bar, O., Nadav, G., Fleischer, N., Gelbman, D., Basel-Salmon, L., Krawitz, P. M., Zenker, M., Bird, L. M., & Gripp, K. W. (2019). Identifying facial phenotypes of genetic disorders using deep learning. *Nature Medicine*, *25*(1), 60–64. https://doi.org/10.1038/s41591-018-0279-0

Klein, H. K., & Kleinman, D. L. (2002). The Social Construction of Technology: Structural Considerations. *Science, Technology, & Human Values*, *27*(1), 28–52.

Liehr, T., Acquarola, N., Pyle, K., St-Pierre, S., Rinholm, M., Bar, O., Wilhelm, K., & Schreyer, I. (2018). Next generation phenotyping in Emanuel and Pallister-Killian syndrome using computer-aided facial dysmorphology analysis of 2D photos. *Clinical Genetics*, *93*(2), 378–381. https://doi.org/10.1111/cge.13087

Lumaka, A., Cosemans, N., Mampasi, A. L., Mubungu, G., Mvuama, N., Lubala, T., Mbuyi-Musanzayi, S., Breckpot, J., Holvoet, M., Ravel, T. de, Buggenhout, G. V., Peeters, H., Donnai, D., Mutesa, L., Verloes, A., Tshilobo, P. L., & Devriendt, K. (2017). Facial dysmorphism is influenced by ethnic background of the patient and of the evaluator. *Clinical Genetics*, *92*(2), 166–171. https://doi.org/10.1111/cge.12948

Mack, H. (2017, March 21). *FDNA launches app-based tool for clinicians using facial recognition, AI and genetic big data to improve rare disease diagnosis and treatment*. MobiHealthNews. https://www.mobihealthnews.com/content/fdna-launches-app-based-tool-clinicians-using-facial-recognition-ai-and-genetic-big-data

Medical products: Facial recognition technology to diagnose rare genetic diseases. (2015, December 1). UK Research and Innovation. https://mrc.ukri.org/news/browse/medical-products-facial-recognition-technology-to-diagnose-rare-genetic-diseases/

Valentine, M., Bihm, D. C. J., Wolf, L., Hoyme, H. E., May, P. A., Buckley, D., Kalberg, W., & Abdul-Rahman, O. A. (2017). Computer-Aided Recognition of Facial Attributes for Fetal Alcohol Spectrum Disorders. *Pediatrics*, *140*(6). https://doi.org/10.1542/peds.2016-2028

Vandeborne, L., van Overbeeke, E., Dooms, M., De Beleyr, B., & Huys, I. (2019). Information needs of physicians regarding the diagnosis of rare diseases: a questionnaire-based study in Belgium. *Orphanet Journal of Rare Diseases*, *14*(1), 99. https://doi.org/10.1186/s13023-019-1075-8

Wallner, S. (2017, October 3). *Precision Medicine Startup FDNA Introduces Next-Generation Phenotyping (NGP)*. https://magazine.startus.cc/precision-medicine-startup-fdna-introduces-next-generation-phenotyping-ngp/

Creating a Temperature-Tracking Door Locking System for COVID-19
(Technical Paper)

Face2Gene: Using Facial Recognition to Aid in Diagnosing Rare Genetic Disorders
(STS Paper)

A Thesis Prospectus Submitted to the

Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements of the Degree
Bachelor of Science, School of Engineering

Amanda Rein
Fall, 2020

Technical Project Team Members
Andy Hui
Greg Vavoso
Jack Shefer
Matthew Bain

On my honor as a University Student, I have neither given nor received
unauthorized aid on this assignment as defined by the Honor Guidelines
for Thesis-Related Assignments

Introduction

1 in 3 patients with rare genetic disorders wait at least two years to receive a diagnosis according to The Genetic Alliance UK. A 2019 study in the United States and United Kingdom estimated that it takes even longer, between 5.6 to 7.6 years on average. Physicians simply cannot keep up with the 6,000 and counting rare genetic disorders, despite many of them possessing characteristic facial cues (Vandeborne et al., 2019). A mobile application called Face2Gene is utilizing facial recognition technology to drastically reduce the amount of time it takes for physicians to diagnose these rare disorders. Not only does this application of facial recognition save physicians time; patients and their loved ones also save enormous amounts of time, and money, that would have been spent traveling to specialists and undergoing expensive testing. This sociotechnical project will analyze Face2Gene as a technology that has been shaped by societal needs, rather than shaping society itself.

