TerraStep
(formerly LLPAT)
Final Report

Megan Harrison, Justin Fortner, Nicholas Rossi
Max Smalley, Cesar Vargas, Zixi Wang

March 16, 2020

Contents

1 Abstract

Lower limb prosthetic users lack underfoot sensory feedback, leading to instabilities that can cause injury. The TerraStep addresses this by introducing force sensors underfoot, and vibrotactile feedback motors on the thigh. The motors give continuous low intensity feedback that allows the user to develop familiarity with it, and high intensity, fast response alerts when a tripping hazard is detected underfoot. We have defined a tripping hazard as a 1.25 inch hard wooden sphere (for testing), with a tolerance of +/- 0.25 inches. This increased sensory communication and immediate alert allows the user to adjust their body weight to remain balanced when encountering a tripping hazard. Our results show that the system can detect and alert the user of underfoot objects within 160ms of initial contact.

2 Background

According to Advanced Amputee Solutions [6], an amputation occurs every 30 seconds. Two million amputees currently live in the United States, 1.7 million of whom are lower-limb prosthetic users, the significance of these statistics can be seen in Figure 1. Advancements in prosthetic technology have allowed these users to remain highly active, however some difficulties remain.

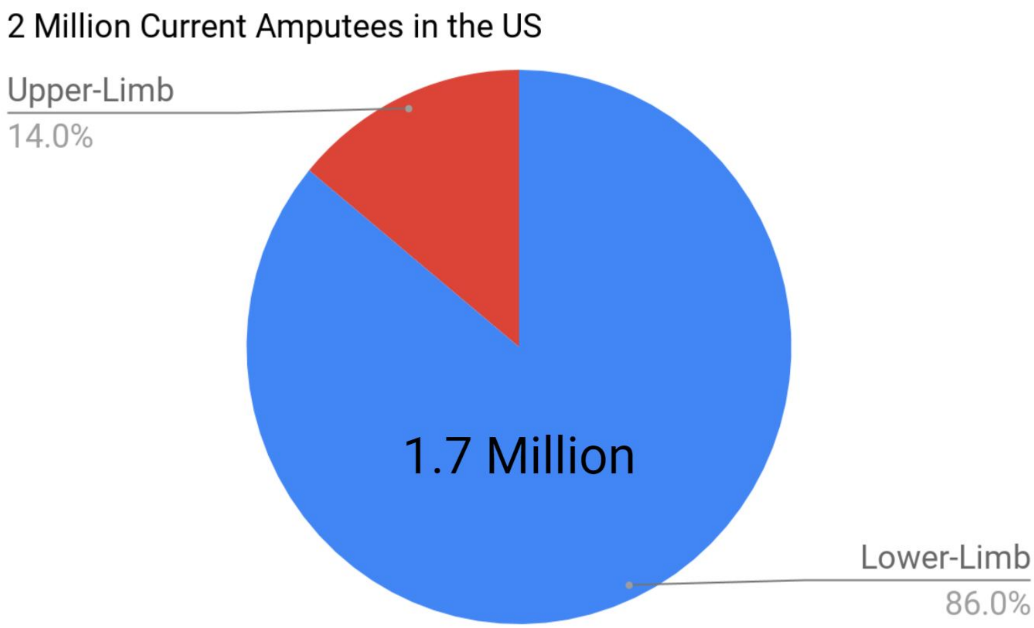

Figure 1: *Shows the breakdown of how many lower limb amputees vs upper limb amputees currently reside in the United States. Made using Microsoft Excel.*

Hiking with a lower limb prosthetic can be problematic because there is a lack of sensory feedback from the bottom of the prosthetic foot. Active prosthetic user and blogger Angelina Boulicault has said that "when there are rocks beneath me I often lose my footing because I cannot feel the instability beneath my leg" [2]. Even with current prosthetic technologies, active users still face a significantly higher risk of injury because they do not know what they are stepping on. There is a lack of sensory communication between the user and their prosthetic. The TerraStep system is aimed at improving this communication. It uses our proprietary algorithm to predict tripping hazards and quickly alert the user. This sensory feedback assists the user by providing fast, accurate information needed so that they can stay upright and avoid injury.

Object size is an important factor when determining a potential tripping hazard. Via experimentation and feedback from prosthetic users, we determined that objects that are approximately 1.25 inches in diameter, and taller than 1 inch are problematic. Anything larger is easily seen, and testing has shown that objects smaller than 1.0 inches in diameter will not destabilize gait. For this document and our testing, we've defined a tripping hazard as a 1.25 inch inelastic sphere with a tolerance of +/- 0.25 inches.

The Terra-Step system is focused on the large market of below-knee amputees. The inspiration for this project came from an active prosthetic user named Jim Brown, also a below-knee amputee. 86% of the two million amputees in the US are lower limb, and of the 200,000 new lower limb amputees annually, 97% are below-knee[5]. This means over one million below-knee amputees in the US today, and that number is growing rapidly. This is a large user base that would directly benefit from the TerraStep sensory feedback system. With some minor modifications, the system could be expanded to above-knee amputees (see Future Work section).

We use vibrotactile motors, similar to those found in phones, to provide feedback to the user's thigh. This feedback is informed by force sensors positioned underneath the prosthetic foot. Normal steps will result in low intensity vibrotactile feedback, allowing the user to develop an intuitive sense for the feedback. When a tripping hazard is is detected, a higher intensity vibrotactile pulse is used to alert the user of the object, and to convey its relative location underfoot.

Please note that the TerraStep is not designed to reproduce the full tactile sensations of a biological foot, but rather to provide the user with enough plantar sensory feedback to quickly inform them of the general surface contours and potential hazards underfoot while walking only. The system is not designed to detect steps while jogging or running.

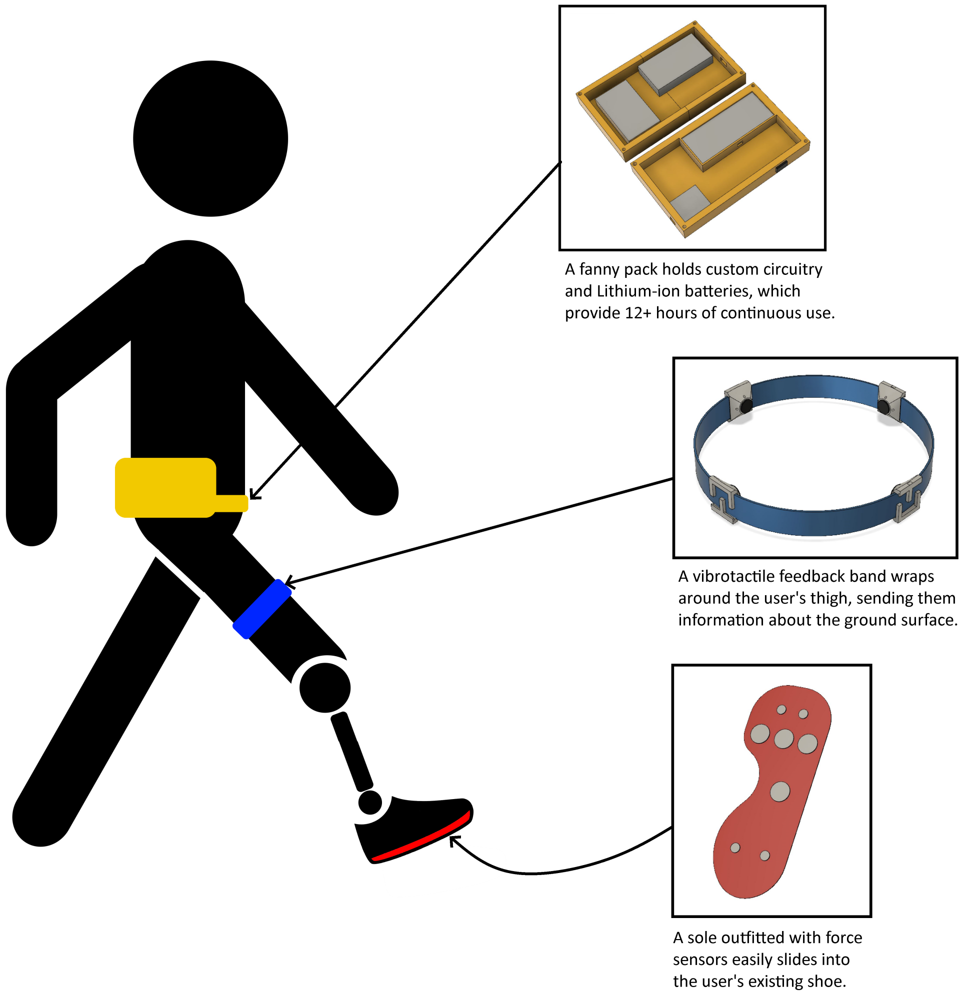

Figure 2: *The TerraStep's sub-systems; seen on a stick figure to better understand how the entire system comes together. Made using Paint.net.*

3 System-Level Overview

The three major TerraStep Sub-systems are the haptic feedback motors around the mid-thigh, the sensors underfoot, and the fanny-pack containing the microcontroller and additional electronics. We have illustrated the relative locations of each sub-system in Figure 2. The haptic feedback band provides vibrotactile feedback based on the forces underfoot, measured by the sensors. The microcontroller parses the sensor data and activates the feedback band accordingly in real time.

Two styles of vibrotactile feedback accomplish our overall goal of alerting the user of hazardous steps (see Detection Algorithm section for details). Normal steps will result in low intensity continuous feedback, allowing the user to develop an intuitive sense for the feedback over time. When a tripping hazard is is detected, a higher intensity vibrotactile pulse is used to quickly alert the user of the object and to convey its relative location underfoot. These pulse alerts are needed to notify the user immediately, allowing time to react. These two feedback styles will be referred to throughout this document as "continuous feedback", and "alert feedback", respectively.

The force sensors are incorporated as an array into a shoe sole insert. This array is easily

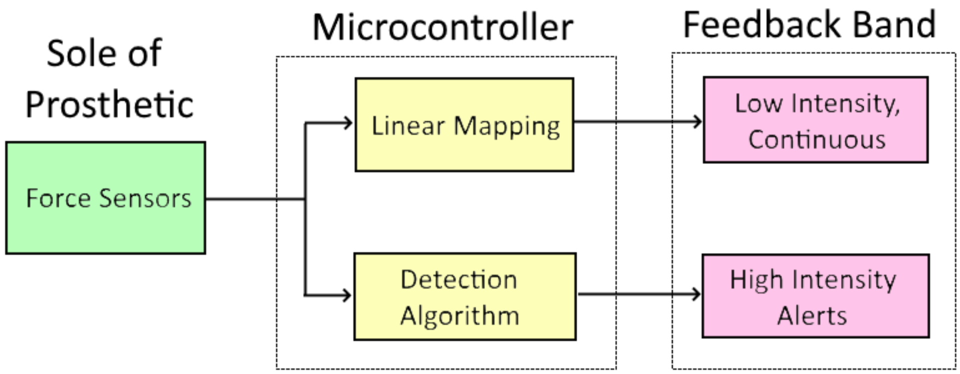

Figure 3: *High level block diagram of the TerraStep operation. Forces are measured by the sensors and parsed by the microcontroller. The algorithm classifies the step as a normal step or a hazardous step, and activates continuous or alert feed back, respectively. Made using Paint.net.*

slipped into the shoe of the prosthetic limb.

To detect hazardous steps, the sensors measure forces underfoot that the microcontroller compares to a normal step (see Figure 4). A "normal" step is defined as a step on a flat, hard surface, while a "hazardous" step is defined as a step on a tripping hazard. We have defined a "tripping hazard" as a hard wooden sphere approximately 1.25 inches in diameter with a deviation of +/- 0.25 inches. When a user has a hazardous step, there is a measurable difference in the force between adjacent sensors in that region underfoot, which would otherwise provide signals of very similar amplitude during a normal step. (See Detection Algorithm section for details).

During a normal step or normal walking, the microcontroller periodically drives the haptic feedback motors on the user's thigh according to which regions underfoot are in contact with the floor. This method of "continuous feedback" develops user familiarity with the new system. After wearing the TerraStep system for a period of time, the user will learn to associate pressure on the front of their prosthetic foot with the vibrotactile feedback on the quadriceps region of the thigh. Conversely, they will associate pressure on the heel with feedback applied on the hamstring region of the thigh. The system will provide feedback unless the user is stationary. Our prosthetic user testing subject Jim reported that it got more intuitive after four hours of use, and that it likely would continue to do so with more time.

During an abnormal step, the system will provide a high-intensity alert feedback to notify the user that they have just stepped on an obstacle. The sensors measure a large difference in adjacent sensors and the microcontroller determines what region the difference is in. High-intensity feedback is provided on the thigh corresponding to the region where the obstacle was detected. The flow of information is presented in Figure 3, where the distinct feedback options are under the same sub-system.

The circuit boards and batteries are contained in the fanny-pack, to be worn on the user's waist. See Appendix for PCB engineering schematics. The batteries have the capacity to run the system at full functionality for at least 12 hours.

The TerraStep needs a method of quantifying its success at its goal, which is primarily to alert the user of hazardous steps. However, the subjective tactile feedback sensation and broad variation of user gait makes quantifying effectiveness difficult. We've split up effectiveness into two areas of focus, which together provide a metric of success. These are, (1) correctly distinguishing the object's location underfoot, and (2) the delay of the system from object contact to alert activation. If the system can alert the user of a hazardous object's location within 40ms, it will give the user enough time to react and prevent injury.

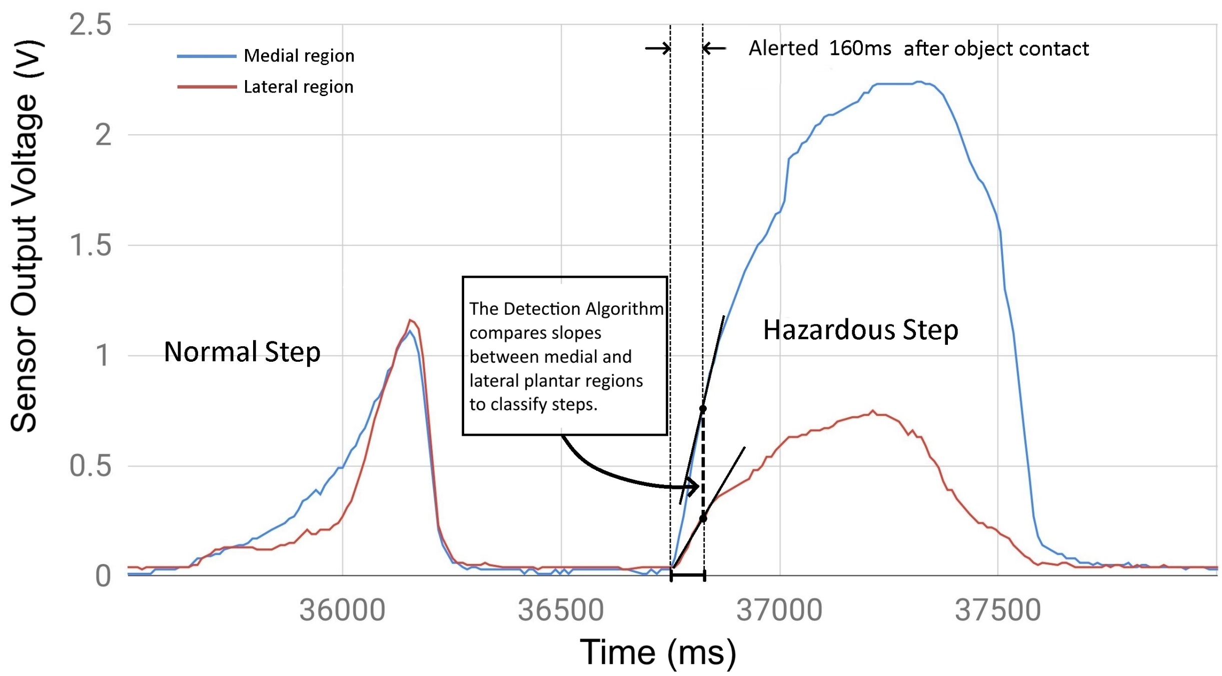

Figure 4: *Shows a typical normal step and hazardous step with the object located in the anterior medial region, and the typical response time from object contact to alert activation. Made using Microsoft Excel and Paint.net.*

Mechanically, compactness and universality are two important constraints of the TerraStep. Considering the intended application of helping active prosthetic users, the system must be comfortable and easy to move around with. An amputee's prosthetic is designed to be lightweight, and the added weight or displacement of our system should not be significantly noticeable to the user. The sensors are attached to the bottom of the foot, the actuators on the lower thigh, and the circuit boards and batteries contained in a fanny-pack. Sensors and actuators are connected to the fanny-pack via a wire bus.

4 Haptic Feedback

Prosthetic users are disconnected from the orienting sensory information that a biological foot provides. The feedback component of the TerraStep system closes that disconnection by communicating to the user details about the surface underfoot.

Two forms of haptic feedback are needed to effectively help the user during a hazardous step, continuous feedback and alert feedback. If only relying on continuous feedback to detect objects, the user will be relying on two potential signals. The first is the increased plantar force in the region contacting the object. This greater force would produce a greater feedback intensity, however it can be difficult to distinguish such a subtle change in intensity especially while walking. Second, the user could notice a difference in their normal feedback pattern, but this would occur far too late in the step.

The continuous feedback not only allows the user to grow familiar with the vibrating feedback, but increases overall sensory communication between prosthetic and user. The potential benefits include shorter learning period for new prosthetic users, and greater attunement with the prosthetic leading to overall improved efficacy of use.

It uses small vibrotactile motors (i.e. buzzing motors). Such tactile communication is categorized under the umbrella term, haptic, which refers to all forms of touch sensation.

The TerraStep's haptic actuators must exert a great enough intensity to alert the user while walking, have an intensity resolution allowing at least three distinct levels (not including zero), have consistent intensity, a delay less than 40ms, be durable, and consume reasonable power. Ideally (but less critically) they will also be inexpensive, and easy to implement.

While there are a few different existing haptic technologies (including electrostimulation, and pneumatic), the most common by far is vibrotactile. This is the same technology used in our cell phones and many other modern devices. The other technologies were ultimately decided against because of their high cost to perform tests (electrostimulation and its FDA approval fee) or their lack of availability and existing knowledge-base (bladders). On the other hand, vibrotactile motors are ubiquitous and inexpensive. The Adafruit 1201 seen in Figure 5 was chosen as our final haptic actuator primarily due to it's greater maximum intensity than comparable actuators. This naturally correlates to greater power consumption. However maximum intensity is a much more critical criteria regarding the functionality of the overall feedback mechanism, so it is an easy trade off. Each 1201 motor draws no more than 120mA at full intensity, a reasonable amount. A secondary unique advantage of the 1201 is it's ease of implementation. It is driven by a simple DC signal. This allows a minimized driver circuit, as well as simplified microcontroller output compared to other vibrotactile motors' input requirements.

The choices of the 1201 and the insert seen in Figure 6 are supported by personally conducted laboratory experiments. The first experiment described here quantifies motor intensity, comparing different motors and configurations. It was developed with the aim of answering three primary questions concerning the haptic motors:

- Are they strong enough to be felt while walking?
- Do they provide enough intensity resolution to be perceived at three distinct levels?
- What will the input and output voltage-to-intensity mapping data be? (needed for the microcontroller subteam)

This experiment is set up by attaching a single motor on a band to the subject's mid-thigh. The band is a standard backpacking strap with the motor held in an insert seen in Figure 7. The

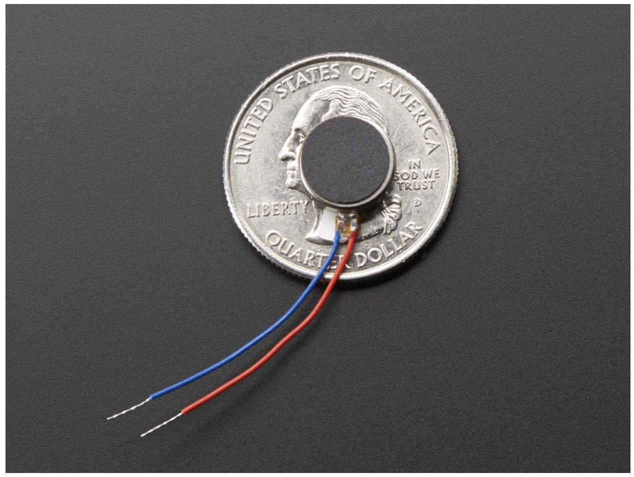

Figure 5: *Adafruit 1201 vibrotactile motor is strong enough to feel while walking, but also very compact.*

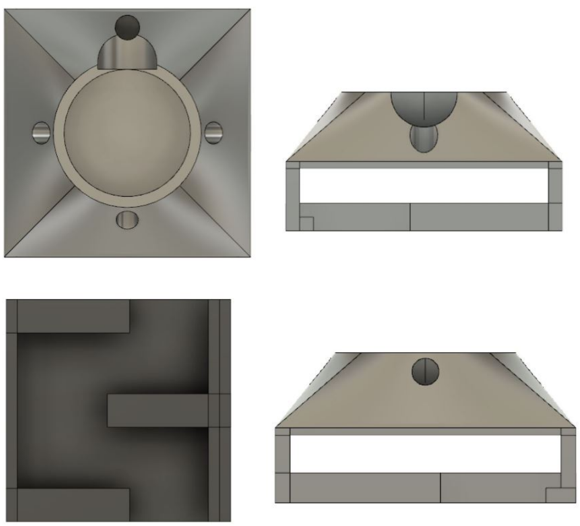

(a) *Four views of our actuator insert design. Securely holds the actuator, allows movement along the band, and protects the fragile motor leads. CAD done in Fusion360.*

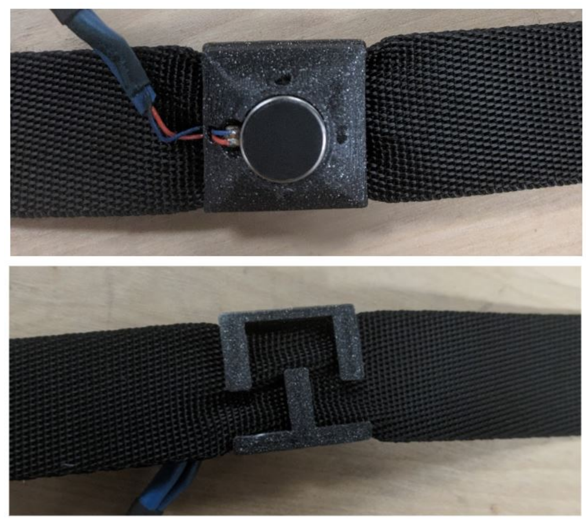

(b) *3D printed actuator insert. Notice how the design allows the insert to be slid along the band. Motor leads are better protected by being glued into place. Made using a Prusa i3 MK3.*

Figure 6: *Shows the actuator inserts we used to increase the range of vibration felt by the user. Our CAD design is shown in (a) while the physical prototype, with actuator, is shown in (b).*

motors have a few yards of cables separating the subject from the experimenter to allow freedom of movement while walking. The experimenter has a normally-open switch that activates the motors when pressed. See Figure 8.

1. The experimenter writes out a random voltage order unknown to subject, i.e. "blind". These are voltages within the operating range of the motor, 0-5V.

2. The subject takes several steps (> 3) at one voltage, while the experimenter activates the motor each time the subject's foot hits the floor. This simulates the motor's operation within the full system.

Figure 7: *Teammate Justin Fortner is shown wearing the current design with the 1201 motors attached to his leg for testing.*

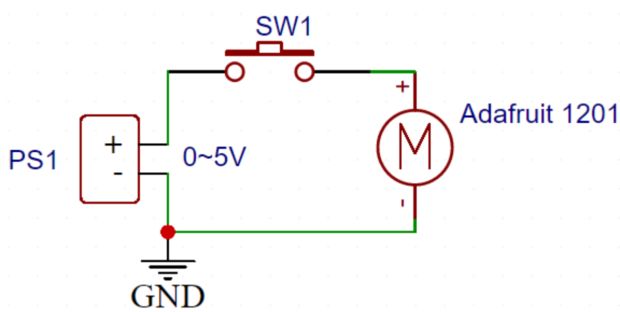

Figure 8: *Schematic for experiment setup. The experimenter varies voltage source. Made using EasyEDA.*

3. The subject then verbally quantifies the perceived intensity of the random voltage on a scale of 0-10, ”1” being the feeling just above threshold, and ”10” being the motor driven by its maximum 5V. The “10” reference intensity is applied to the subject at the very beginning of the experiment.

4. The next random voltage is set by the experimenter, and the subject begins walking again, repeating steps 2-4 until all desired voltages are tested.

The Adafruit 1201 motors were proven to be the only actuators with a strong enough intensity to be felt while walking. This critical criteria made them the clear first choice. They also have a decent resolution intensity, with results suggesting that there are approximately three distinct levels of intensity perceived by subjects. This should be sufficient for our needs, which are not to reproduce a foot's entire sensation, but rather to alert the user of potential unseen tripping hazards underfoot. Thus the lower intensity level is employed for normal steps, and the highest for alert status. However, users' tactile sensitivity varies widely, therefore some adjustment is required. For example, one test subject did not feel the low intensity continuous feedback at all, while another felt it easily. For future work, this adjustment could be implemented into a knob or calibration procedure.

Here we should comment on the perceived intensity resolution (see Figure 9). It is unclear whether the three levels described above are particular to these motors (no other motors had a comparable intensity range), or whether it is a matter of inherent perceivable intensity resolution of human beings. It seems probable that within the intensity range of these motors, each division of the 1-10 scale was so small that it is all but indistinguishable. It may have been more effective to have a scale of 1-5, or even 1-3. Since we only really need two or three levels of distinguishable intensity, that would be an appropriate scale for this test.

In addition to proving the 1201 as the best motor choice, the experiment results also proved that the custom insert significantly increased the maximum intensity level, by an average 42 percent (see Figure 9), thus increasing the intensity range as desired.

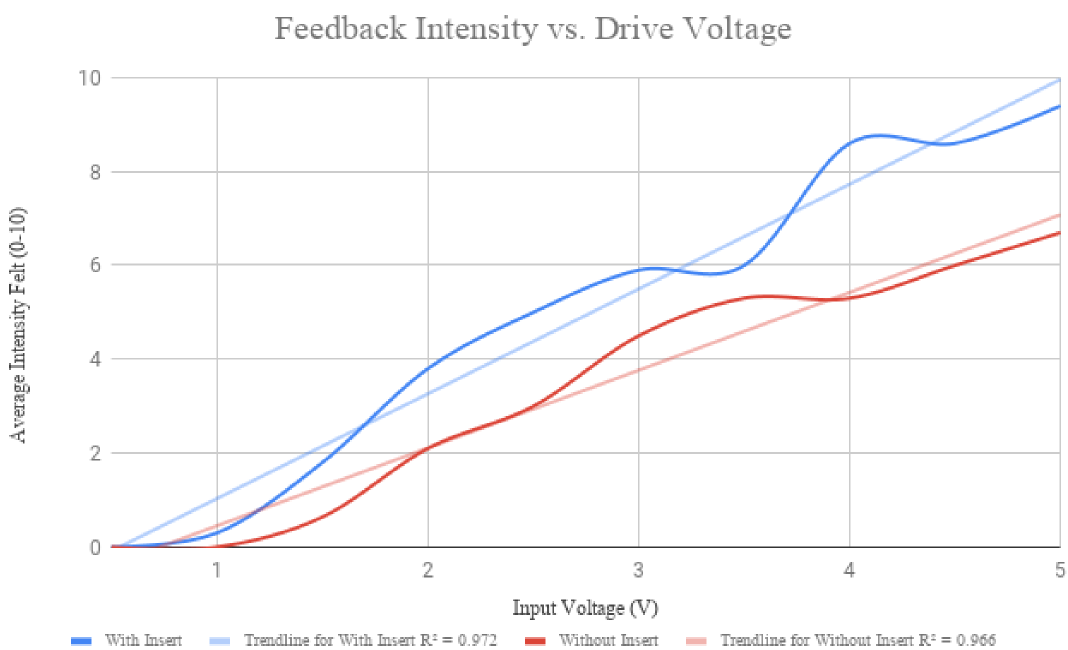

Figure 9: *Intensity with insert is approximately 42% greater than without. Made using Microsoft Excel.*

The feedback band contains five vibrotactile motors positioned around the circumference of the user's thigh. These motors need to be individually distinguishable, therefore we need a minimum effective distance between them. Tactile sensitivity varies at different points around the thigh, and the perceived motor intensities are subjective from user to user. If the motors are in an insensitive region or are too close together on the thigh, we are concerned that the user won't be able to distinguish the separate motors effectively. This experiment will clarify the minimum distance between two Adafruit 1201 motors, while considering the different regions of varying sensitivity along the users thigh. For example, a result might be "motors must be spaced 2 inches apart on outer thigh region". These results allow us to optimize the motor positioning on the band.