Automation technology is being applied to save time in other areas of healthcare, as well. During the COVID-19 pandemic, individuals not trained in medical professions have had to donate their time to enforcing public health policies. Office managers, security guards, and retail workers are just a few types of people that are manning doorways to scan entrants' temperatures in an attempt to prevent a contagious individual from infecting the space. Automation technology in the form of this proposed Capstone project will free up these individuals' time by automating the process of scanning a person's temperature and controlling their access to the space accordingly.

Technical Topic (Capstone)

The goal of this technical project is to help combat the spread of the novel coronavirus (COVID-19), a deadly illness with over eight million reported cases and more than two hundred thousand deaths in 2020 (CDC, 2020b). As economies start to reopen, concern for public health has given rise to the demand for technology that automates finding and monitoring cases (Morrissey, 2020). Contact tracing mobile applications have emerged to try and meet this need (COVIDWISE, 2020). These applications keep track of nearby devices in order to build network of which users have been in close proximity to others. Thus, when a user tests positive, they can privately and anonymously notify all other devices (a.k.a. users) that were nearby. However, these technologies are limited by the number of users they can attract and the willingness of these users to report positive test results. One of the major goals of this project is to create a technology that does not rely on user motivation. Rather than having users actively enroll in public safety measures, the default state should be participation.

There are many stakeholders in this project, incentivized to ensure public safety and automate enforcement of these measures. Brick-and-mortar businesses, such as restaurants and retail stores, are likely candidates. Their motivations are two-fold. First, in order for customers to return, they must feel safe; therefore, increasing safety steps may increase in-person traffic and subsequently profit. Second, business owners are willing to implement these safety measures to protect themselves and their employees from illness and potential legal liability (Elejalde-Ruiz, 2020). This growing desire for new safety measures applies to office spaces, as well, as employers look to keep themselves and their employees healthy. Additionally, since these brick-and-mortar businesses and office spaces are often privately-owned businesses, they are free to enforce limits on who can enter their property. Other possible stakeholders are healthcare centers and educational institutions. Healthcare centers, such as the UVa Elson Student Health Center, have been asking visitors to wait outside until someone arrives to manually measure their temperature before admittance. This process of waiting to be scanned before entering relies upon visitors having the patience to wait outside, rather than entering through one of the other entrances. This is an inefficient system, especially for the staff of the healthcare center who must constantly operate the entrance.

This project will prototype an automated temperature screening system to control who can enter a space. This system will be attached to a door-frame. Visitors will approach door and scan their temperatures by holding their wrists under a non-contact temperature sensor. If their temperature is within a predetermined bound, the user will be admitted to the facility; if not, the door will be automatically locked to prevent their entrance. Their temperature reading will be sent via Bluetooth to a server that will store the reading. This collection of stored data will be rendered on a web application that tracks usage and temperature measurements of people attempting to enter the space through that doorway, enabling the owners of the space (and other stakeholders) to have an overall idea of the health of its visitors. Since this web application will display the number of people that have been admitted to the space, it may also be helpful for determining the number of people in the space at a particular time (i.e. for students wanting to avoid peak times at the gyms or libraries). While temperature is only one indicator of a person's wellbeing, it is still one of the most powerful tools for detecting if someone carries a communicable disease, making this system a highly effective way to automate screening the health of individuals before they are admitted to a space (CDC, 2020a).