1. The mobile TerraStep prototype including the band with X motors is attached to the subjects leg. Each motor is held in an insert along the inside of the band.

2. All but two motors are disabled, to narrow test focus. Place motors X inches apart on thigh region of interest.

3. Experimenter randomly activates one motor, or both, while the subject is walking at a normal pace. The subject states which motor(s) they perceived after each activation, which is recorded by experimenter.

4. Repeat from step 2, changing the distance between motors.

After the test, the subjects reported data is compared with the actual sequence of motor activation. The minimum distance is based on the accuracy of the subject's perceived intensities. Minimum distance is defined as that when the subject's accuracy falls below 75%.

Each motor requires a drive circuit. This allows the microcontroller's PWM output to individually vary the motor intensities, while their current is sourced from a separate power supply[1]. The microcontroller used (Teensy 3.5) cannot supply adequate current to drive the 1201 motors, therefore we developed a buffer and FET configuration that allows the current to be drawn directly from the battery (see Figure 10).

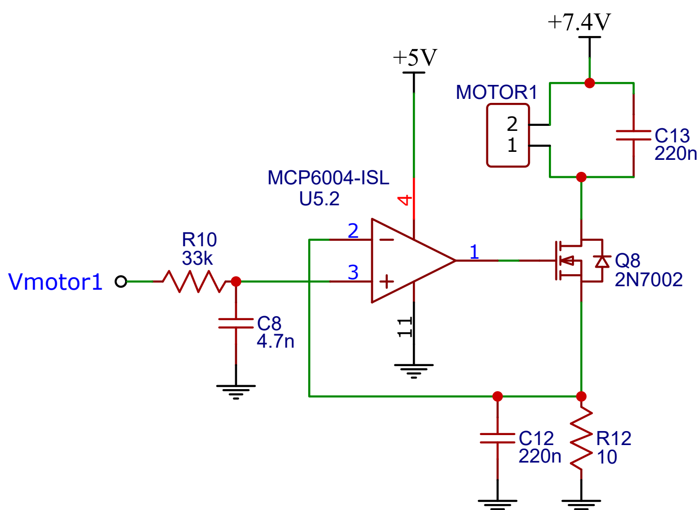

Figure 10: *Driver circuit for Adafruit 1201 motor. Made using EasyEDA.*

The first stage of the circuit is a low pass filter to change the PWM input to dc. The opamp acts as a buffer to the MOSFET, with negative feedback from the drain resistor. This ultimately makes the circuit transistor-independent, an important feature considering the multiple motors involved system-wide. The FET essentially controls motor current by operating as a variable resistor. It is an open circuit when the opamp input is below $150mV$, and fully closed above $1.3V$ (or a $3.3V$ PWM signal with a duty cycle of 4 to 40 percent, respectively). The shunt capacitor is for snubbing purposes, required at shut off due to the inductance of the motor.

[1] Adafruit has a driver available specifically for these 1201 motors, DRV2605L. However, after testing we determined that they significantly reduced the overall intensity and were insufficient for our application.

4.1 Haptic Feedback Mechanical

The feedback band must be comfortable and durable enough to wear during a long period of time over the course of several months and is hold the vibrating motors tightly to the users leg without any slippage during athletic movements, achieved through the band seen in Figure 11. All of the requirements are met through the use of a UV and water resistant, non slip nylon rated to hold a dynamic force of 175 pounds.

The 1201 vibrating motors provided the greatest amount of vibrational force out of the motors we rested. These motors still did not meet the required amount of vibrational force needed to effectively transfer feedback to the users thigh. Thus we needed to increase this force and attach these motors to the band. This feat was accomplished through the use of custom 3D printed housings, as seen in Figure 12 and Figure 13. The force behind the band transfers to the housing and the motor to force the motor into the users skin. This increase in force allows for the maximum transfer of vibrations from the motor to the users leg. These 3D printed housings use set crews to hold one motor in the housing and a three prong system to secure the housing to the band. The set screws do not dampen the vibrations coming from the motor and allow for a longer lasting hold than a traditional adhesive would. The three prong system allows for the motors to be secured in proximity to the users leg, but can be slid around the circumference of the thigh to adjust to the users preference. This was achieved by using a 3D printer capable of printing 20% carbon fibre fill filament at an accuracy of .02mm. The housings take roughly 25 minutes to produce.

Figure 11: *Full System Band And Actuator Layout. CAD done in Fusion360.*

5 Sensors

5.1 Force Sensor Considerations

Quite possibly the most fundamental component of the TerraStep system, the force sensors act to reconstitute the user's underfoot sensation in a new way. The ground surface applies forces directly to the sensors, which are then converted to electrical signals. It is with these raw signals that

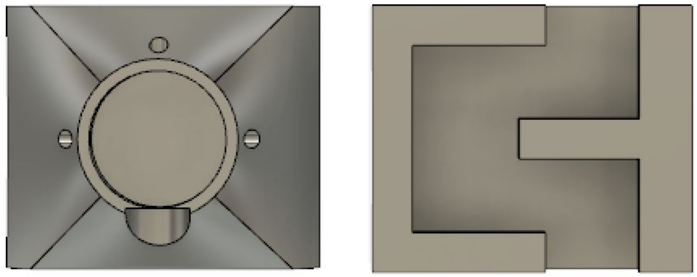

Figure 12: *Front and Back View of Actuator Inserts. CAD done in Fusion360.*

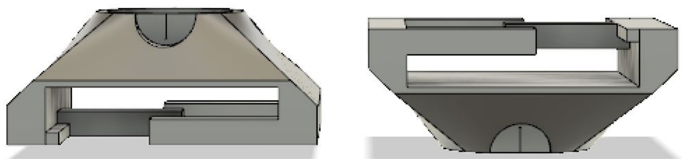

Figure 13: *Side View of Actuator Inserts. CAD done in Fusion360.*

the microcontroller activates corresponding vibrotactile feedback, and can quickly alert the user of tripping hazards. Ultimately, this improves balance and safety.

The primary criteria for a TerraStep force sensor are its thickness, and durability. A thick sensor will affect the overall length of the prosthetic leg, and consequently the user's gait. Compensation for the additional length would be possible, but is undesirable. The sensor must be less than 0.6 inches thick so that the user is unware of the added height[2]. The sensor must be able to remain functional for at least six months of average use, which can be approximated by 10,000 steps a day or 900,000 total steps. Six months between sensor replacement was deemed reasonable for the user.

Tekscan A301 and A401 piezo-resistive sensors were chosen as the best sensors for the TerraStep for their thin profile and one-million-step lifetime.

Drift is a critical parameter of the sensors. An excessively drifting sensor can either decrease feedback intensity, or make all feedback high intensity (no gradient). It will also affect the detection algorithm, giving false alerts, or no alerts at all. The sensor voltages must not drift by more than +/- 5% of their maximum force within 12 hours (determined by looking at the graph in Figure 9, where, assuming linear sensor-to-motor mapping, 5% drift translates to 0.5 drive on the intensity scale, y-axis). The datasheet does state this +/- 5% tolerance (per logarithmic time scale), therefore the sensor is within the expected tolerance.

Regarding linearity, the Tekscan sensors are manufactured to have a linearly proportional response to force (within +/-3% of maximum force within chosen range). The force range is specified by the drive circuit. A body weight of 160 lbs means each sensor will be receiving approximately 20 lbs. 3% translates to a +/- 0.6 lbs error in linearity, entirely acceptable, and unnoticeable at the feedback intensity.

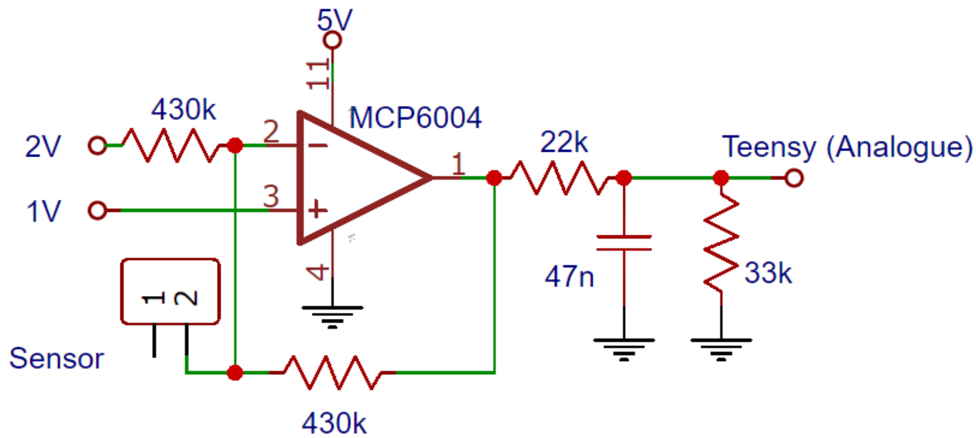

Figure 14: *Shows the inverting-summing gain circuit used to convert the sensor's resistance to a voltage between* $0V - 3.3V$*. The MCP6004 IC was chosen due to it's recommendation from Tekscan, ubiquity, and* $0V$*-*$5V$ *range. Made using EasyEDA.*

5.2 Sensor Circuit Design

The full sensor circuit is constrained by the following system level requirements: low time delay and long battery life. The low time delay is addressed by using higher corner frequencies while the battery life is addressed by using larger resistors. The circuit between the sensors and micro-controller provides a reference voltage to drive the sensors, utilize the full $0V$ - $3.3V$ range of the micro-controller for maximum resolution, and keep any time delay imperceptible to a user.

We opted for an inverting gain circuit because it can produce a gain lower than one, meaning the signal can effectively go to 0V. We superimposed a summing circuit on the inverting gain circuit to remove the necessity of a dual-ended power supply that would have been needed to produce a reverse voltage across the sensor, as recommended by Tekscan. This is best illustrated by Figure 34 in Appendix: Section 13.1. The voltage from terminal 1 and terminal 2 of Sensor_1 should be negative. That is to say, terminal 2 should be at a lower voltage potential than terminal 1. This reverse voltage is inverted by the gain stage and produces a positive voltage at the output of the op-amp. The circuit in Figure 14 provides an output voltage proportional to the resistance of the sensor while subtracting the $1V$ bias introduced by the reverse voltage needed to run the sensor. Without removing the $1V$ bias, the op-amp would output a signal with a range from $1V$ - $5V$ and would have removed 20% of our total possible signal resolution. Due to the $3.3V$ input range of the micro-controller, we added a voltage divider at the end to down-step the voltage range from $0V$ - $5V$ to $0V$ - $3.3V$, eliminating the possibility of overloading any micro-controller pins. A reference voltage circuit was designed to provide the $1V$ and $2V$ reference rails in Figure 14 while a separate circuit was designed to provide the reverse voltage to the sensor.

In order to keep any time delays imperceptible to the user, we designed the filter in Figure 14 with a corner frequency of approximately $250Hz$. This converts to $4ms$, approximately two orders of magnitude smaller than the average human reaction time of $250ms$.

The reference circuit was designed while considering the system-level constraint of a 12+ battery life. The circuit in Figure 15 produces a $1V$ and $2V$ stable rail that does not draw more than $200\mu V$ per rail, excluding the current draw of any op-amps used. The circuit utilizes one voltage divider connected to a $5V$ supply to produce the $2V$ rail. The $1V$ voltage divider is placed at the output

[2] https://orthoinfo.aaos.org/en/diseases–conditions/limb-length-discrepancy/

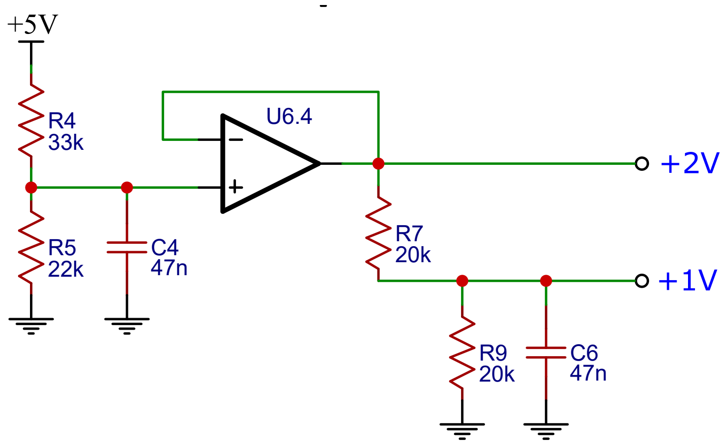

Figure 15: *Circuit designed to produce voltage rails with less than $200\mu V$ per voltage divider. Provides voltage rails to inverting-summing circuit in Figure 14. Reference Appendix: Section 13.1 for full schematic. Made using EasyEDA.*

of the $2V$ rail to allow for an easier supply voltage change, as only one voltage divider would have to be redesigned. Only one op-amp was needed for this circuit, as the $1V$ rail is only connected to the positive terminal of the op-amps, which draw no current.

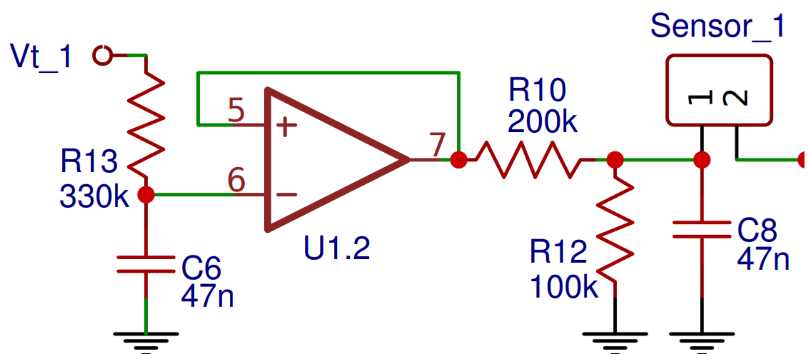

Figure 16: *Shows the Sensor Drive-Voltage schematic, providing a negative voltage in Figure 14 using an integrator circuit with a $f_c = 15mHz$. Made using EasyEDA.*

The circuit in Figure 16 was designed to supply a reverse voltage, biased at $1V$, to the sensor that controls the effective range of the circuit in Figure 14. The reverse voltage directly proportional to the duty cycle of a PWM signal from the micro-controller. This is accomplished with an integrator circuit that has a corner frequency of $15mHz$. The speed of the circuit is irrelevant, as it will only be used for an initial calibration and will not need to update in real time. As long as the DC voltage does not surpass $1V$, it will produce a reverse voltage across the sensor when referenced to the $1V$ rail supplied by the circuit in Figure 15.

5.3 Sensor Mechanical Implementation

The more data collected by the sensors, the more accurately we can analyze a users step and relay the correct feedback. In order to reach the maximum amount of data to be collected by the sensors, all of the sensing area must be used. The most ideal way we found to achieve this is through the

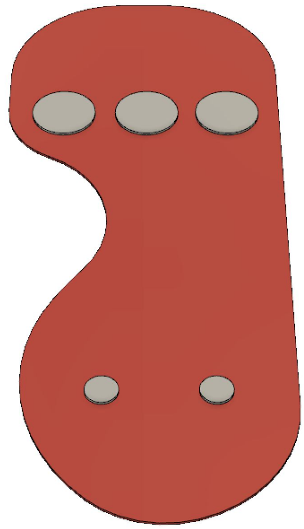

Figure 17: *5 Sensor Array Layout. Allows for detection of forces underfoot in five specific locations denoted by the grey circles. CAD done in Fusion360.*

use of sensor pucks. The sensor picks need to be the exact diameter of the sensing area. If smaller not enough of the sensing area is used and if bigger the pucks experience a rocking motion that causes force to be placed outside of the sensing area. The thickness of the pucks needs to strike a balance between thickness, rigidity and strength. The thickness of the sensing pads needs to remain as this as possible to allow for the use of other components in the sole insert and keeping the overall thickness under 7mm. After this threshold user will begin to experience hip issues due to the discrepancy in leg height. The pucks need to remain rigid when stepped on in order to maintain consistent and accurate transfer of force to the sensors,m regardless of foot pressure or placement. If the puck begins to flex under the pressure exerted on it the readings will be inconsistent and the surface area used by the sensor will decrease as well. Leading to a reduction in quality of the feedback. Lastly, the pucks need to be strong. The pucks need to handle thousands of steps per day for month on end. If the pucks lack strength to remain rigid over this extended period of time they will become brittle and break. Thus reducing the longevity of our system. We found that we were able to exactly match the sensing area and achieve balance of strength and rigidity at just a 1mm thickness. This was achieved by using a 3D printer capable of printing 20% carbon fibre fill filament at an accuracy of .02mm. The pucks take roughly 10-15 minutes to produce depending on the diameter required by the sensor.

5.4 Sensor Array Design

The sensor array is comprised of three different layers serving three distinct purpose. The second layer, or center, of the array is the vacuum seal plastic enclosed sensor array shown in Figure **??**. The vacuum seal plastic was used in order to make sure the force sensors were fixed in position, and to allow for a more durable sensing array. The plastic is designed to be heat sealed using a press. We utilized this method to fix each sensor in place, preventing unwanted sensor movement. In addition, the vacuum seal plastic provides a water resistant enclosure for the sensors. However, the current implementation is limited in its ability to prevent standing water from entering the housing due to unsealed openings for wires.

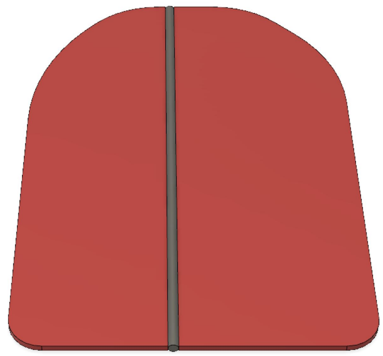

Figure 18: *Sensor Array Foot Plate. CAD done in Fusion360.*

Enclosing the array is a layer comprised of a thin and soft foam insole and a layer with two hard plastic discs hinged at the center shown in Figure 18. The soft insole is what comes in contact directly with the users prosthetic foot while the hinged discs sit between the sensors and the inside of the users shoe. Each disc is shaped to match the contour representing each of the two toe region sub-quadrants. The purpose of the plastics discs is to extend the sensing area of each sub-quadrant of the toe region, while allowing them to remain isolated. The hard plastic allows for pressures exerted off center of the pressure sensors sensing area to be transferred more towards the center. Without the plastic disks, the detection algorithm would only be able to sense a hazardous step if the obstacle encountered made direct contact with a pressure sensor.

On the other side of the array the soft insole is sandwiched, which helps to keep the sensor array in the correct position relative to the plantar region of the foot. This is necessary as to provide consistent hazardous step detection accuracy during extend use of the TerraStep system. If the insole is not present, the more pliable plastic of the sensor array layer tends to move around and contact different parts of the plantar region of the foot. In addition to holding the sensor array layer in place, the soft insole acts as a sort of "wedge" helping to keep plantar pressure transfer to sensors consistent.

6 Micro-controller Considerations

We chose the Teensy 3.5, from PJRC, to fulfill the following system-level and circuit-level features: eight distinct sensor channels, seven distinct motor channels, battery monitoring and data logging, easy PCB integration, and a low-cost, easy-to-use package. These high-level requirements break down into more specific micro-controller constraints.

The Teensy 3.5 has enough peripheral pins to run our hardware without the need for external circuitry. Each sensor requires one PWM pin to produce the reverse voltage on the sensor and one ADC pin to read the sensor circuit output. Each motor requires one PWM pin to drive the motor. The battery monitoring function requires one ADC pin to read the battery voltage. These I/O specifications require at least 15 PWM peripheral pins and 9 ADC pins. The Teensy 3.5 boasts a maximum of 20 PWM output pins and 27 ADC input pins. These specs fulfill our peripheral needs while allowing for future expansion of sensor and motor channels.

In order to facilitate the data logging feature of our system, the micro-controller would need to have either internal or external memory to save data files. The Teensy 3.5 has a micro SD port that will accept SD cards of varying storage sizes to fit our needs. Initial use of the SD card allows the team to parse the data easily by transferring files to a computer easily. The Teensy 3.5 software also has functions that can streamline the data logging function of the system, requiring less time spent coding functions for this feature.

The Teensy 3.5 is meant to be used on a breadboard for prototyping, making it's integration onto the PCB relatively easy. Once the layout was determined, the Teensy 3.5 is attached using female header pins that allow the board to be removed easily, should a replacement be necessary. The board is 62.3mm x 18.0mm, making it a compact solution.

The final reason for choosing the Teensy 3.5 is its relatively low cost and ease of use. The Teensy 3.5 development board is $24.25 from PJRC. The company also provides many references and guides for using the Teensy 3.5, allowing for less time on learning the system and more time coding. While a stand-alone micro-controller unit may be less expensive with a smaller PCB footprint, the development board, which allows for easy prototyping without external circuitry, and easy to use IDE make the Teensy 3.5 a great choice for our current applications.

7 Software Implementation

Our software implements four basic functions of the TerraStep: minimal delay, data logging, user acceptance and a hazardous-step detection algorithm that will alert the user. User acceptance and the detection algorithm inform different aspects of the feedback provided to the user. We use a default continuous feedback method, linearly mapping force underfoot to motor vibration intensity, to increase the user's familiarity and acceptance of the system, while checking step data to predict hazardous steps. If a hazardous step is predicted, the system will alert the user with high-intensity vibrations, overriding the continuous feedback. The delay between predicting a hazard and alerting the user must be small, as we are hoping to allow the user time to react and prevent potential injury. We also implemented data logging for research purposes, which will allow us to more easily test the efficacy of our system.

7.1 Data Logging

Data from each sensor is logged during operation of the TerraStep system. The total number of sensor which are logged can be configured to match the array type which is currently in use. During operation, samples between each sensor have a delay of about 2ms. When all sensors for a specific array have been sampled, then a sensing cycle has been completed. Each sensing cycle takes between 10-16ms depending on how many sensors should be logged. The array which data is logged from can be seen in Figure 17.

The raw voltage is the only data logged at this time. The voltage was chosen because it is the most basic form of the sensor response, and can be modified in matlab to obtain any other metrics: slope, slope difference, etc. Once each sensing cycle is complete, the voltage of each sensor sampled in the cycle is stored as entries in a row of a comma eliminated file (.csv). Each row is marked with a time stamp which is used to analyze the timing of the system.

7.2 Continuous Feedback

The default feedback provided, when hazardous step are undetected, is the continuous style. Continuous feedback is implemented for familiarity, allowing for the alert style feedback to be learned more rapidly. The feedback is controlled by the raw voltage response of the force sensors. The force sensors respond with a operating range of [0.0V, 3.3V] and is linearly mapped to the control range of the actuators, being [5%, 35%] duty cycle of a 1Khz PWM signal. Whenever a hazardous step is detected, existing continuous feedback is disable, and alert feedback takes over.

7.3 Detection Algorithm: Alerting the User

The detection algorithm uses the difference in slopes of adjacent sensors to determine whether an alert is necessary. The slope difference is computed as the difference between the anterior lateral positive slope and the anterior medial positive slope:

$$\text{Slope Difference} = \text{positive}(\text{Anterior Lateral Slope}) - \text{positive}(\text{Anterior Medial Slope}) \tag{1}$$

Figure 19 shows the positions of the anterior lateral and medial regions.

During a hazardous step, the slope difference can exhibit either a positive or negative response. The positive response occurs when a tripping hazard is under the anterior lateral quadrant, and a negative response occurs when one is under the anterior medial quadrant. This is because the

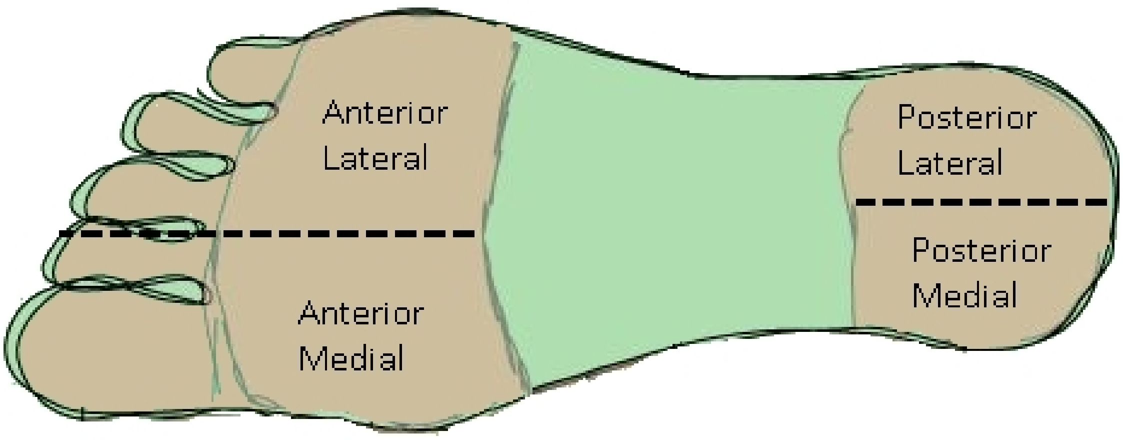

Figure 19: *Depicts the different regions of the foot we defined for easy classification. Our detection algorithm currently makes use of the difference in slopes between the anterior lateral and medial regions to detect whether a hazardous object underfoot. Made in Paint.net.*

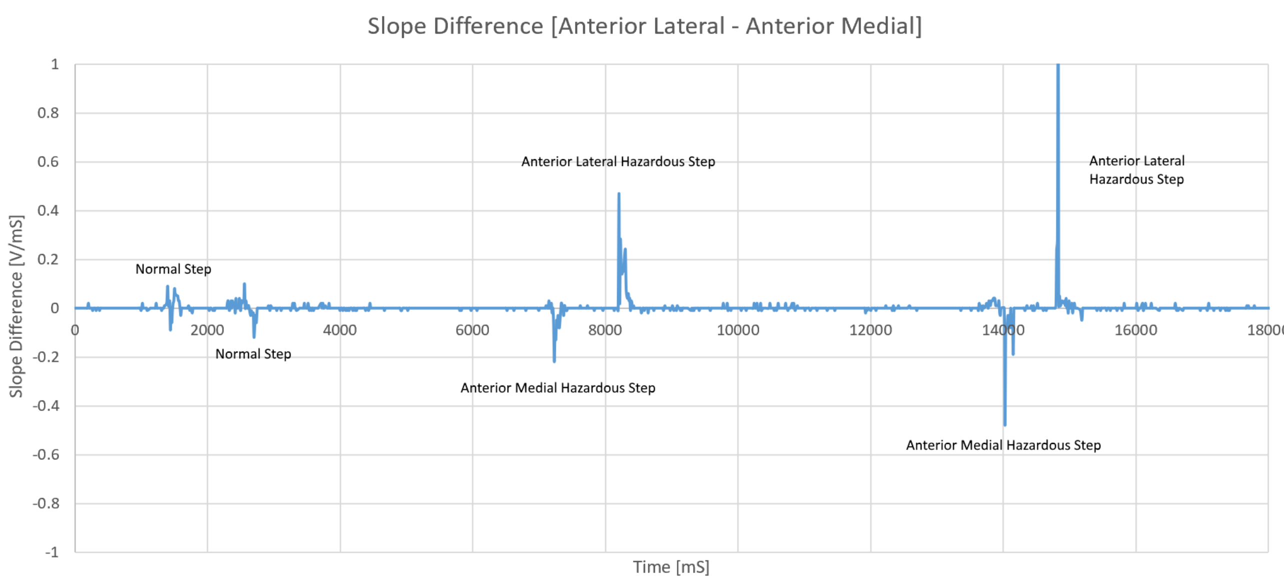

Figure 20: *The two steps farthest to the left are normal. Notice how these exhibit a positive and negative response of similar magnitude. The four addition steps are hazardous in nature. The slope difference is significantly (when compared to a normal step) negative during a anterior medial hazardous step. On the other hand it is significantly positive during a anterior lateral hazardous step. Made in Paint.net*

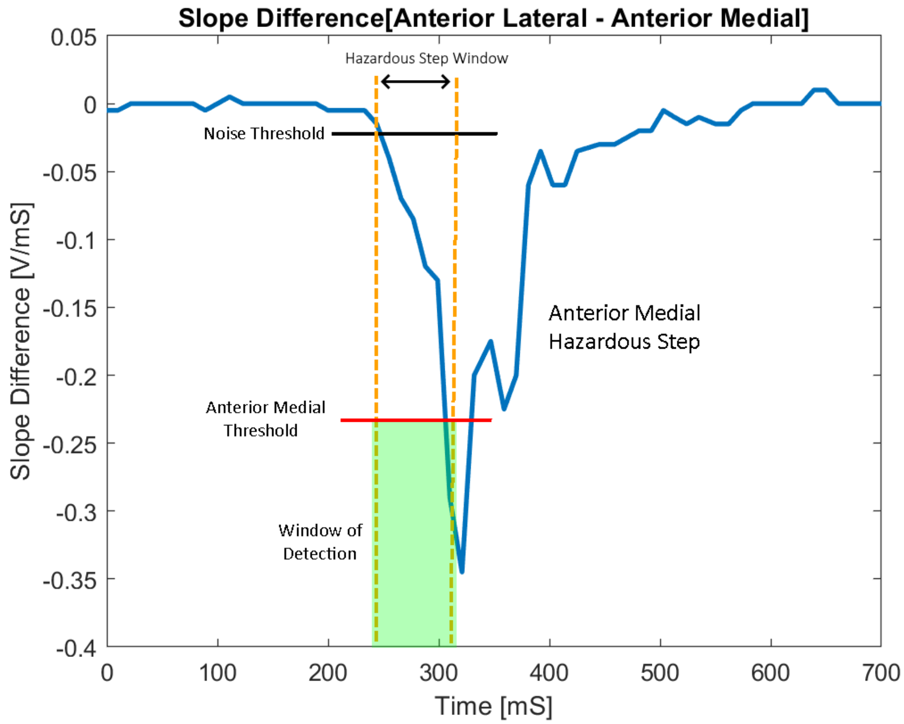

Figure 21: *The window of detection is the green region above. The Noise Threshold (black line) is used to generate the window of detection for each step. The Hazardous Step Window encompasses the interval of time in which a hazardous step can occur. The Hazardous Step Window is triggered by the Noise Threshold. Within the time allocated by the Hazardous Step Window, the Anterior Medial Threshold is used to determine the magnitude of slope difference which corresponds to a hazardous step. Made in Paint.net*

quadrant experiencing the tripping hazard produces a significantly larger positive slope response. This response will either make the right or left terms of Equation (1) much larger,facilitating a positive or negative slope difference depending on what type of hazardous step has occurred. This behavior is depicted in Figure 20 which shows the slope difference response during six various steps.