STS Topic

Over 300 million people currently live with a rare genetic disorder. While this statistic only amounts to approximately 5% of the world population, these disorders significantly reduce the quality of life for hundreds of millions of people. Diagnosis of these disorders is incredibly difficult due to their uncommon nature, as well as their deceptively-common symptoms

(Nguengang et al., 2020). Patients with rare genetic disorders do not receive a proper diagnosis for several years on average, as physicians cannot maintain a knowledge of the symptoms associated with the ever-growing number of rare disorders (Evans, 2018). As such, it is not just patients and their loved ones that are stakeholders in the development of faster, more accurate ways to diagnose rare genetic disorders; medical professionals are also eager to correctly and efficiently diagnose their patients' disorders. While genetic testing is an option for diagnosing genetic disorders, it is very expensive and time-consuming, especially without an idea of the few most likely culprits. There are rare disorder search engines on the Internet that aid in the diagnosis process, as well, but they rely upon the physician's ability to detect all symptoms accurately.

Recent advances in facial recognition technology have brought great hope to increasing the speed and accuracy of diagnosis of these rare genetic disorders. While 30 to 40% of genetic disorders involve detectable changes to the face or skull, these facial cues may be subtle and not well-known to medical professionals (*Medical products: Facial recognition technology to diagnose rare genetic diseases*, 2015). One particular application, developed in 2016, has become incredibly useful in detecting these facial cues and subsequently suggesting likely diagnoses: Face2Gene. This application has been trained on tens of thousands of case files, growing even more accurate every time that it is used due to the nature of computer vision models (Wallner, 2017).

The adoption of Face2Gene is an example of social construction of technology (SCOT) as a technology that has been shaped by its users. SCOT is a popular STS framework which posits that human action shapes technology. The framework was designed with four key pillars: interpretive flexibility, relevant social group, closure and stabilization, and wider context. SCOT is critiqued by a variety of scholars, including its original authors, for its grouping of society into "relevant social groups" that are assumed to be of equal importance and to have equal voice in the design process. Authors of SCOT have attempted to address these critiques of the framework being overly simplistic, adding another major tenant to SCOT: technological frame. While SCOT is still an imperfect framework, it provides for a rich analysis of the rapid growth of Face2Gene (Klein and Kleinman, 2002).

In summation, this paper will explain how Face2Gene's development and adoption by the medical community aligns with the major components of SCOT: interpretive flexibility, relevant social group, closure and stabilization, wider context, and technological frame. By exploring these components in the context of Face2Gene, it will become clear how human action has shaped this technology.

Research Question and Methods

Research Question: "How has the facial recognition technology Face2Gene improved diagnosis of rare genetic disorders?"

This research question will be answered using an STS framework entitled social construction of technology (SCOT). First, the research question will be framed with the appropriate background information about the challenges of diagnosing rare genetic disorders, as well as the spread of facial recognition technology as a whole. After that, an explanation will be provided on how the adoption of the Face2Gene application is an example of social construction of technology. To do a full SCOT analysis, it will be necessary to first define the relevant social groups: patients with rare genetic disorders and their loved ones, medical professionals searching for diagnoses for these patients, and the scientists responsible for creation of Face2Gene. Then, the analysis will launch into how the application's development and adoption fits into the main pillars of SCOT: interpretive flexibility, relevant social groups, closure and stabilization, wider context, and technological frame.

For interpretive flexibility, it will be necessary to include a timeline for the development and adoption of facial recognition as a whole, to point out how facial recognition has been around for decades and yet has only applied to diagnosing rare genetic disorders in the past few years ("A brief history of facial recognition", 2020). Interpretive flexibility also involves *how* relevant social groups assign meaning to the technology. Addressing this component of interpretive flexibility will require an overview of general sentiment around Face2Gene, as well as the ways in which it has supported its relevant social groups. This will necessitate the use of sources like news and magazine articles, or even quotations from those that have benefited from the application (Etter, 2019) (Grifantini, 2020) (Mjoseth, 2017) (Vincent, 2019).