Ideally, the feedback would be nearly instantaneous to provide the user with the most time to regain their balance. However, this is a significant challenge and thus a maximum allowable delay was set at 40ms. This was was chosen as a reasonable maximum delay by analyzing the average amount of time it takes for a force sensor to achieve peak weight distribution during a sample of steps. This was found to be an average of 370ms over 15 sample steps as seen in Figure 23. Approximately one order of magnitude below the average was then used to arrive on a maximum allowable delay of 40 ms. With a delay at 40ms the system will provide the feedback in one tenth of the time it takes to go from initial contact to fully weight distribution. A delay of this magnitude still allows the user enough time to redistribute their weight in order to avoid fully stepping on a hazardous object.

7.4 Detection Algorithm Tuning Parameters

Four tuning parameters are used within the detection algorithm to calibrate the system to a specific user. These are described and visualized in Figure 21. Control over the sensitivity and accuracy of

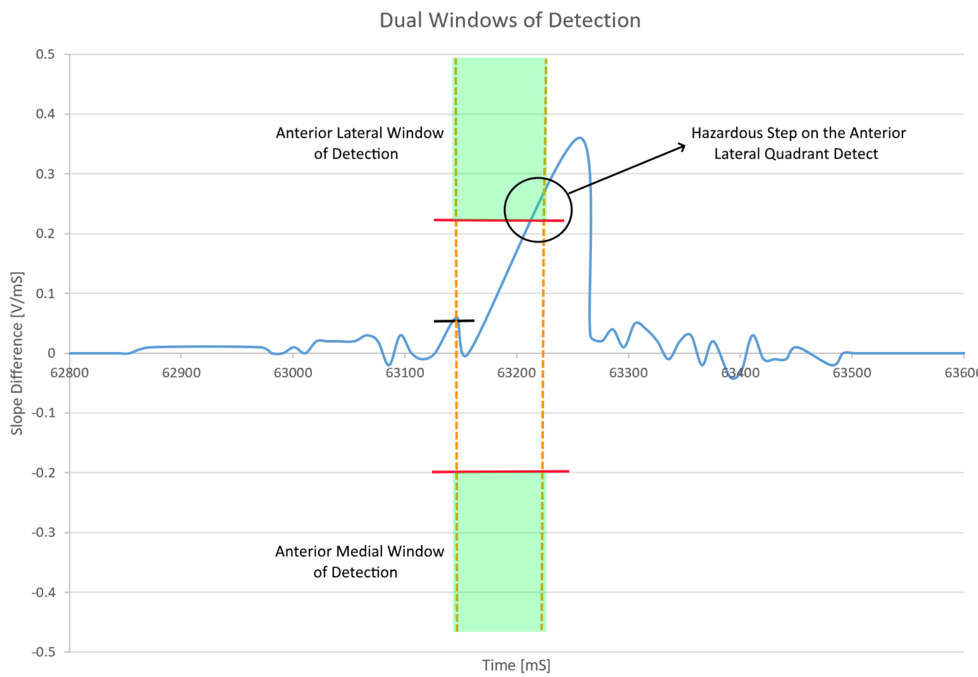

Figure 22: *Two distinct windows o detection are generate for each step. The positive window (green rectangle) depicted above is used to detect a hazardous step in the anterior lateral quadrant. The negative window is used to detect a hazardous step in the anterior medial quadrant. Made in Paint.net*

the detection algorithm can be achieved by changing any of the four tuning parameters. Changing each has a different effect to scale and shift the window of detection. Control over the location of the window of detection is important in making the TerraStep system effective for a specific user, and must be manually changed.

The Hazardous Step Window tuning parameter defines the period of time during a step in which a hazardous step can occur. The current implementation of the Hazardous Step Window is set to 60 ms. The original goal was to set this value to 40 ms which is the maximum allowable feedback delay constraint on the system. However, for an unknown reason, setting the Hazardous Step Window to this value was ineffective to allowing the algorithm to distinguish a normal from hazardous step. To solve this, the number was slowly increased until hazardous steps were able to be detected, while normal steps were ignored. This problem mostly like is caused by the say issue which causes the detection algorithm has an overall delay of 120ms when tested.

When the slope difference crosses the Noise Threshold (described below) the Hazardous Step Window is triggered. From the point in time in which the trigger occurs, the current step has 70 ms to achieve a slope difference response which crosses either of the hazardous step thresholds (described below). If neither of these thresholds are crossed within the allotted window, the detection algorithm makes no hazardous step classification, ignores the remainder of the step, and reverts to continuous feedback until another step has occurred.

The Noise Threshold parameter is utilized to filter out inherent system noise to prevent a false triggering of the Hazardous Step Window. This parameter works closely with the Hazardous Step Window to define the specific period of time during which a step can be determined to be hazardous. Each of the anterior sensing region's slope response signals are filtered using a five sample moving

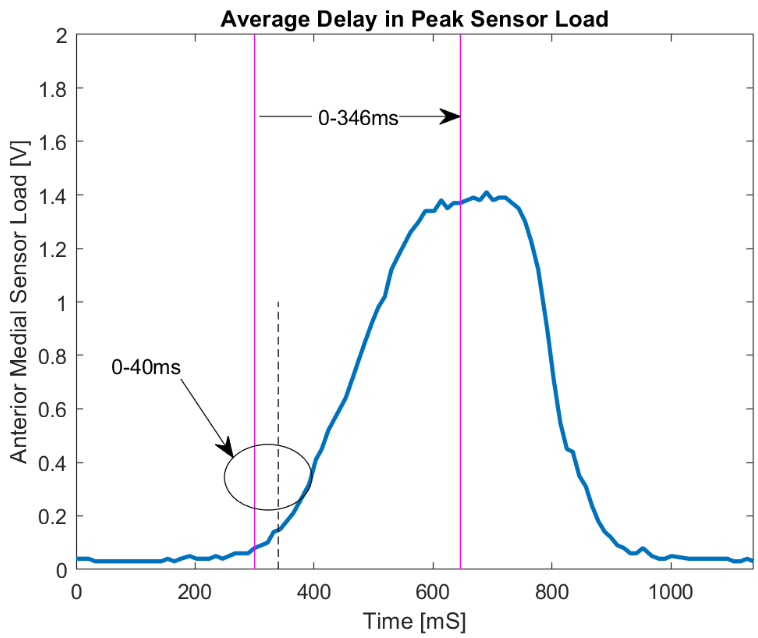

Figure 23: *The median response of the anterior medial sensor is depicted above. The magenta line farthest to the left defines the point of initial contact between the plantar region and the ground. Feedback at this point in time would have a delay of 0ms. The magenta line farthest to the right defines the average time it takes for the peak load on the sensor to occur. This average of 346ms was found from a sampled of 15 different normal steps. The dashed black line signifies a feedback delay of 40ms after initial contact. 40ms is the maximum allowable feedback delay required for the TerraStep system to be fully effective. Made in Paint.net*

average. Although this helps to reduce noise, this method is incapable of completely doing so without attenuating the sensitivity of the detection algorithm significantly. The combination of the five sample moving average and Noise Threshold tuning parameter solves this problem. Together, they filter out undesired noise while retaining a responsive slope difference signal. Adjusting the Noise Threshold to a lower value will result in more false positives while achieving a higher rate of successful hazardous step detection. In contrast, increasing the threshold works to eliminate false positives while reducing the rate of detection. An example of the noise encountered in the system can be seen in Figure 24.

Although this noise is small in comparison to the response seen during any step, it is still capable of a false trigger. To prevent this from happening, the Noise Threshold tuning parameter is used as an entry condition to the detection algorithm. If this threshold is never exceeded, then the core logic of the detection algorithm is never executed, and alert style feedback is not provided to the user. This prevents the algorithm from needlessly running during the period of time between steps.

The Anterior Medial and Anterior Lateral Slope Difference Thresholds are the final tuning parameters. These thresholds correspond to the magnitude of slope difference response which defines a hazardous step for that region. The toe sensing region is only implemented currently, therefore these two thresholds are used as triggers on the slope difference response of the anterior region.

The two distinct thresholds are used to create dual windows of detection, one for a positive response and another for a negative response. The separate windows are needed to capture both

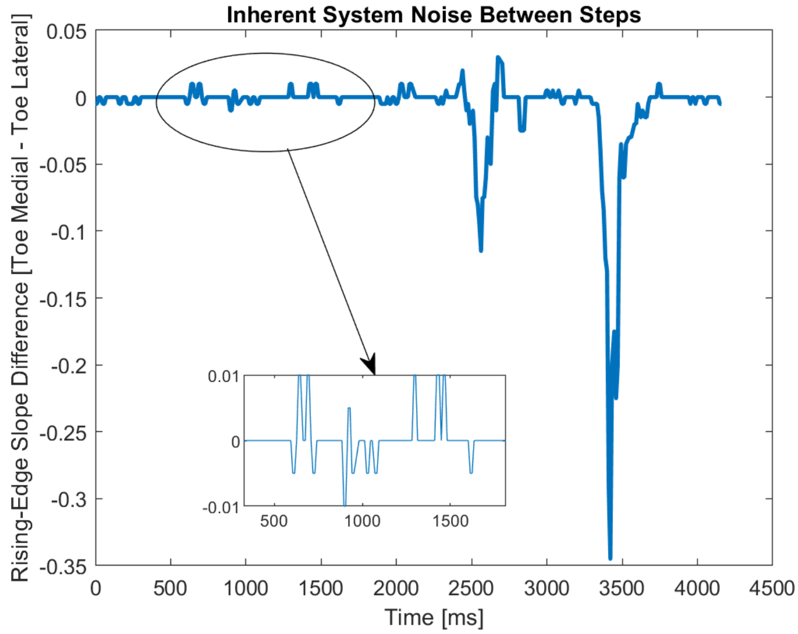

Figure 24: *The zoomed in graph above shows the magnitude of noise present in the slope difference signal. This noise is present when the user is in between steps, and there is no force is being applied to sensing array. The Noise Threshold is used to filter out this noise. Made in Paint.net*

extremes of the slope difference response, being positive during a hazardous step on the anterior medial quadrant, and negative during a hazardous step on the anterior lateral quadrant. The two windows are shown in Figure 22.

The tuning parameters can be independently tuned to adjust the accuracy of hazardous step detection for the anterior medial or anterior lateral quadrants, depending on the gait profile of the current user. For example, the system may be highly effective at detecting a hazardous step for the anterior lateral quadrant, but may struggle at detecting a hazardous step in the adjacent anterior medial quadrant. To solve this, just the Anterior Medial Threshold can be manually adjusted down towards a level which achieves the rate of successful hazardous step detection desired. At this time, there is no automatic calibration process for a user to run.

7.5 Confirming Slope Difference Efficacy

The slope difference was analyzed to confirm that it is the most effective metric at classifying any given step. The usefulness of the metric depends on how quick it is able to make a step classification. Also, the metric must have a significantly different hazardous step response when compared to a normal step. Data on a normal step, and both an anterior medial and posterior lateral hazardous step, were gathered using the three anterior sensors of the five array shown in Figure 17. The data in question is the rise-edge slope response of each sensor in the array with three force sensors. The slope difference between the anterior lateral and anterior medial quadrants were then computed. The average of this metric over 30 of each step type can be seen in Figure 25.

The average and standard deviation of the slope difference peak was computed for 30 normal step and 30 hazardous step on the anterior lateral quadrant. Only the positive slope difference peaks of the normal steps were averaged. This is to allow for a legitimate comparison between a normal step and a hazardous step on the anterior lateral quadrant which exhibits a positive

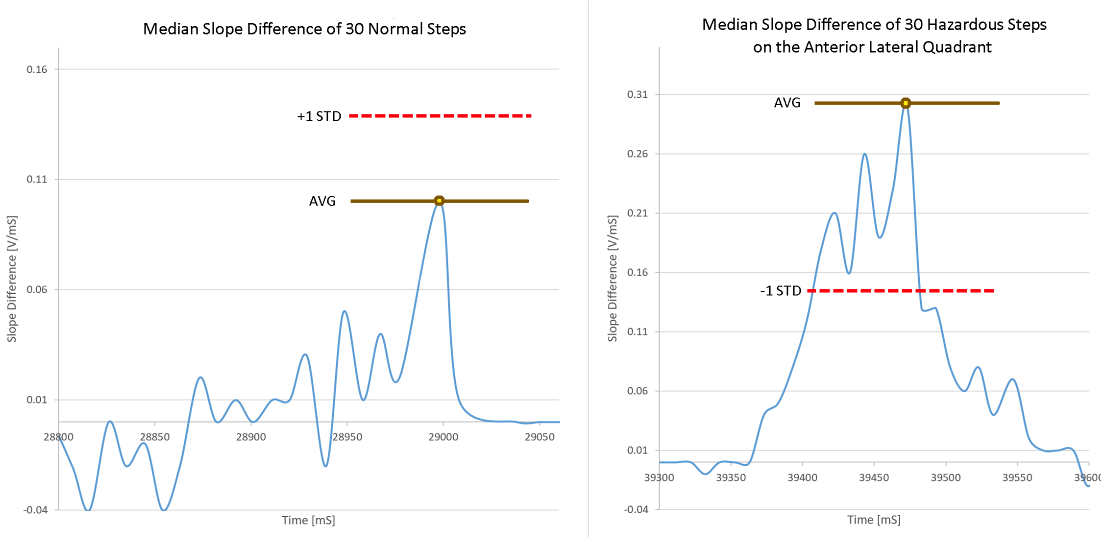

Figure 25: *The left response is that of a normal step. The right response is hazardous step on the anterior lateral quadrant. The graphs shown are of a single response of each step type, however the depicted peak for each corresponds to its respective average. Made in Paint.net*

response. The average peak for a normal step is 0.34 [V/ms] with a standard deviation of 0.17 [V/ms]. A hazardous step on the anterior lateral quadrant has an average peak of 0.11 [V/ms] with a standard deviation of 0.04 [V/ms]. Using these statistics, it is known that the hazardous step to normal step slope difference ratio will be 3:1 approximately 50% of the time. The significantly larger slope difference response caused by a hazardous step is the key to step classification preformed by the detection algorithm.

From a physical perspective, the slope difference is so much more effective because what it represents. The slope difference describes the difference in rate of weight distribution between the anterior lateral and anterior medial quadrants. This phenomenon is at the core of what we are trying to understand about the users step. It doesn't matter how "much" weight is on any given sensor, what matters is how fast weight is being applied to one sensor compared to another. When the user is taking a normal step, the anterior lateral and anterior medial quadrants will experience relatively similar rates of weight distribution. Thus, the slope difference will be near zero. This is illustrated in Figure 20 when comparing the slope difference of normal steps and hazardous steps. In contrast, when a user is taking a bad step with an obstacle encountering the anterior medial quadrant of the foot, there is a significant increase in the rate of weight distribution in this quadrant. Subsequently, there is a significant reduction in the rate of weight distribution in the anterior lateral quadrant. This is because while the anterior medial quadrant is first contacting the obstacle, the anterior lateral quadrant is for the most part floating in the air. The reduction of the lateral, and increase of the medial rate of weight distribution facilitates a significant spike in the slope difference between these two adjacent quadrants. This spike is significantly greater than that seen during a normal step 20, even if the step is somewhat irregular, quick in pace, or otherwise different then an ideal normal step. The slope difference response of a normal step as compared to a hazardous step is statistically significant as seen in Figure 25. Therefore the slope difference was chosen as the metric for the Detection Algorithm to use in the classification of steps.

7.6 Detection Algorithm Test Results

On May 26th, 2019, Jim Brown, a trans-tibial amputee, participated in a test trial of the TerraStep. The test procedure followed is described in detail in Section 10.1: Procedure for Plantar Regional Detection. Data was gathered showing the TerraStep system has an average hazardous step detection rate of 60.6% with the range [48% to 72%] expressing a 95% confidence interval. The statistic was gathered over 66 sample steps and is thus of a statistically significant sample size. Our current accuracy rate is lower than 80%, our desired accuracy rate.

We believe increasing the number of sensors underfoot, thereby increasing resolution, will increase the accuracy of the system. The current iteration of the sensor array only features a single force sensor in each sub-quadrant of the anterior sensing region. When the underfoot obstacles does not make direct, or near direct contact with the the force sensor in the quadrant experiencing the hazardous obstacle, the detection is often missed by the algorithm, resulting in a lower accuracy. We believe that increasing sensors underfoot will increase our system's accuracy to meet our goal of 80%.

8 Power System

The system-level constraints relevant to the power system are battery life, rechargeability, size, and weight. The system must last for 12 hours of walking, and be recharged within 12 hours. The inner dimensions of the fanny-pack housing are 210mm×100mm×50mm (L×W×H), where the battery will be affixed. The battery displacement must be be smaller than that, and less than two pounds to maintain fanny-pack comfort.

Twelve hours of battery life is a generous estimate of the maximum amount of time an average active user will walk per day. It is a worst case scenario, and therefore specified as a goal for our team. The microcontroller draws $16mA$ at $5V(80mW)$, each sensor and its driver draw $2mA$ at $5V$ (eight sensors, $80mW$), and each motor driver draws $120mA$ maximum at a voltage within the range of $6.0V - 10V$ (four motors, $1.8W$). These current consumption values were determined by measuring current from the Agilent E3631A power supply, hooked up separately to each subsystem.

The efficiency of the 5V regulator must be accounted for. The chosen regulator, SPAN02E-05, has a 78% efficiency, therefore the actual power consumption for the microcontroller and sensors from the battery is 200mW, instead of their originally measured $160mW$. This is accounted for in the values in Figure 8. This efficiency doesn't affect the primary consumers, the motors, because they are drawing power directly from the battery. The sensors and microcontroller draw only 9% of the total power, therefore their $5V$ regulator's efficiency isn't critical, and 78% is acceptable, adding only 40mW.

The most significant power consumer in the system by far are the feedback motors. We must consider their worst case mode of operation, power-wise, to size the battery capacity correctly. For this analysis, first we assume a 75% gait duty cycle[3], accounted for in the following values. The next consideration is how many motors will be driven simultaneously. There are two worst case scenarios to be analyzed here, alert feedback mode and continuous feedback mode. A hazardous step alert activates only two motors maximum, in the case of a tripping hazard at the center of a foot region. This corresponds to two motors each drawing $120mA$, or $1.3W$. However the feedback will be in continuous mode for the majority of operation, with the actual amount depending on the surface the user is walking. For this worst case power scenario, all four motors are on at medium intensity, drawing $80mA$ each or $320mA$ total, totalling $1.8W$. Clearly the continuous mode case

[3] Gait duty cycle average gathered from normal step data recorded in lab.

draws greater power, and therefore $1.8W$ was used for the motors' power budget.

The microcontroller, sensors, and motors total power consumption is $2.0W$. Twelve hours of this continuous draw is $24Wh$. This is the minimum capacity required for our battery. The two Adafruit 353's chosen have a total of $48Wh$. This excessive capacity was chosen early on while the design included up to eight Adafruit 1201 motors. A future prototype should have a smaller battery pack, however the additional capacity was helpful during testing. A summarizing breakdown of total power consumption during worst case use is seen in Figure 8:

Subsystem	Operating Conditions	Power Consumption
Sensors and associated drivers:	2.0mA each at 5V	13mW x8
Microcontroller:	16mA at 5V	100mW x1
Motors and associated drivers:	80mA each at Vbatt=7.4V	450mW x4
Total	N/A	2.0W
Capacity req'd for 12 hours of operation	N/A	24Wh

The minimum required battery energy density is $46Wh/L$ ($24Wh/525L$). This is the amount of energy needed for a twelve hour battery life and the size constraint of the fanny-pack housing's internal displacement. The battery technology with greatest energy density is lithium-ion-C, cylindrical type[1] (see Figure 26). It is for this reason that we chose the Adafruit 353, which has a $350Wh/L$ density. However, after determining the actual number of motors on the feedback band (four) and thus a more accurate, lower power budget, a different battery will be better for the application. This high density is unnecessary for the system. Also, two lithium ion batteries should not be used in series. They worked for our system because each had its own charge controller, and that prevented either of them from discharging to a level that would be dangerous. However, it is advised for a future iteration to use either a single lithium ion battery, or even another chemistry type with less density, and likely less expensive.

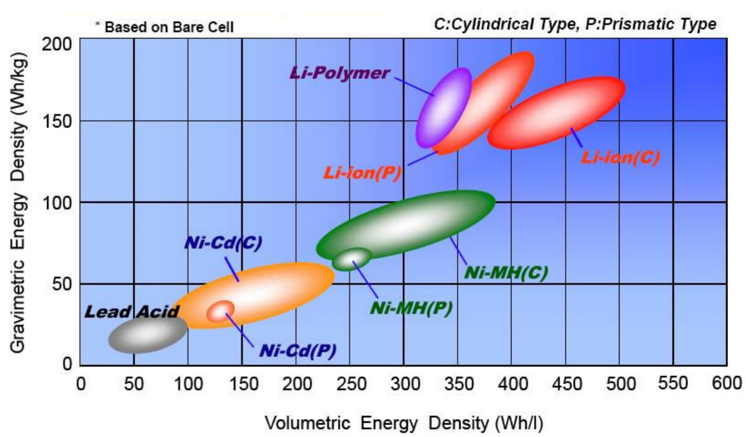

Figure 26: *Graph of various battery technologies and their energy densities and galvanic densities.*

The Adafruit 353 was ultimately chosen due to its cost, capacity, high energy density, and the availability of its corresponding charge controller, the Adafruit 259. It is the least expensive battery compared to similar lithium-ion-C batteries with its capacity ($6.6Ah$ at $3.7V$).

Two of these $3.7V$ $6.6Ah$ batteries are connected in series to provide a voltage of $6.0V$-$8.4V$ (voltage is not constant throughout a discharge cycle). That range works well to drive the motor driver circuit. It allows the motor to have $5.0 - 6.0V$ across it when commanding full intensity, with the rest of the voltage drop occuring across the MOSFET channel and 10Ω source resistor. The 1201 motors have a nominal drive voltage of $5.0V$, so further testing should be done to determine if driving at $6.0V$ will significantly affect its average lifetime. This voltage range also works for our chosen voltage regulator.

A simple and efficient way to produce and regulate $5V$ from the battery voltage is to use a switching regulator, i.e. a DC-DC converter. Converter choices considered are listed below from most to least expensive:

- Meanwell SPAN02E-05
- Recom R785.0-0.5
- Maxim Integrated MAX15062
- Texas Instruments TPS62173

The Meanwell SPAN02E-05 was chosen for it's ease of implementation (zero external components), and input voltage range. The MAX15062 and TPS62173, while inexpensive, required several other components, as well as a complex PCB layout due to their high switching frequency and small size. Although the Recom regulator was less expensive and has a smaller footprint, its input voltage is limited at $6.5V$, which conflicts with our battery's voltage range. The battery, being Lithium ion, has a charge controller that shuts off at $6.0V$. Using the Recom regulator would decrease our systems operational battery life because the $5V$ regulator would shut off before the battery is fully depleted. Therefore, the SPAN02E-05 was the best choice with the proper input voltage range, and which didn't require external components.

8.1 Battery Housing

The battery housing was influenced by a large number of factors. However with its primary function being housing and protecting the batteries and their boards, they were the first influential piece considered. The internal dimensions needed to be large enough to comfortably house all four components. Thus the internal long dimension of the housing is 210.631 mm. The boards naturally fit flush to the floor of the housing. Thus 3mm of layers below where the board is to be places is a solid fill rather than the normal 20% fill. Giving the screws something to grip onto. The next constraint is provided by the batteries themselves. The batteries are attached to the housing through the use of a strong electrical adhesive. This adhesive will ensure the batteries remain in place for the long term. In addition to the, walls also surround the battery in order to minimize possible movement to up. This increasing the lifetime of the adhesive. A power switch is also included in this portion of the housing so that during the day if the user chooses to power off their system, this is done quickly and easily. As for the external dimensions, this was influenced based upon how much protection was needed around the PCBs. We found that a wall 5 mm thick and 20% fill from the 3D printer provided the optimum ratio between rigidity and flexibility. Thus optimizing absorption during sudden impact and resistance over long term stress. Additionally, this 5 mm allows for bolts to be used to fasten the PCB housing and the battery housing together. The battery housing can be seen in Figure 27.

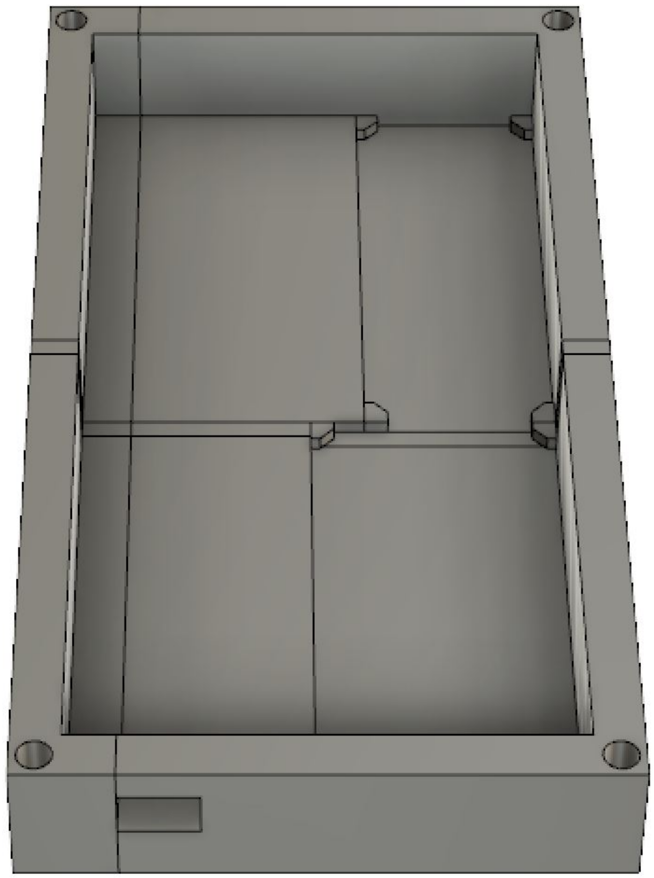

Figure 27: *Battery Side of Fanny Pack Housing. This just hold the batteries and the Adafruit 259 battery charger . It is half of the box that is screwed together that is held in the fanny pack seen in Figure 31 . CAD done in Fusion360.*

9 PCB Design

A compact, durable system is necessary because the target market includes active lower-limb amputees who will be less willing to use a cumbersome system during high-intensity activities. The PCB is mounted to a 3D-printed frame, housed inside a fanny-pack. The PCB size must be kept to a minimum to meet the fanny-pack's allotted PCB space constraint, which is 190mm x 80mm x 27mm (LxWxH). Another important constraint is the durability of the boards regarding the sensor and motor connectors. These will receive quite a bit of stress during testing, when motors and sensors are inevitably interchanged with others. This means a reliable connection must be maintained.