To address relevant social group, motivations of the social groups defined previously will be outlined, in order to show how they align with the adoption of Face2Gene. Again, this connection requires a general idea of the attitude around Face2Gene rather than sources that provide quantitative data (Etter, 2019) (Grifantini, 2020) (Mjoseth, 2017) (Vincent, 2019).

To address closure and stabilization, it will be important to note how Face2Gene received HIPAA and EU data privacy compliance via sources published by the founders of Face2Gene. Additionally, Face2Gene's method of "de-identifying" the facial photos in their data set in order to maintain individuals' privacy will be highlighted to make it clear how Face2Gene received acceptance by relevant social groups and therefore was able to be adopted by the medical community ("Data Sharing & Protection Policy, 2017) (Martinez-Martin, 2019).

For the wider context, relevant sociocultural and political background at the time that Face2Gene was being released will be laid out. This part of the analysis will include background from news articles and blog posts about the climate surrounding Face2Gene. Most importantly, the analysis of Face2Gene as a form of SCOT will be solidified by case analyses of how the application has been used to aid in diagnosis of a wide variety of genetic disorders, extending beyond its original intent and abilities. A multitude of studies published in medical journals regarding the utilization of Face2Gene to aid in diagnosis of rare genetic disorders have already been collected (Basel-Vanagaite et al., 2016) (Jin et. al, 2020) (Kosilek et al., 2015) (Liehr et al., 2018) (Nguengang et al., 2020) (Stephen et al., 2017) (Valentine et al., 2017).

Finally, resources that highlight the launch and marketing strategy of Face2Gene will detail the technological frame. Including this particular pillar of SCOT will mitigate criticisms of the framework by providing additional detail about the technological strategy utilized by Face2Gene's founders that enabled its success ("2017 Year in Review", 2018) ("2019 Highlights", 2020) ("FDNA and Microsoft Collaborate to Enhance Genomics Technology Through AI", 2018) ("GDNA and PerkinElmer Announce Collaboration to Offer Enhanced Genetic Testing Augmented by AI and Facial Analysis", 2019) ("FDNA Partners with Blueprint Genetics to Spotlight RASopathies During the Year of Discovery", 2017) (Wallner, 2017) ("What to do and See at ACMG 2017 in Phoenix, AZ", 2017) ("What to Do and See at ESHG 2017 in Copenhagen", 2017). Through the analysis of Face2Gene from each of the pillars of SCOT, it will be clear that Face2Gene has been shaped by its relevant social groups.

Conclusion

These two projects address critical applications of technology in healthcare, the first an analysis of facial recognition as a tool to aid in diagnosing rare disorders and the second a prototype for an automated door locking system that prevents people from entering a room if their body temperature is irregular. The analysis of a facial recognition mobile application that aids in diagnosing rare genetic disorders focuses on a technology that provides answers to millions of people suffering from undiagnosed (and therefore untreated) disorders. The analysis of Face2Gene provides evidence for how technology can be shaped by society to fit its needs. Similarly, the team creating the door locking system will prioritize mitigating public health concerns, in this case focused on the current COVID-19 pandemic. The invention will ensure that indoor spaces do not include individuals with high body temperatures, as high fevers are often indicative of infectious diseases such as COVID-19. In summation, both of the endeavors included in this portfolio focus on technology as a means to improve global public health.

References

Capstone

CDC. (2020a, February 11). *Coronavirus Disease 2019 (COVID-19)*. Centers for Disease Control and Prevention. https://www.cdc.gov/coronavirus/2019-ncov/hcp/guidance-risk-assesment-hcp.html

CDC. (2020b, March 28). *COVID-19 Cases, Deaths, and Trends in the US | CDC COVID Data Tracker*. Centers for Disease Control and Prevention. https://covid.cdc.gov/covid-data-tracker

COVIDWISE. (2020). Virginia Department of Health. https://www.vdh.virginia.gov/covidwise/