The circuit is split into two small boards, with dimensions shown in Figure 28, in order to decrease the size and the time needed to manufacture a functional board. Two boards allows two teammates to design and solder simultaneously, respectively. This can save a significant amount of time, considering the design and soldering of one board can take +10 hours. Other advantages include: greater flexibility in board placement, and ease of troubleshooting.

The three main subcircuits, the motor drivers, the sensor drivers, and the microcontroller, were laid out according to Appendinx Section 13.1 Figure 33 and Figure 34. The sensor drivers subcircuit was placed on the same board as the microcontroller because they have 19 connections between them, compared to 12 connections between the motor drivers subcircuit and microcontroller. This layout required a smaller ribbon cable connecting the two boards. A 14-pin connector was chosen due to time constraints and availability, but a minimum of 12 pins would also work. The $5V$ regulator, as well as the $1V$ and $2V$ reference voltage circuits were placed on the same board as the motor drivers. For brevity, these two boards will be called the motors board, and the sensors board, respectively, for the rest of this document.

In the interest of a small overall footprint, surface mount components were used where possible, including the MCP6004 op-amp chips. Each sensor and motor has its own driver circuit which are arranged into modular footprints; ideal for efficient board layout. The compact modular footprint can be seen in Figure 29, for both the sensors and the motors board. Each sensor drive circuit uses two op amps, meaning one MCP6004 chip can support up to two sensors. Each motor drive circuit uses only one op amp meaning one MCP6004 chip can support up to four motor driver circuits.

Board-mount connector durability is of utmost importance, especially during the testing and prototyping phase while connectors are frequently being inserted and removed. Once a connection fails the device will no longer operate properly. Time constraints led us to use the M60 computer controlled circuit board router to make our boards initially. However the lack of through-hole plating creates a structurally weaker connection. This is important for components that may endure extra stress, such as board-mount connectors. PCBs from a board-house are preferable for this structural advantage. They also make assembly easier with silk-screened component names.

EasyEDA is the PCB design tool we used, chosen for its cloud based accessibility, its low cost for ordering, and extensive component library. The program can be accessed from any computer connected to the internet, which is helpful when switching between lab computers and a home computer. This also makes sharing the project files much easier. The cost for five PCBs with a 24 hour turnaround is $6 plus six day shipping from China. This cost does not even compare to American board-houses, which charge much more for 24 hour turnaround. If enough time is available, American board-houses could be a good option. The component library is vast, including many manufacturers footprints, along with user-created footprints. We did not need to create any footprints for these PCBs. Trace width was calculated using the online tool from Advanced Circuits [3]. It uses formulas from IPC-2221 to calculate width from a given current. For the motors board, the signal traces were determined to be seven millimeters, and the driver circuits to be 12 millimeters for a maximum current of 120mA. With seven motors, that could get up to 840mA in some trace lengths, so the trace coming from the battery was increased to 15 millimeters. These values are greater than the calculated ones, to allow for some potential spike in current draw. The signal traces could have been less than one millimeter, according to the calculator. We opted for seven millimeters instead because it is more structurally sound and is thin enough for an easy layout.

9.1 PCB Connections

During prototyping, modularity was the biggest factor when deciding the connectors to be used by the TerraStep. The 2 circuit JST connectors (Figure 30) allow for easy removal, rearranging and color coding of our sensing and motor arrays. Their easy snap in design clips in allows for ease of removal and replacement for rapid prototyping of various sensor and actuator arrays.

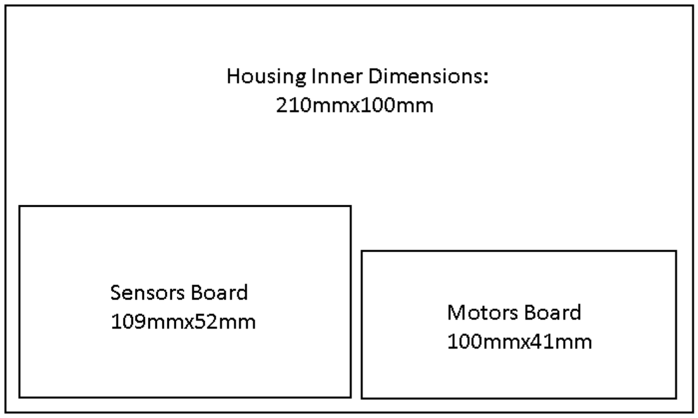

Figure 28: *Fanny-pack housing inner dimensions.*

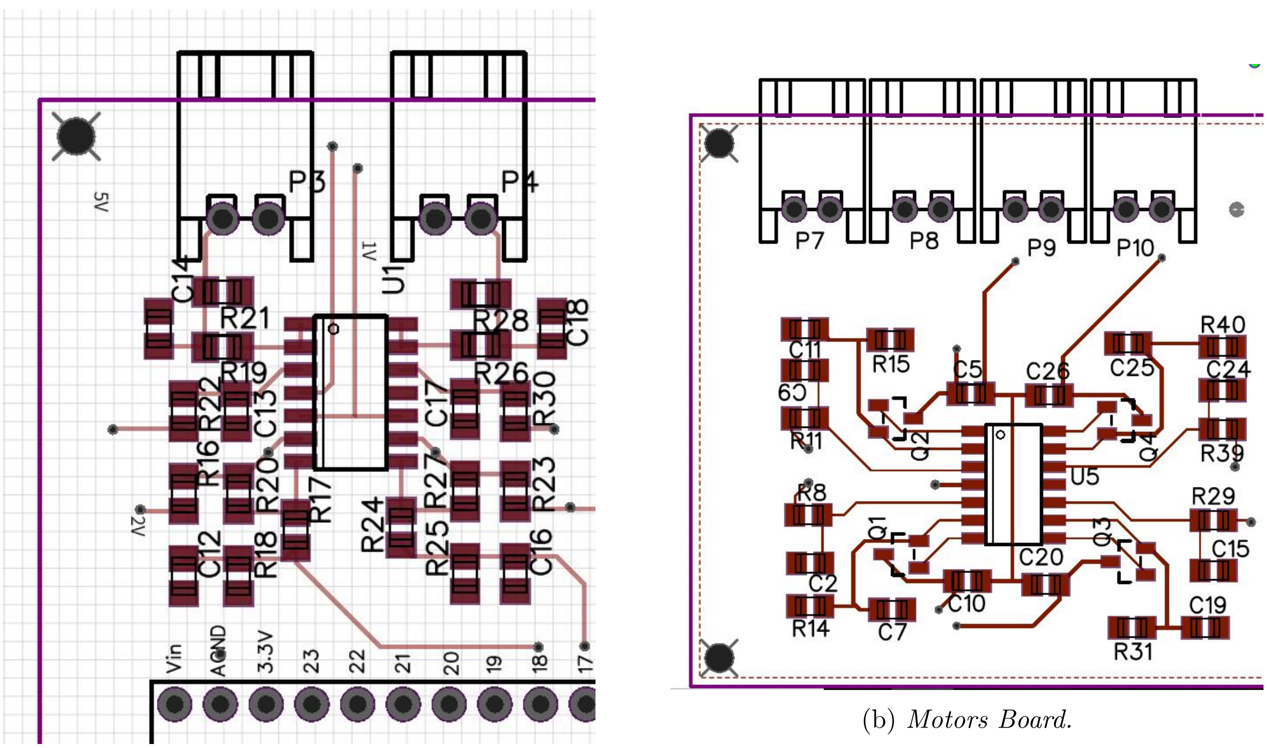

(a) *Sensors Board. CAD done in Fusion360.*

(b) *Motors Board.*

Figure 29: *Modular footprint around MCP6004 SMD. Made using EasyEDA.*

9.2 PCB Housing

The PCB housing (Figure 31 and Figure 32) was primarily influenced by housing and protection of the PCBs. The internal long dimension of the housing is 210.631 mm to accommodate both boards side by side. The boards needed to fit flush to the housing so they would be able to screw into the housing itself. This was accomplished by sinking half of the housing floor by 2mm and leaving just the portions where the board would screw into the housing raised. The bottom of this lowered floor to the top of the PCB housing is 27 mm and 25 mm for the raised portion.

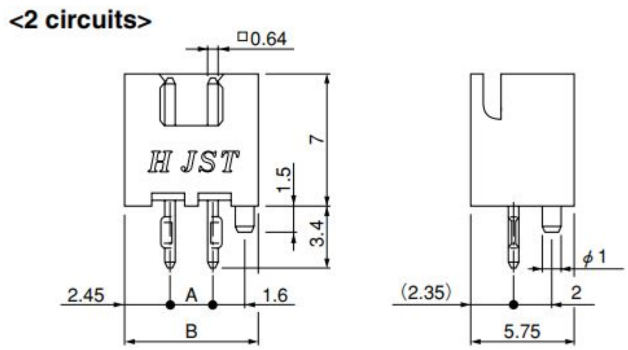

Figure 30: *Connectors Used In System Prototype.*

The next constraint provided by the PCBs was the need for the wires to be connected to the housing without interfering with one another, the functionality of the housing or the PCB itself. The ideal balance between space and footprint seemed to be 52.3 mm for the PCBs and 48.65 mm for the wires for a total of 100.95 mm. The wires are then fed through the housing in a hole just big enough to allow for the connectors to slip though. This hole is located on the short side of the housing and measures 16 mm long by 10 mm tall. As for the external dimensions, this was influenced based upon how much protection was needed around the PCBs. We found that a wall 5 mm thick and 20% fill from the 3D printer provided the optimum ratio between rigidity and flexibility. Thus optimizing absorption during sudden impact and resistance over long term stress. Additionally, this 5 mm allows for bolts to be used to fasten the PCB housing and the battery housing together.

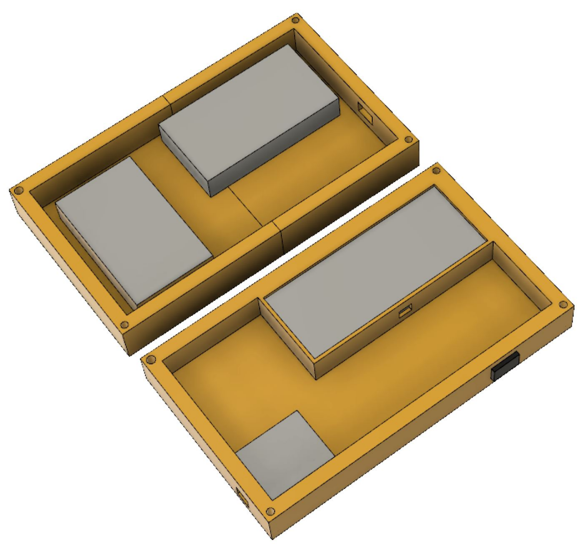

Figure 31: *Full board housing that holds the PCB boards, including the microcontroler, the batteries, and the Adafruit 259 battery charger. CAD done in Fusion360.*

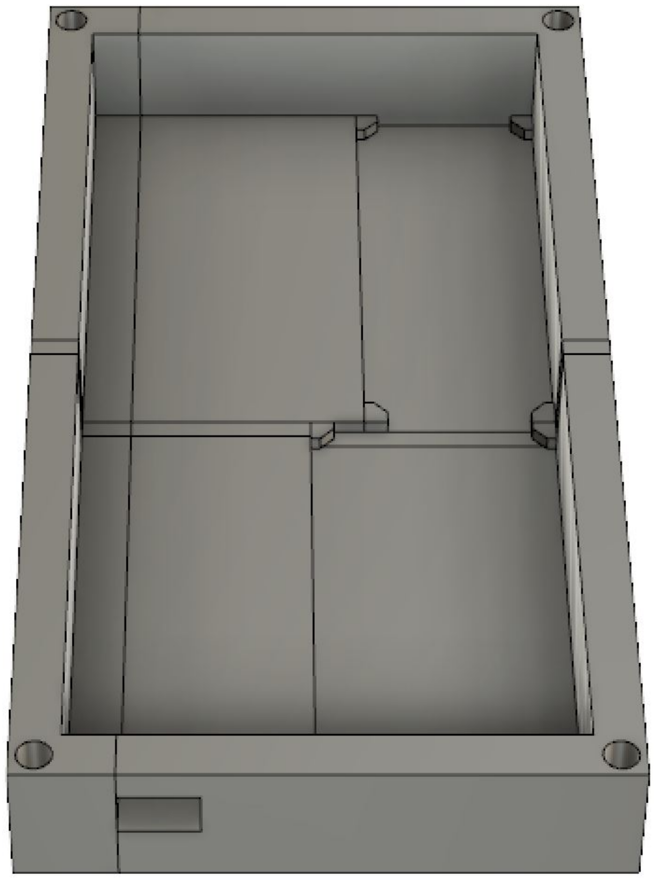

Figure 32: *Board Side of Fanny Pack Housing. This just hold the PCB boards and microcontroler. It is half of the box that is screwed together that is held in the fanny pack as seen in Figure 31. CAD done in Fusion360.*

10 System Efficacy Test

The TerraStep system provides low intensity, continuous feedback for normal steps, as well as high intensity alert feedback for hazardous object detection. The primary goal however is the alert feedback, in order to allow the user to react and prevent injury. However the subjective tactile feedback sensation and broad variation of user gait makes quantifying effectiveness difficult. We've split up the definition of effectiveness into two areas of focus, which together provide a metric of success. These are, (1) correctly distinguishing the object's location underfoot, and (2) the delay of the system from object contact to alert activation being less than 40ms. If the system can alert the user of a hazardous object's location within 40ms, it will give the user enough time to react and prevent injury.

- Plantar Regional Detection: The user needs to be able to distinguish a normal step from a hazardous step, and for the latter, needs to be able to determine where the hazardous object is (i.e. anterior or posterior, and lateral or medial side). The user will state where she feels the feedback after each object is stepped on.

- Response Time: The user needs to receive the feedback in less than 40ms in order to have time to react and adjust their balance. After the test, the user will be asked whether the alert feedback responds fast enough to help them maintain balance.

The tests are repeated once with feedback while only logging detected hazardous steps, and again without feedback while logging each sensor's raw data. Continuous data logging adversely affects the sampling frequency of the detection algorithm, therefore affecting the feedback response.

10.1 Procedure for Plantar Regional Detection

1. **Normal Steps**
   Setup: The user is outfitted with the TerraStep system, including five-sensor array, fanny pack, and feedback band. The feedback style applied is "continuous with alert", meaning low intensity continuous feedback, and high intensity alerts for objects detected. He is given time to get familiar with the new feedback system. The microcontroller is set to log hazardous steps only. Alert thresholds will need to be adjusted based on his feedback, to minimize false alerts without negatively affecting hazardous step detection sensitivity. Jim walks 165 ft hallway at a medium pace, while number of steps is recorded. He lets us know when he receives a false alert, and it is recorded (however they are being logged as well). Repeat with fast pace.