Elejalde-Ruiz, A. (2020, May 4). If you get sick with COVID-19, is your employer liable? As businesses prepare to reopen, worker safety is a priority. *Chicago Tribune*. https://www.chicagotribune.com/coronavirus/ct-coronavirus-employer-liability-workplace-exposure-20200501-dye6husnszchpnpaadiensn2ja-story.html

Morrissey, J. (2020, June 16). Fighting the Coronavirus With Innovative Tech. *The New York Times*. https://www.nytimes.com/2020/06/16/business/fighting-covid-19-innovative-tech.html

STS

2017 Year in Review. (2018, January 31). https://www.fdna.com/blog/2017-year-in-review/

2019 Highlights. (2020, January 20). *FDNA*. https://www.fdna.com/blog/2019-highlights/

A brief history of facial recognition - NEC New Zealand. (2020, May 26). NEC. https://www.nec.co.nz/market-leadership/publications-media/a-brief-history-of-facial-recognition/

Basel-Vanagaite, L., Wolf, L., Orin, M., Larizza, L., Gervasini, C., Krantz, I. D., & Deardoff, M. A. (2016). Recognition of the Cornelia de Lange syndrome phenotype with facial dysmorphology novel analysis. *Clinical Genetics*, *89*(5), 557–563. https://doi.org/10.1111/cge.12716

Data Sharing & Protection Policy. (2017, March 28). *Face2Gene*. https://www.face2gene.com/data-sharing-protection-policy/

Etter, N. S. (2019, April 9). *Facial recognition and the future of diagnosis*. Children's Wisconsin. https://childrenswi.org/newshub/stories/face2gene-facial-recognition-diagnosis

Evans, W. (2018). Dare to think rare: diagnostic delay and rare diseases. *British Journal of General Practice*, *68*, 145. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5916061/

FDNA and Microsoft Collaborate to Enhance Genomics Technology Through AI. (2018, October 16). *FDNA*. https://www.fdna.com/blog/fdna-and-microsoft-collaborate-to-enhance-genomics-technology-through-ai/

FDNA and PerkinElmer Announce Collaboration to Offer Enhanced Genetic Testing Augmented by AI and Facial Analysis. (2019, June 25). https://www.businesswire.com/news/home/20190625005184/en/FDNA-and-PerkinElmer-

Announce-Collaboration-to-Offer-Enhanced-Genetic-Testing-Augmented-by-AI-and-Facial-Analysis

FDNA Partners with Blueprint Genetics to Spotlight RASopathies During the Year of Discovery. (2017, February 15). *FDNA*. https://www.fdna.com/blog/rare-diseases-blog-post/

FDNA Presents Rare Disease Technologies at the 2017 Precision Medicine Summit at Boston Children's Hospital. (2017, September 18). *FDNA*. https://www.fdna.com/blog/fdna-presents-rare-disease-technologies-2017-precision-medicine-summit-boston-childrens-hospital/

Grifantini, K. (2020, May 13). Detecting Faces, Saving Lives: How facial recognition software is changing health care. *IEEE Pulse*, *March/April 2020*. https://www.embs.org/pulse/articles/detecting-faces-saving-lives/

Jin, B., Qu, Y., Zhang, L., & Gao, Z. (2020). Diagnosing Parkinson Disease Through Facial Expression Recognition: Video Analysis. *Journal of Medical Internet Research*, *22*(7). https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7382014/

Klein, H. K., & Kleinman, D. L. (2002). The Social Construction of Technology: Structural Considerations. *Science, Technology, & Human Values*, *27*(1), 28–52.

Kosilek, R. P., Frohner, R., Wurtz, R. P., Berr, C. M., Schopohl, J., Reincke, M., & Schneider, H. J. (2015). Diagnostic use of facial image analysis software in endocrine and genetic disorders: review, current results, and future perspectives. *European Society of Endocrinology*, *173*(4). https://eje.bioscientifica.com/downloadpdf/journals/eje/173/4/M39.pdf

Liehr, T., Acquarola, N., Pyle, K., St-Pierre, S., Rinholm, M., Bar, O., Wilhelm, K., & Schreyer, I. (2018). Next generation phenotyping in Emanuel and Pallister-Killian syndrome using computer-aided facial dysmorphology analysis of 2D photos. *Clinical Genetics*, *93*(2), 378–381. https://doi.org/10.1111/cge.13087

Lumaka, A., Cosemans, N., Mampasi, A. L., Mubungu, G., Mvuama, N., Lubala, T., Mbuyi-Musanzayi, S., Breckpot, J., Holvoet, M., Ravel, T. de, Buggenhout, G. V., Peeters, H., Donnai, D., Mutesa, L., Verloes, A., Tshilobo, P. L., & Devriendt, K. (2017). Facial dysmorphism is influenced by ethnic background of the patient and of the evaluator. *Clinical Genetics*, *92*(2), 166–171. https://doi.org/10.1111/cge.12948

Martinez-Martin, N. (2019). What Are Important Ethical Implications of Using Facial Recognition Technology in Health Care? *AMA Journal of Ethics*, *21*(2), 180–187. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6634990/

Medical products: Facial recognition technology to diagnose rare genetic diseases. (2015, December 1). UK Research and Innovation. https://mrc.ukri.org/news/browse/medical-products-facial-recognition-technology-to-diagnose-rare-genetic-diseases/

Mjoseth, J. (2017, March 23). *Facial recognition software helps diagnose rare genetic disease*. National Human Genome Research Institute. https://www.genome.gov/news/news-release/Facial-recognition-software-helps-diagnose-rare-genetic-disease

Nguengang Wakap, S., Lambert, D. M., Olry, A., Rodwell, C., Gueydan, C., Lanneau, V., Murphy, D., Le Cam, Y., & Rath, A. (2020). Estimating cumulative point prevalence of rare diseases: analysis of the Orphanet database. *European Journal of Human Genetics*, *28*(2), 165–173. https://doi.org/10.1038/s41431-019-0508-0

Stephen, I., Hiew, V., Coetzee, V., Tiddeman, B., & Perrett, D. (2017). Facial Shape Analysis Identifies Valid Cues to Aspects of Physiological Health in Caucasian, Asian, and African Populations. *Frontiersin Psychology*, *8*, 1883. https://www.frontiersin.org/articles/10.3389/fpsyg.2017.01883/full

Valentine, M., Bihm, D. C. J., Wolf, L., Hoyme, H. E., May, P. A., Buckley, D., Kalberg, W., & Abdul-Rahman, O. A. (2017). Computer-Aided Recognition of Facial Attributes for Fetal Alcohol Spectrum Disorders. *Pediatrics*, *140*(6). https://doi.org/10.1542/peds.2016-2028

Vandeborne, L., van Overbeeke, E., Dooms, M., De Beleyr, B., & Huys, I. (2019). Information needs of physicians regarding the diagnosis of rare diseases: a questionnaire-based study in Belgium. *Orphanet Journal of Rare Diseases*, *14*(1), 99. https://doi.org/10.1186/s13023-019-1075-8

Vincent, J. (2019, January 15). *Facial recognition and AI could be used to identify rare genetic disorders*. The Verge. https://www.theverge.com/2019/1/15/18183779/facial-recognition-ai-algorithms-detect-rare-genetic-disorder-fdna

Wallner, S. (2017, October 3). *Precision Medicine Startup FDNA Introduces Next-Generation Phenotyping (NGP)*. https://magazine.startus.cc/precision-medicine-startup-fdna-introduces-next-generation-phenotyping-ngp/

What to Do and See at ACMG 2017 in Phoenix, AZ. (2017, February 15). *FDNA*. https://www.fdna.com/blog/events-blog-post/

What to Do and See at ESHG 2017 in Copenhagen. (2017, May 22). *FDNA*. https://www.fdna.com/blog/what-to-do-and-see-at-eshg/