2. **Hazardous Steps**
   Setup: Jim is still wearing the TerraStep system. The test course with single tripping hazard object is placed 4 ft from a marked starting line. Two cameras are positioned to record the object, one capturing the lateral axis, the other the longitudinal axis. Use time marker to allow synchronization with other logged data. Anterior Lateral Region: Starting at the marked line, Jim takes a normal left step, then a right step onto the object, aiming for the lateral front side of his prosthetic foot. He states whether he felt the alert feedback, and if so, in which region. Repeat 30 times. Anterior Medial Region: Same as "Outer Front Region", but aim for inner front region.

3. **Raw Data**
   Still using five-sensor array, repeat steps 1 and 2 above, with feedback, and microcontroller set to record only the hazardous steps. Now with eight-sensor array, repeat steps 1 and 2 above, with feedback removed, and microcontroller set to record each sensor's raw data only.

11 Future Work

Even though much progress was made during the 2018-19 academic year, there are still areas that could be improved. These areas include posterior plantar detection, improved response time, increased sensor area, a calibration procedure, circuit simplification, and a low battery indicator.

For the TerraStep to be fully functional and effective, the detection algorithm will need to be applied to the posterior region of the foot. At the current prototype stage, the algorithm only detects objects at the anterior region. While the same algorithm will very likely apply to the posterior region, it will need some modifications. However, since the heel has a smaller area than the anterior region, fewer sensors are required. We expect fewer sensors to require a simpler detection scheme.

Detection algorithm response time in the current prototype is approximately 160ms, which is approximately five times our goal. The alert should be occurring within one tenth of the rise time of a step, which would constitute a negligible delay and could rightfully be called "immediate response". For an average step at a normal pace the rise time is 350ms. An immediate response time would therefore be 35ms. The breakdown of the delay is approximately 120ms for the detection algorithm, and 38ms for the motor to provide a significant vibrotactile intensity. Our current device is technically functional as reported by Jim Brown, our prosthetic user test subject. He was able to discern the extra delay and would prefer it to be decreased.

A significant issue we faced was providing the desired plantar regions with the proper resolution. The current prototype has small sensing areas, and objects underfoot are not detected by the algorithm as well as they should be. These areas are generally around the edges of the sensors and between them. Even with five sensors in the anterior region, there are areas not being properly measured. A potential solution to this problem could be a sensor matrix. While the detection algorithm would require significant modifications, the same foundations of monitoring slope difference will likely still be effective.

The four tuning parameters used to define the Hazardous Step Window are manually adjusted for each particular user on the current prototype, as detailed in Section **??**, and could be improved by the addition of an automatic tuner. The tuning process takes about 10 minutes to complete and currently requires a trained technician. This solution is neither user friendly nor practical when it comes to the TerraStep system as a consumer product. Future work on the software side of the project would focus heavily on implementing a calibration protocol. The calibration protocol would involve a short interval where the user performs a predefined set of steps, while the system collects data on relevant gait metrics. Once the calibration interval is completed for a particular user, the system would automatically tune the four detection algorithm parameters to allow the system to function for any user without technical intervention.

A neural network could be implemented, allowing the calibration algorithm to tune the hazardous step thresholds based on normal steps, to eliminate the need for a user to step on hazardous objects purposefully. As more data is collected, a neural network could be used to better detect hazardous steps and better calibrate the tuning parameters. This would add a valuable dimension to the product, reducing or eliminating calibration periods, and possibly detecting other important features or concerns related to user gait and posture.

The TerraStep holds much potential in markets outside of trans-tibial prosthetic users. The TerraStep could be introduced to athletes and physical therapists as a data acquisition system. The motors would be omitted, greatly reducing the power needs for the system. Ideally it would be the size of a simple shoe insole, and would have wireless capability, transmitting the force data to a phone application. More research needs to be done as to what data is relevant for each user base. This model has been used successfully by companies such a FitBit, and could be a similar endeavor for TerraStep.

A less critical area of improvement is a battery indicator. The current iteration of the Motors PCB does include a voltage divider for battery voltage monitoring and a $3.3V$ zener diode for redundant protection of the microcontroller. However, we were not able to implement this monitoring into the microcontroller due to time constraints. We did determine that there will be approximately one hour left of battery life when the voltage reaches $6.2V$. At this point, an LED on the fanny-pack would turn on, and an auditory alert would inform the user.

12 Conclusion

TerraStep allows prosthetic users to partially reclaim their planar sensory feedback, reducing tripping and injury. Continuous feedback is provided to the user for familiarity, and will facilitate more rapid understanding of feedback [4]. Alert feedback is provided to inform the user quickly of an underfoot tripping hazard. The device was designed to detect objects that are approximately 1.25 inches in diameter, and taller than 1 inch. The system was proven to detect objects of this size with about 61% accuracy.

Although our current prototype is still in the early stages of development, we gained a vast amount of knowledge about prosthetic users, human gait, discrete monitoring and classification algorithms, and troubleshooting systems. The two-style feedback scheme proved to be more effective than simply continuous feedback, as reported by our prosthetic user test subject Jim Brown.

Acknowledgements

Our team would like to thank Professor Stephen Petersen and Tela Favaloro for being the staff that took our project under their wing, and for support and advice on this project. We would also like to thank Diablo Prosthetics and Orthotics for donating the current prosthetic limbs we are using to roughly estimate the dimensions for the housing to keep the system on the leg. Pat Crowley is the practitioner that put our team in contact with Jim Brown from the beginning and encouraged us to work on this project for him. Lastly we would like to thank Jim Brown for the constant communication, willingness to prototype our system and for asking for this device to be built.

13 Appendix

13.1 Engineering Schematics

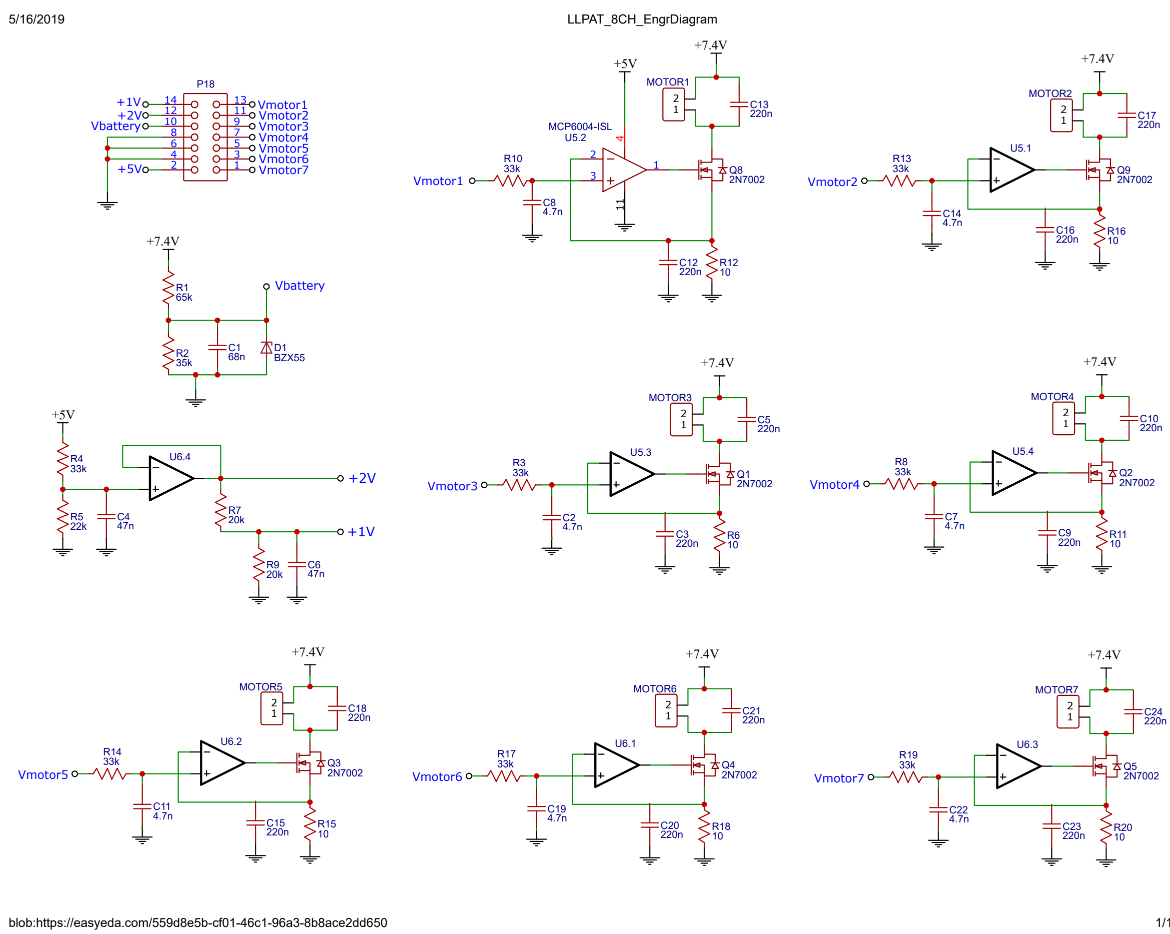

Figure 33: *PCB 1 of 2: Motors Board. This board contains all motor drive circuits and reference voltages. 5V is produced by SPAN02E-05 voltage regulator, which is powered by battery.*

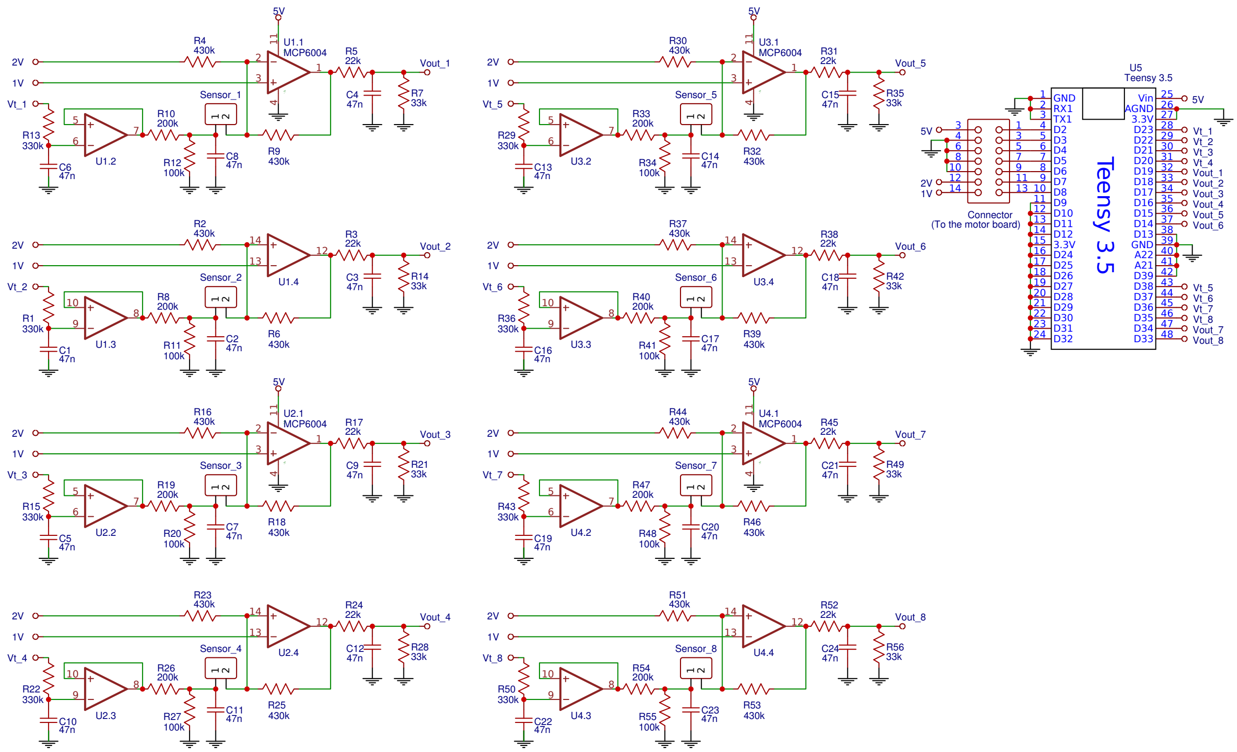

Figure 34: *PCB 2 of 2: Sensors Board. This board contains all sensor drive circuits and the microcontroller.*

13.2 Glossary

Here are some words that appear throughout the document that should be defined, as they are in this specific application.

Plantar - (adj.) Of or relating to the sole of the foot.

Medial - (adj.) Regarding plantar regions, medial refers to the inner side (toward the center of the body). Example: The front-medial region of a right foot would be the front left region.

Lateral - (adj.) Regarding plantar regions, lateral refers to the outer side (away from the center of the body). Example: The front-lateral region of a right foot would be the front right region.

Anterior - (adj.) Regarding plantar regions, anterior refers to the front side (away from the center of the body). Example: The anterior-lateral region of a right foot would be the front right region.

Posterior - (adj.) Regarding plantar regions, posterior refers to the back side (away from the center of the body). Example: The posterior-lateral region of a right foot would be the back right region.

Vibrotactile - (adj.) relating to or involving the perception of vibration through touch.

Transtibial - (adj.) Below the knee, regarding lower limb amputations.

13.3 Code

The code which implements the functionality described in the Software Implementation section can be found at: https://bitbucket.org/llpat/llpatuc/src/master/

References

[1] Keith Araujo. *Battery Cell Comparison*. 2019. URL: www.epectec.com/batteries/cell-comparison.html.

[2] Angela Boulicault. *The Truth about Hiking as an Amputee*. 2018. URL: www.boundless-journey.com/2017/01/24/hiking-as-an-amputee/.

[3] Advanced Circuits. *Trace Width Calculator*. 2018. URL: https://www.4pcb.com/trace-width-calculator.html.

[4] Lauretti. *A vibrotactile stimulation system for improving postural control and knee joint proprioception in lower limb amputees*. Universita Campus Bio-Medico di Roma, 2017.

[5] Center for Orthotic Prosthetic Care. *Amputation Statistics*. 1988-1996. URL: http://www.centeropcare.com/Portals/COPC/Amputation%5C%20Statistics.pdf.

[6] Advanced Amputee Solutions. *Amputee Statistics You Ought to Know*. 2012. URL: www.advancedamputees.com/amputee-statistics-you-ought-know.